Differentiable Discrete Event Simulation for Queuing Network Control

Ethan Che Jing Dong Hongseok Namkoong

Columbia Business School

{eche25, jing.dong, namkoong}@gsb.columbia.edu

Abstract

Queuing network control is essential for managing congestion in job-processing systems such as service systems, communication networks, and manufacturing processes. Despite growing interest in applying reinforcement learning (RL) techniques, queueing network control poses distinct challenges, including high stochasticity, large state and action spaces, and lack of stability. To tackle these challenges, we propose a scalable framework for policy optimization based on differentiable discrete event simulation. Our main insight is that by implementing a welldesigned smoothing technique for discrete event dynamics, we can compute PATHWISE policy gradients for large-scale queueing networks using auto-differentiation software (e.g., Tensorflow, PyTorch) and GPU parallelization. Through extensive empirical experiments, we observe that our policy gradient estimators are several orders of magnitude more accurate than typical REINFORCE-based estimators. In addition, We propose a new policy architecture, which drastically improves stability while maintaining the flexibility of neural-network policies. In a wide variety of scheduling and admission control tasks, we demonstrate that training control policies with pathwise gradients leads to a 50-1000x improvement in sample efficiency over state-of-theart RL methods. Unlike prior tailored approaches to queueing, our methods can flexibly handle realistic scenarios, including systems operating in non-stationary environments and those with non-exponential interarrival/service times.

1 Introduction

Queuing models are a powerful modeling tool to conduct performance analysis and optimize operational policies in diverse applications such as service systems (e.g., call centers [1], healthcare delivery systems [6], ride-sharing platforms [9], etc), computer and communication systems [45, 76], manufacturing systems [86], and financial systems (e.g., limit order books [24]). Standard tools for queuing control analysis involve establishing structural properties of the underlying Markov decision process (MDP) or leveraging analytically more tractable approximations such as fluid [28, 18] or diffusion approximations [47, 89, 48, 68]. These analytical results often give rise to simple control policies that are easy to implement and interpret. However, these policies only work under restrictive modeling assumptions and can be highly sub-optimal outside of these settings. Moreover, deriving a good policy for a given queuing network model requires substantial queuing expertise and can be theoretically challenging.

Recent advances in reinforcement learning (RL) have spurred growing interest in applying learning methodologies to solve queuing control problems, which benefit from increased data and computational resources [26, 98, 64]. These algorithms hold significant potential for generating effective controls for complex, industrial-scale networks encountered in real-world applications, which typically fall outside the scope of theoretical analysis. However, standard model-free RL algorithms [84, 82, 72] often under-perform in queuing control, even when compared to simple queuing



Figure 1. Improvements in sample efficiency of our proposed PATHWISE policy gradient estimator over a standard model-free RL estimator, REINFORCE. (Left) Samples of policy gradient estimators for a parameterized MaxPressure policy in a criss-cross network with traffic intensity $\rho = 0.9$ (see Example 3). Each draw of the REINFORCE estimator is averaged over $B = 10^3$ trajectories and is equipped with a value function baseline, which is fitted using 10^6 state transitions. The PATHWISE estimator uses only a single trajectory, and no value function. Despite using less data, it is more closely aligned with the true gradient. (Right) Average cosine similarity (higher is better) of policy gradient estimators with the true policy gradient (see (10) for more details) across different levels of traffic intensity for the criss-cross network. For REINFORCE, we plot the cosine similarity of the estimator under different batch sizes $B = 1, ..., 10^4$. We see that the efficiency advantages of PATHWISE, with only 1 trajectory, are greater under higher traffic intensities, even outperforming REINFORCE with a value function baseline and $B = 10^4$ trajectories.

policies [78, 64], unless proper modifications are made. This under-performance is primarily due to the unique challenges posed by queuing networks, including (1) high stochasticity of the trajectories, (2) large state and action spaces, and (3) lack of stability guarantees under sub-optimal policies [26]. For example, when applying policy gradient methods, typical policy gradient estimators based on coarse feedback from the environment (observed costs) suffer prohibitive error due to high variability (see, e.g., the REINFORCE estimators in Figure 1).

To tackle the challenges in applying off-the-shelf RL solutions for queuing control, we propose a new scalable framework for policy optimization that incorporates domain-specific queuing knowledge. Our main algorithmic insight is that queueing networks possess key structural properties that allow for several orders of magnitude more accurate gradient estimation. By leveraging the fact that the dynamics of discrete event simulations of queuing networks are governed by observed exogenous randomness (interarrival and service times), we propose a differentiable discrete event simulation framework. This framework enables the computation of a PATHWISE gradient of a performance objective (e.g., cumulative holding cost) with respect to actions.

Our proposed gradient estimator, denoted as the PATHWISE estimator, can then be used to efficiently optimize the parameters of a control policy through stochastic gradient descent (SGD). By utilizing the known structure of queuing network dynamics, our approach provides finer-grained feedback on the sensitivity of the performance objective to any action taken along the sample path. This offers an infinitesimal counterfactual analysis: how the performance metric would change if the scheduling action were slightly perturbed. Rather than relying on analytic provess to compute these gradients, we utilize the rapid advancements in scalable auto-differentiation libraries such as PyTorch [77] to efficiently compute gradients over a single sample path or a batch of sample



Figure 2. (Left) Pseudo-code of a single gradient step of our proposed PATHWISE estimator. Computing the estimator requires only a few lines of code to compute the cost incurred by the policy. Once this cost is calculated, the sample path gradient is computed automatically via reverse-mode auto-differentiation. Unlike standard methods such as infinitesimal perturbation analysis or likelihood-ratio estimation, we can apply the same code for any network without any bespoke modifications. Unlike model-free gradient estimators like REINFORCE, our method does not need a separate value function fitting step, managing a replay buffer, feature/return normalization, generalized advantage estimation, etc., as it has a low variance without any modification. (Right) A sample path of the total queue length (light blue) for a multi-class queuing network (see Example 2) under a randomized priority scheduling policy. Along the path, we display the gradients (dark blue) computed using our framework of the average cost with respect to each action produced by the policy: $\nabla_{u_k} \frac{1}{N} \sum_{k=0}^{N-1} c(x_k, u_k) \tau_{k+1}^*$.

paths. Our proposed approach supports very general control policies, including neural network policies, which have the potential to improve with more data and computational resources. Notably, our method seamlessly handles large-scale queuing networks and large batches of data via GPU parallelization. Unlike off-the-shelf RL solutions whose performance is exceedingly sensitive to implementation details [56, 58], our method is easy to implement (see e.g., Figure 2) and requires minimal effort for parameter tuning.

Across a range of queueing networks, we empirically observe that our PATHWISE estimator substantially improves the sample efficiency and stability of learning algorithms for queuing network control while preserving the flexibility of learning approaches. In Figure 1, we preview our main empirical findings which show that PATHWISE gradients lead to a 50-1000x improvement in sample efficiency over model-free policy gradient estimators (e.g., REINFORCE [101]). Buoyed by the promising empirical results, we provide several theoretical insights explaining the observed efficiency gains.

Our proposed approach draws inspiration from gradient estimation strategies developed in the stochastic modeling and simulation literature, particularly infinitesimal perturbation analysis (IPA) [35, 52, 60]. While IPA has been shown to provide efficient gradient estimators for specific small-scale queuing models (e.g., the G/G/1 queue), it is well-known that unbiased IPA estimates cannot be obtained for general multi-class queuing networks due to non-differentiability of the sample path [15, 32, 33]. Our framework overcomes this limitation by proposing a novel smoothing technique based on insights from fluid models/approximations for queues and tools from the machine learning (ML) literature. To the best of our knowledge, our method is the first to provide a gradient estimation framework capable of handling very general and large-scale queuing networks and various control policies. Our modeling approach is based on discrete-event simulation models, and as a result, it can accommodate non-stationary and non-Markovian inter-arrival and service times, requiring only samples instead of knowledge of the underlying distributions.

Our second contribution is a simple vet powerful modification to the control policy architecture. It has been widely observed that training a standard RL algorithm, such as proximal policy optimization [84] (PPO), may fail to converge due to instabilities arising from training with random initialization. To address this issue, researchers have proposed either switching to a stabilizing policy when instability occurs [64] or imitating (behavior cloning) a stabilizing policy at the beginning [26]. However, both methods limit policy flexibility and introduce additional complexity in the training process. We identify a key source of the problem: generic policy parameterizations (e.g., neural network policies) do not enforce work conservation, leading to scenarios where even optimized policies often assign servers to empty queues. To address this, we propose a modification to standard policy parameterizations in deep reinforcement learning, which we refer to as the 'work-conserving softmax'. This modification is compatible with standard reinforcement learning algorithms and automatically guarantees work conservation. Although work conservation does not always guarantee stability, we empirically observe across many scenarios that it effectively eliminates instability in the training process, even when starting from a randomly initialized neural network policy. This modification not only complements our gradient estimator but is also compatible with other model-free RL approaches. We find that while PPO without any modifications fails to stabilize large queuing networks and leads to runaway queue lengths, PPO with the work-conserving softmax remains stable from random initialization and can learn better scheduling policies than traditional queuing policies.

Since rigorous empirical validation forms the basis of algorithmic progress, we provide a thorough empirical validation of the effectiveness of the differentiable discrete event simulator for queuing network control. We construct a wide variety of benchmark control problems, ranging from learning the $c\mu$ -rule in a simple multi-class queue to scheduling and admission control in large-scale networks. Across the board, we find that our proposed PATHWISE gradient estimator achieves significant improvements in sample efficiency over model-free alternatives, which translate to downstream improvements in optimization performance.

- In a careful empirical study across 10,800 parameter settings, we find that for **94.5**% of these settings our proposed PATHWISE gradient estimator computed along a single sample path achieves greater estimation quality than REINFORCE with **1000x** more data (see section 5.1).
- In a scheduling task in multi-class queues, gradient descent with PATHWISE gradient estimator better approximates the optimal policy (the $c\mu$ -rule) and achieves a smaller average cost than REINFORCE with a value function baseline and **1000x** more data (see section 5.2).
- In an admission control task, optimizing the buffer sizes with PATHWISE gradient estimator achieves smaller costs than randomized finite differences (SPSA [88]) with 1000x more data, particularly for higher-dimensional problem instances (see section 5.3).
- For large-scale scheduling problems, policy gradient with PATHWISE gradient estimator and work-conserving softmax policy architecture achieves a smaller long-run average holding cost than traditional queuing policies and state-of-the-art RL methods such as PPO, which use **50x** more data (see section 7). Performance gains are greater for larger networks with non-exponential noise.

These order-of-magnitude improvements in sample efficiency translate to improved computational efficiency when drawing trajectories from a simulator and improved data efficiency if samples of event times are collected from a real-world system.

Overall, these results indicate that one can achieve significant improvements in sample efficiency by incorporating the specific structure of queuing networks, which is under-utilized by model-free reinforcement learning methods. In section 8, we investigate the M/M/1 queue as a theoretical case study and show that even with an optimal baseline, REINFORCE has a sub-optimally large variance under heavy traffic compared to a pathwise policy gradient estimator. This analysis identifies some of the statistical limitations of REINFORCE, and illustrates that a better understanding of the transition dynamics, rather than narrowly estimating the value-function or Q-function, can deliver large improvements in statistical efficiency. Given the scarcity of theoretical results comparing the statistical efficiency of different policy gradient estimators, this result may be of broader interest.

Our broad aim with this work is to illustrate a new paradigm for combining the deep, structural knowledge of queuing networks developed in the stochastic modeling literature with learning and data-driven approaches. Rather than either choosing traditional queuing policies, which can be effective for certain queueing control problems but do not improve with data, or choosing model-free reinforcement learning methods, which learn from data but do not leverage known structure, our framework offers a favorable midpoint: we leverage structural insights to extract much more informative feedback from the environment, which can nonetheless be used to optimize black-box policies and improve reliability. Beyond queuing networks, our algorithmic insight provides a general-purpose tool for computing gradients in general discrete-event dynamical systems. Considering the widespread use of discrete-event simulators with popular modeling tools such as AnyLogic [94] or Simio [87] and open-source alternatives such as SimPy [69], the tools developed in this work can potentially be applied to policy optimization problems in broader industrial contexts.

The organization of this paper is as follows. In section 2, we discuss connections with related work. In section 3, we introduce the discrete-event dynamical system model for queuing networks. In section 4, we introduce our framework for gradient estimation. In section 5, we perform a careful empirical study of our proposed gradient estimator, across estimation and optimization tasks. In section 6, we discuss the instability issue in queuing control problems and our proposed modification to the policy architecture to address this. In section 7, we empirically investigate the performance of our proposed pathwise gradient estimation and work-conserving policy architecture in optimizing scheduling policies for large-scale networks. In section 8, we discuss the M/M/1queue as a theoretical case study concerning the statistical efficiency of REINFORCE compared to PATHWISE estimators. Finally, section 9 concludes the paper and discusses extensions.

2 Related Work

We discuss connections to related work in queuing theory, reinforcement learning, and gradient estimation in machine learning and operations research.

Scheduling in Queuing Networks Scheduling is a long-studied control task in the queuing literature for managing queues with multiple classes of jobs [48, 70]. Standard policies developed in the literature include static priority policies such as the $c\mu$ -rule [25], threshold policies [80], policies derived from fluid approximations [7, 20, 71], including discrete review policies [46, 67], policies that have good stability properties such as MaxWeight [89] and MaxPressure [27]. Many of these

policies satisfy desirable properties such as throughput optimality [93, 5], or cost minimization [25, 68] for certain networks and/or in certain asymptotic regimes. In our work, we aim to leverage some of the theoretical insights developed in this literature to design reinforcement learning algorithms that can learn faster and with less data than model-free RL alternatives. We also use some of the standard policies as benchmark policies when validating the performance of our PATHWISE policy gradient algorithm.

Reinforcement Learning in Queueing Network Control Our research connects with the literature on developing reinforcement learning algorithms for queuing network control problems [73. 85, 79, 26, 64, 100, 78]. These works apply standard model-free RL techniques (e.g. Q-learning, PPO, value iteration, etc.) but introduce novel modifications to address the unique challenges in queuing network control problems. Our work differs in that we propose an entirely new methodology for learning from the environment based on differentiable discrete event simulation, which is distinct from all model-free RL methods. The works [85, 64, 26, 78] observe that RL algorithms tend to be unstable and propose fixes to address this, such as introducing a Lyapunov function into the rewards, or behavior cloning of a stable policy for initialization. In our work, we propose a simple modification to the policy network architecture, denoted as the work-conserving softmax as it is designed to ensure work-conservation. We find empirically that work-conserving softmax ensures stability with even randomly initialized neural network policies. In our empirical experiments, we primarily compare our methodology with the PPO algorithm developed in [26]. In particular, we construct a PPO baseline with the same hyper-parameters, neural network architecture, and variance reduction techniques as in [26], although with our policy architecture modification that improves stability.

Differentiable Simulation in RL and Operations Research While differentiable simulation is a well-studied paradigm for control problems in physics and robotics [50, 55, 90, 53, 81], it has only recently been explored for large-scale operations research problems. For instance, [66, 2] study inventory control problems and train a neural network using direct back-propagation of the cost, as sample paths of the inventory levels are continuous and differentiable in the actions. In our work, we study control problems for queuing networks, which are discrete and non-differentiable, preventing the direct application of such methods. To address this, we develop a novel framework for computing pathwise derivatives for these non-differentiable systems, which proves highly effective for training control policies. Another line of work, including [3, 4], proposes differentiable agent-based simulators based on differentiable relaxations. While these relaxations have shown strong performance in optimization tasks, they also introduce unpredictable discrepancies with the original dynamics. We introduce tailored differentiable relaxations in the back-propagation process only, ensuring that the forward simulation remains true to the original dynamics.

Gradient Estimation in Machine Learning Gradient estimation [74] is an important sub-field of the machine learning literature, with applications in probabilistic modeling [61, 59] and reinforcement learning [101, 92]. There are two standard strategies for computing stochastic gradients [74]. The first is the score-function estimator or REINFORCE [101, 92], which only requires the ability to compute the gradient of log-likelihood but can have high variance [44]. Another strategy is the reparameterization trick [61], which involves decomposing the random variable into the stochasticity and the parameter of interest, and then taking a pathwise derivative under the realization of the stochasticity. Gradient estimators based on the reparameterization trick can have much smaller variance [74], but can only be applied in special cases (e.g. Gaussian random variables) that enable this decomposition. Our methodology makes a novel observation that for queuing networks, the structure of discrete-event dynamical systems gives rise to the reparameterization trick. Nevertheless, the function of interest is non-differentiable, so standard methods cannot be applied. As a result, our framework also connects with the literature on gradient estimation for discrete random variables [59, 65, 11, 96]. In particular, to properly smooth the non-differentiability of the event selection mechanism, we employ the straight-through trick [11], which has been previously used in applications such as discrete representation learning [97]. Our work involves a novel application of this technique for discrete-event systems, and we find that this is crucial for reducing bias when smoothing over long time horizons.

Gradient Estimation in Operations Research There is extensive literature on gradient estimation for stochastic systems [35, 36, 39, 14, 33], some with direct application to queuing optimization [63, 42, 33]. Infinitesimal Perturbation Analysis (IPA) [35, 52, 60] is a standard framework for constructing pathwise gradient estimators, which takes derivatives through stochastic recursions that represent the dynamics of the system. While IPA has been applied successfully to some specific queuing networks and discrete-event environments more broadly [91], standard IPA techniques cannot be applied to general queuing networks control problems, as has been observed in [15]. There has been much research on outlining sufficient conditions under which IPA is valid, such as the commuting condition in [35, 36] or the perturbation conditions in [14], but these conditions do not hold in general. Several extensions to IPA have been proposed, but these alternatives require knowing the exact characteristics of the sampling distributions and bespoke analysis of event paths [33, 32]. Generalized likelihood-ratio estimation [39] is another popular gradient estimation framework, which leverages an explicit Markovian formulation of state transitions to estimate parameter sensitivities. However, this requires knowledge of the distributions of stochastic inputs, and even with this knowledge, it may be difficult to characterize the exact Markov transition kernel of the system. Finally, finite differences [31] and finite perturbation analysis [52, 15] are powerful methods, particularly when aided with common random numbers [41, 38], as it requires minimal knowledge about the system. However, it has been observed that performance can scale poorly with problem dimension [37, 41], and we also observe this in an admission control task (see Section 5.3).

Our contribution is proposing a novel, general-purpose framework for computing pathwise gradients through careful smoothing, which only requires samples of random input (e.g., interarrival times and service times) rather than knowledge of their distributions. Given the negative results about the applicability of IPA for general queuing network control problems (e.g., general queuing network model and scheduling policies), we introduce bias through smoothing to achieve generality. It has been observed in [29] that biased IPA surrogates can be surprisingly effective in simulation optimization tasks such as ambulance base location selection. Our extensive empirical results confirm this observation and illustrate that while there is some bias, it is very small in practice, even over long time horizons (> 10^5 steps).

3 Discrete-Event Dynamical System Model for Queuing Networks

We describe multi-class queuing networks as discrete-event dynamical systems. This is different from the standard Markov chain representation, which is only applicable when inter-arrival and service times are exponentially distributed. To accommodate more general event-time distributions, the system description not only involves the queue lengths, but also auxiliary information such as residual inter-arrival times and workloads. Surprisingly, this more detailed system description leads to a novel gradient estimation strategy (discussed in Section 4) for policy optimization.

We first provide a brief overview of the basic scheduling problem. We then describe the discreteevent dynamics of multi-class queuing networks in detail and illustrate with a couple of wellknown examples. While queuing networks have been treated as members of a more general class of Generalized Semi-Markov Processes (GSMPs) that reflect the discrete-event structure of these systems [40], we introduce a new set of notations tailored for queuing networks to elaborate on some of their special structures. In particular, we represent the discrete event dynamics via matrix-vector notation that maps directly to its implementation in auto-differentiation frameworks, allowing for the differentiable simulation of large-scale queueing networks through GPU parallelization.

3.1 The Scheduling Problem

A multi-class queuing network consists of n queues and m servers. The core state variable is the queue lengths associated with each queue, denoted as $x(t) \in \mathbb{N}^n_+$, which evolves over continuous time. As a discrete-event dynamical system, the state also includes auxiliary data denoted as z(t)— consisting of residual inter-arrival times and workloads at time t—which determines state transitions but are typically not visible to the controller.

The goal of the controller is to route jobs to servers, represented by an assignment matrix $u \in \{0,1\}^{m \times n}$, to manage congestion. More concretely, the problem is to derive a policy $\pi(x)$, which only depends on the observed queue lengths and selects scheduling actions, to minimize the integral of some instantaneous costs c(x, u). A typical instantaneous cost is a linear holding/waiting cost:

$$c(x,u) = h^{\top}x$$

for some vector $h \in \mathbb{R}^n_+$. The objective is to find a policy π that minimizes the cumulative cost over a time horizon:

$$\min_{\pi} \mathbb{E}\left[\int_0^T c(x(t), \pi(x(t)))dt\right].$$
(1)

Optimizing a continuous time objective can be difficult and may require an expensive discretization procedure. However, discrete-event dynamical systems are more structured in that x(t) is piecewise constant and is only updated when an **event** occurs. For the multi-class queuing networks, events are either arrivals to the network or job completions, i.e., a server finishes processing a job.

It is then sufficient to sample the system only when an event occurs, and we can approximate the continuous-time objective with a performance objective in the discrete-event system over N events,

$$\min_{\pi} \left\{ J_N(\pi) := \mathbb{E} \left[\sum_{k=0}^{N-1} c(x_k, \pi(x_k)) \tau_{k+1}^* \right] \right\}$$

$$\tag{2}$$

where x_k is the queue lengths after the kth event update. τ_{k+1}^* is an inter-event time that measures the time between the kth and (k+1)th event, and N is chosen such that the time of the Nth event, a random variable denoted as t_N , is "close" to T.

The dynamics of queuing networks are highly stochastic, with large variations across trajectories. Randomness in the system is driven by the random arrival times of jobs and the random workloads (service requirements) of these jobs. We let $\xi_{1:N} = {\{\xi_i\}_{i=1}^N}$ denote a single realization, or 'trace', of these random variables over the horizon of N events. We can then view the expected cost (2) more explicitly as a policy cost averaged over traces. In addition, we focus on a parameterized family of policies ${\pi_{\theta} : \theta \in \Theta}$, for some $\Theta \subseteq \mathbb{R}^d$, in order to optimize (2) efficiently. In this case, we utilize the following shorthand $J_N(\theta; \xi_{1:N})$ for the policy cost over a single trace and $J_N(\theta)$ for the average policy cost under π_{θ} , which leads to the parameterized control problem:

$$\min_{\theta} \left\{ J_N(\theta) := \mathbb{E}\left[J(\theta; \xi_{1:N}) \right] := \mathbb{E}\left[\sum_{k=0}^{N-1} c(x_k, \pi_\theta(x_k)) \tau_{k+1}^* \right] \right\}.$$
(3)

We now turn to describe the structure of the transition dynamics of multi-class queuing networks, to elaborate how scheduling actions affect the queue lengths.

3.2 System Description

Recall that the multi-class queuing network consists of n queues and m servers, where each queue is associated with a job class, and different servers can be of different compatibilities with various job classes. Recall that $x(t) \in \mathbb{N}^n_+$ denotes the lengths of the queues at time $t \in \mathbb{R}_+$. The queue lengths x(t) are updated by one of two types of events: job arrivals or job completions. Although the process evolves in continuous time, it is sufficient to track the system only when an event occurs. We let $k \in \mathbb{N}_+$ count the kth event in the system, and let t_k denote the time immediately after the kth event occurs. By doing so, we arrive at a discrete-time representation of the system. Given that we do not assume event times are exponential, the queue lengths x_k alone are not a Markovian descriptor of the system. Instead, we must consider an **augmented** state $s_k = (x_k, z_k)$, where $x_k \in \mathbb{N}^n_+$ is the vector of **queue lengths** and $z_k = (\tau^A_k, w_k), \in \mathbb{R}^{2n}_+$ is an auxiliary state vector that includes **residual inter-arrival times** $\tau^A_k = {\tau^A_{k,j}}_{j=1}^n \in \mathbb{R}^n_+$ and **residual workloads** $w_k = {w_{k,j}}_{j=1}^n \in \mathbb{R}^n_+$ of the 'top-of-queue' jobs in each queue. The auxiliary state variables determine the sequence of events.

More explicitly, for each queue $j \in [n]$, the residual inter-arrival time $\tau_{k,j}^A$ keeps track of the time remaining until the next arrival to queue j occurs. Immediately after an arrival to queue j occurs, the next inter-arrival time is drawn from a probability distribution F_j^A . When a job arrives to queue j, it comes with a workload (service requirement) drawn from a distribution F_j^S . We allow the distributions F_j^A 's and F_j^S 's to vary with time, i.e., the interarrival times and service requirements can be time-varying. For notational simplicity, we will not explicitly denote the time dependence here. We refer to the residual workload at time t_k of the top-of-queue job in queue j as $w_{k,j}$, which specifies how much work must be done before the job completion. A job is only processed if it is routed to a server $i \in [m]$, in which case the server processes the job at a constant service rate $\mu_{ij} \in \mathbb{R}_+$. We refer to $\mu \in \mathbb{R}^{m \times n}_+$ as the matrix of service rates. Under this scheduling decision, the **residual processing time**, i.e., the amount of time required to process the job, is $\tau_{k,j}^S = w_{k,j}/\mu_{ij}$.

The augmented state s_k is a valid Markovian descriptor of the system and we now describe the corresponding transition function f such that

$$s_{k+1} = f(s_k, u_k, \xi_{k+1}),$$

where u_k is an action taken by the controller and ξ_{k+1} contains external randomness arising from new inter-arrival times or workloads drawn from F_j^{A} 's or F_j^{S} 's depending on the event type.



Figure 3. One step of the dynamics for the criss-cross network (see Example 3 and Figure 13). There are 3 queues and 2 servers. Beginning with queue-lengths $x_k = (3, 1, 4)$ and workloads w_k , the action u_k assigns server 1 to queue 3 and server 2 to queue 2. The workloads of the selected queues are highlighted in light green. As a result, the valid events are arrivals to queue 1 and queue 3 (queue 2 has no external arrivals) and job completions for queue 2 and queue 3 (queue 1 cannot experience any job completions because no server is assigned). The arrival event to queue 1 has the minimum residual time (highlighted in red) so it is the next event, and e_{k+1} is a one-hot vector indicating this. Since an arrival occurred, the queue-lengths are updated as $x_{k+1} = (4, 1, 4)$.

The transition is based on the next event, which is the event with the minimum residual time. The controller influences the transitions through the processing times, by deciding which jobs get routed to which servers. We focus on scheduling problems where the space of controls \mathcal{U} are feasible assignments of servers to queues. Let $\mathbf{1}_n \in \mathbb{R}^n$ denote an *n*-dimensional vector consisting of all ones. The action space is,

$$\mathcal{U} := \left\{ u \in \{0, 1\}^{m \times n} : u \mathbf{1}_n = \mathbf{1}_m, \mathbf{1}_m^\top u = \mathbf{1}_n, u \le M \right\},\tag{4}$$

where $M \in \{0, 1\}^{m \times n}$ is the **topology** of the network, which indicates which job class can be served by which server. Following existing works on scheduling in queuing networks [70], we consider networks for which each job class has exactly 1 compatible server.

Assumption A. For every queue j, there is 1 compatible server, i.e., $\sum_{i=1}^{m} M_{ij} = 1$.

Given an action u, the residual processing time is $w_{k,j}/\mu_{ij}$ when $u_{ij} = 1$ and ∞ when $u_{ij} = 0$. This can be written compactly as

$$\tau_{k,j}^S \equiv \frac{w_{k,j}}{\sum_{i=1}^m u_{ij}\mu_{ij}} = \frac{w_{k,j}}{\operatorname{diag}(u^\top \mu)_{j,j}},\tag{5}$$

where $\operatorname{diag}(u^{\top}\mu) \in \mathbb{R}^{n \times n}$ extracts the diagonal entries of the matrix $u^{\top}\mu \in \mathbb{R}^{n \times n}$.

As a result, at time t_k the **residual event times** $\tau_k \in \mathbb{R}^{2n}_+$ consists of the residual inter-arrival and processing times,

$$au_k \equiv (au_k^A, au_k^S) = (au_k^A, \mathbf{diag}(u_k^\top \mu)^{-1} w_k)$$

We emphasize that τ_k depends on the action u. The core operation in the transition dynamics is the **event selection** mechanism. The next event is the one with the minimum residual time in τ_k . We define $e_{k+1} \in \{0,1\}^{2n}$ to be a one-hot vector representing the argmin of τ_k – the position of the minimum in τ_k :

$$e_{k+1}(z_k, u_k) \equiv \operatorname{argmin}(\tau_k) \in \{0, 1\}^{2n}$$
 (Event Select)

 $e_{k+1}(z_k, u_k)$ indicates the type of the (k+1)th event. In particular, if the minimum residual event time is a residual inter-arrival time, then the next event is an arrival to the system. If it is a residual job processing time, then the next event is a job completion. We denote τ_{k+1}^* to be the **inter-event time**, which is equal to the minimum residual time:

$$\tau_{k+1}^*(z_k, u_k) = \min\{\tau_k\}$$
 (Event Time)

 $\tau_{k+1}^*(z_k, u_k)$ is the time between the kth and (k+1)th event, i.e. $t_{k+1} - t_k$.

After the job is processed by a server, it either leaves the system or proceeds to another queue. Let $R \in \mathbb{R}^{n \times n}$ denote the **routing matrix**, where the jth column, R_j details the change in the queue lengths when a job in class j finishes service. For example, for a tandem queue with two queues, the routing matrix is

$$R = \begin{bmatrix} -1 & 0 \\ 1 & -1 \end{bmatrix}$$

indicating that when a job in the first queue completes service, it leaves its own queue and joins the second queue. When a job in the second queue completes service, it leaves the system.

We define the **event matrix** D as a block matrix of the form

$$D = \begin{bmatrix} I_n & R \end{bmatrix}$$

where I_n is the $n \times n$ identity matrix. The event matrix determines the update to the queue lengths, depending on which event took place. In particular, when the (k + 1)th event occurs, the update to the queue lengths is

$$x_{k+1} = x_k + De_{k+1}(z_k, u_k)$$
 (Queue Update)

Intuitively, the queue length of queue j increases by 1 when the next event is a class j job arrival; the queue lengths update according to R_j when the next event is a queue j job completion.

The updates to the auxiliary state $z_k = (\tau_k^A, w_k) \in \mathbb{R}^{2n}_+$ is typically given by

$$\begin{bmatrix} \tau_{k+1}^{A} \\ w_{k+1} \end{bmatrix} = \begin{bmatrix} \tau_{k}^{A} \\ w_{k} \end{bmatrix} - \underbrace{\tau_{k+1}^{*} \begin{bmatrix} \mathbf{1}_{n} \\ \mathbf{diag}(u_{k}^{\top}\mu) \end{bmatrix}}_{\text{reduce residual times}} + \underbrace{\begin{bmatrix} T_{k+1} \\ W_{k+1} \end{bmatrix} \odot e_{k+1}}_{\text{draw new times / workloads}} \quad (\text{Aux Update})$$

where \odot is the element-wise product and $T_{k+1} = \{T_{(k+1),j}\}_{j=1}^n \in \mathbb{R}^n$ are new inter-arrival times $T_{(k+1),j} \sim F_j^A$ and $W_{k+1} = \{W_{k+1,j}\}_{j=1}^n \in \mathbb{R}^n$ are workloads $W_{k+1,j} \sim F_j^S$. Intuitively, after an event occurs, we reduce the residual inter-arrival times by the inter-event time. We reduce workloads by the amount of work applied to the job, i.e., the inter-event time multiplied by the service rate of the allocated server. Finally, if an arrival occurred we draw a new inter-arrival time; if a job was completed, we draw a new workload for the top-of-queue job (if the queue is non-empty).

There are two boundary cases that make the update slightly different from (Aux Update). First, if a new job arrives at an empty queue j (either an external arrival or a transition from a job completion), we also need to update $w_{k+1,j}$ to $W_{k+1,j}$. Second, if a queue j job completion leaves an empty queue behind, we set $w_{k+1,j} = \infty$, indicating that no completions can occur for an empty queue.



Figure 4: M/M/1 queue.

Let ξ_k denote exogenous noise in the environment, which consists of the sampled inter-arrival times and workloads for resetting the time of a completed event,

$$\xi_k = (T_k, W_k) \in \mathbb{R}^{2n}_+.$$

We finally arrive at the stated goal of describing the transition dynamics of $s_k = (x_k, z_k)$ in terms of a function $f(s_k, u_k, \xi_{k+1})$. Notably, all the stochasticity is captured by ξ_k 's, which are independent of the states and actions.

It is worth mentioning a few features of this discrete-event representation.

- While auxiliary data $z_k = (\tau_k^A, w_k)$ is necessary for $s_k = (x_k, z_k)$ to be a valid Markovian system descriptor, this information is typically not available to the controller. We assume the controller only observes the queue lengths, i.e., the control policy π only depends on x_k .
- The representation can flexibly accommodate non-stationary and non-exponential event-time distributions, i.e., F_i^{A} 's and F_i^{S} 's can be general and time-varying.
- This model enables purely data-driven simulation, as it only requires samples of the event times ξ_k . One does not need to know the event time distributions F_j^A 's and F_j^S 's to simulate the system if data of these event times are available.
- The matrix-vector representation enables GPU parallelism, which can greatly speed up the simulation of large-scale networks.
- As we will explain later, this representation enables new gradient estimation strategies.

Queuing Network Examples As a concrete illustration, we show how a few well-known queuing networks are described as discrete-event dynamical systems.

Example 1: The M/M/1 queue (see Figure 4) with arrival rate $\lambda > 0$ and service rate $\mu \ge \lambda$ features a single queue n = 1 and a single server m = 1, and exponentially distributed inter-arrival times and workloads, i.e., $T_k \sim \mathsf{Exp}(\lambda)$ and $W_k \sim \mathsf{Exp}(1)$ respectively. The network topology is M = [1], the service rate is μ , and the routing matrix is R = [-1], indicating that jobs leave the system after service completion. The scheduling policy is work-conserving, the server always serves the queue when it is non-empty, i.e. $u_k = 1\{x_k > 0\}$. The state update is,

$$\begin{aligned} x_{k+1} &= x_k + \begin{bmatrix} 1 & -1 \end{bmatrix}^\top e_{k+1} \\ e_{k+1} &= \arg\min\left\{\tau_k^A, \tau_k^S\right\} = \arg\min\left\{\tau_k^A, \frac{w_k}{\mu \cdot 1\{x_k > 0\}}\right\} \in \{0, 1\}^2 \\ \tau_{k+1}^* &= \min\left\{\tau_k^A, \tau_k^S\right\} \\ \frac{\tau_{k+1}^A}{w_{k+1}} &= \begin{bmatrix} \tau_k^A \\ w_k \end{bmatrix} - \tau_{k+1}^* \begin{bmatrix} 1 \\ \mu 1\{x_k > 0\} \end{bmatrix} + \begin{bmatrix} T_{k+1} \\ W_{k+1} \end{bmatrix} \odot e_{k+1}. \end{aligned}$$

 \diamond

Example 2: The **multi-class singer-server queue** features an n queues and a single server m = 1 (see Figure 5). While the inter-arrival times and workloads, i.e., (T_k, W_k) 's are usually exponentially distributed, they can also follow other distributions. The network topology is $M = [1, ..., 1] \in \mathbb{R}^n$, the service rates are $\mu = [\mu_1, ..., \mu_n]$, and the routing matrix is $R = [-1, ..., -1] \in \mathbb{R}^n$, indicating that jobs leave the system after service completion. A well-known scheduling policy for

this system is the $c\mu$ -rule, a static priority rule. Let $h = (h_1, ..., h_n) \in \mathbb{R}^n$ denote the holding costs. The $c\mu$ -rule sets

$$u_k = \operatorname*{argmax}_{j \in [n]} \{h_j \mu_j 1\{x_j > 0\}\} \in \{0, 1\}^n.$$

The state update is,

 \diamond

$$\begin{aligned} x_{k+1} &= x_k + \begin{bmatrix} \mathbf{1}_n & -\mathbf{1}_n \end{bmatrix}^\top e_{k+1} \\ e_{k+1} &= \arg\min\left\{\tau_{k,1}^A, ..., \tau_{k,n}^A, \tau_{k,1}^S, ..., \tau_{k,n}^S\right\} = \arg\min\left\{\tau_{k,1}^A, ..., \tau_{k,n}^A, \frac{w_{k,1}}{\mu_1 u_{k,1}}, ..., \frac{w_{k,n}}{\mu_n u_{k,n}}\right\} \\ \tau_{k+1}^* &= \min\left\{\tau_{k,1}^A, ..., \tau_{k,n}^A, \tau_{k,1}^S, ..., \tau_{k,n}^S\right\} \\ \begin{bmatrix} \tau_{k+1}^A \\ w_{k+1} \end{bmatrix} &= \begin{bmatrix} \tau_k^A \\ w_k \end{bmatrix} - \tau_{k+1}^* \begin{bmatrix} \mathbf{1}_n \\ \mu \odot u_k \end{bmatrix} + \begin{bmatrix} T_{k+1} \\ W_{k+1} \end{bmatrix} \odot e_{k+1}. \end{aligned}$$

Example 3: The criss-cross network [49] features n = 3 queues and m = 2 servers (see Figure 13). External jobs arrive to queues 1 and 3. The first server can serve queues 1 and 3 with service rates μ_{11} and μ_{13} respectively, while the second server is dedicated to serving queue 2 with service rate μ_{22} . After jobs from queue 1 are processed, they are routed to queue 2; jobs from queues 2 and 3 exit the system after service completion. The inter-arrival times and workloads, i.e., (T_k, W_k) 's, can follow general distributions. The network topology M, service rate matrix, and the routing matrix R are:

$$M = \begin{bmatrix} 1 & 0 & 1 \\ 0 & 1 & 0 \end{bmatrix}, \qquad \mu = \begin{bmatrix} \mu_{11} & 0 & \mu_{13} \\ 0 & \mu_{22} & 0 \end{bmatrix}, \qquad R = \begin{bmatrix} -1 & 0 & 0 \\ 1 & -1 & 0 \\ 0 & 0 & -1 \end{bmatrix}$$

Harrison and Wein [49] develop a work-conserving threshold policy for this system. For a threshold $a \in \mathbb{N}_+$, server 1 prioritizes jobs in queue 1 if the number of jobs in queue 2 is below a. Otherwise, it prioritizes queue 3. This gives the scheduling action

$$u_{k,11} = 1\{x_{k,2} \le a\}, \quad u_{k,22} = 1\{x_{k,2} > 0\}, \quad u_{k,13} = (1 - u_{k,11})1\{x_{k,3} > 0\},$$

and the transition dynamics

$$\begin{aligned} x_{k+1} &= x_k + \begin{bmatrix} I_3 & R \end{bmatrix} e_{k+1} \\ e_{k+1} &= \arg\min\left\{\tau_{k,1}^A, \infty, \tau_{k,3}^A, \tau_{k,1}^S, \tau_{k,2}^S, \tau_{k,3}^S\right\} = \arg\min\left\{\tau_{k,1}^A, \infty, \tau_{k,3}^A, \frac{w_{k,1}}{\mu_{11}u_{k,11}}, \frac{w_{k,2}}{\mu_{22}u_{k,22}}, \frac{w_{k,3}}{\mu_{13}u_{k,13}}\right\} \\ \tau_{k+1}^* &= \min\left\{\tau_{k,1}^A, \infty, \tau_{k,3}^A, \tau_{k,1}^S, \tau_{k,2}^S, \tau_{k,3}^S\right\} \\ \begin{bmatrix} \tau_{k+1}^A \\ w_{k+1} \end{bmatrix} = \begin{bmatrix} \tau_k^A \\ w_k \end{bmatrix} - \tau_{k+1}^* \begin{bmatrix} \mathbf{1}_3 \\ \mathbf{diag}(u_k^\top \mu) \end{bmatrix} + \begin{bmatrix} T_{k+1} \\ W_{k+1} \end{bmatrix} \odot e_{k+1}. \end{aligned}$$

Here, $\tau_{k,2}^A = \infty$ since queue 2 has no external arrivals. \diamond

4 Gradient Estimation

In this section, we introduce our proposed approach for estimating the gradient of the objective (3), $\nabla J_N(\theta)$. We start with a brief discussion of existing methods for gradient estimation, including their

advantages and limitations. We then outline the main challenges for computing pathwise derivatives in multi-class queuing networks, and introduce our strategy for overcoming these challenges. Finally, we formally define our gradient estimation framework and discuss its computational and statistical properties. Later in section 5, we perform a comprehensive empirical study and find that our gradient estimation framework is able to overcome many of the limitations of existing methods in that (1) it is capable of estimating gradients for general queuing networks, (2) it provides stable gradient estimations over very long horizons (> 10^5 steps), (3) it provides greater estimation accuracy than model-free policy gradient methods with 1000x less data, and (4) when applying to policy optimization, it drastically improves the performance of the policy gradient algorithm for various scheduling and admission control tasks.



Figure 5. Multi-class, single-server queue.

Our goal is to optimize the parameterized control problem (3). A standard optimization algorithm is (stochastic) gradient descent, which has been considered for policy optimization and reinforcement learning [92, 8]. The core challenge for estimating policy gradient $\nabla J_N(\theta) = \nabla \mathbb{E} \left[J(\theta; \xi_{1:N}) \right]$ from sample paths of the queuing network is that the sample path cost $J(\theta, \xi_{1:N})$ is in general not differentiable in θ . As a consequence, one cannot change the order of differentiation and expectation, i.e.,

$$abla J_N(heta) =
abla \mathbb{E}\left[J_N(heta; \xi_{1:N})
ight]
eq \mathbb{E}\left[
abla J_N(heta; \xi_{1:N})
ight]$$

where $\nabla J_N(\theta; \xi_{1:N})$ is not even well-defined. The nondifferentiability of these discrete-event dynamical systems emerges

from two sources. First, actions $u_{1:N}$ are discrete scheduling decisions, and small perturbations in the policy can result in large changes in the scheduling decisions produced by the policy. Second, the actions affect the dynamics through the event times. The ordering of events is based on the 'argmin' of the residual event times, which is not differentiable.

In the stochastic simulation literature, there are two popular methods for gradient estimation: infinitesimal perturbation analysis (IPA) and generalized likelihood ratio (LR) gradient estimation. To illustrate, consider abstractly and with a little abuse of notation a system following the dynamics $s_{k+1} = f(s_k, \theta, \xi_{k+1})$, where $s_k \in \mathbb{R}$ is the state, $\theta \in \mathbb{R}$ is the parameter of interest, ξ_k is exogenous stochastic noise, and f is a differentiable function. Then, the IPA estimator computes a sample-path derivative estimator by constructing a derivative process $D_k = \partial s_k / \partial \theta$ via the recursion:

$$D_{k+1} = \frac{\partial}{\partial \theta} f(s_k, \theta, \xi_{k+1}) + \frac{\partial}{\partial s_k} f(s_k, \theta, \xi_{k+1}) \cdot D_k$$
(IPA)

Likelihood-ratio gradient estimation on the other hand uses knowledge of the distribution of ξ_k to form the gradient estimator. Suppose that s_k is a Markov chain for which the transition kernel is parameterized by θ , i.e., $s_{k+1} \sim p_{\theta}(\cdot|s_k)$. For a fixed θ_0 , let

$$\frac{\partial}{\partial \theta} \mathbb{E}_{\theta}[s_k] = \frac{\partial}{\partial \theta} \mathbb{E}_{\theta_0}[s_k L_k(\theta)] = \mathbb{E}_{\theta_0} \left[s_k \frac{\partial}{\partial \theta} L_k(\theta) \right] \quad \text{where} \quad L_k(\theta) := \frac{\prod_{j=1}^{k-1} p_{\theta}(s_{j+1}|s_j)}{\prod_{j=1}^{k-1} p_{\theta_0}(s_{j+1}|s_j)}$$

This allows one to obtain the following gradient estimator:

$$D_k = s_k \sum_{j=1}^{k-1} \frac{\frac{\partial}{\partial \theta} p_{\theta}(s_{j+1}|s_j)}{p_{\theta_0}(s_{j+1}|s_j)} L_k(\theta), \text{ where } s_{j+1} \sim p_{\theta_0}(\cdot|s_j), \forall j \le k.$$
(LR)

Despite their popularity, there are limitations to applying these methods to general multi-class queuing networks. While IPA has been proven efficient for simple queuing models, such as the G/G/1 queue through the Lindley recursion, it is well-known that unbiased IPA estimates cannot be obtained for general queuing networks [15, 32, 33]. The implementation of LR gradient estimation hinges on precise knowledge of the system's Markovian transition kernel [39]. This requires knowledge of the inter-arrival time and workload distributions, and even with this knowledge, it is non-trivial to specify the transition kernel of the queue lengths and residual event times in generic systems. Modifications to IPA [32, 33] also require precise knowledge of event time distributions and often involve analyzing specific ordering of events which must be done on a case-by-case basis. As a result, none of these methods can reliably provide gradient estimation for complex queuing networks under general scheduling policies and with possibly unknown inter-arrival and service time distributions. Yet, the ability to handle such instances is important to solve large-scale problems arising in many applications.

Due to the challenges discussed above, existing reinforcement learning (RL) approaches for queueing network control mainly rely on model-free gradient estimators, utilizing either the REINFORCE estimator and/or Q-function estimation. As we will discuss shortly, these methods do not leverage the structural properties of queuing networks and may be highly sample-inefficient, e.g., requiring a prohibitively large sample for gradient estimation.

To address the challenges discussed above, we propose a novel gradient estimation framework that can handle general, large-scale multi-class queuing networks under any differentiable scheduling policy, requiring only samples of the event times rather than knowledge of their distributions. Most importantly, our approach streamlines the process of gradient estimation, leveraging autodifferentiation libraries such as PyTorch [77] or Jax [13] to automatically compute gradients, rather than constructing these gradients in a bespoke manner for each network as is required for IPA or LR. As shown in Figure 2, computing a gradient in our framework requires only a few lines of code. To the best of our knowledge, this is the first scalable alternative to model-free methods for gradient estimation in queuing networks.

4.1 The standard approach: the REINFORCE estimator

Considering the lack of differentiability in most reinforcement learning environments, the standard approach for gradient estimation developed in model-free RL is the score-function or REINFORCE estimator [101, 92]. This serves as the basis for modern policy gradient algorithms such as Trust-Region Policy Optimization (TRPO) [82] or Proximal Policy Optimization (PPO) [84]. As a result, it offers a useful and popular baseline to compare our proposed method with.

The core idea behind the REINFORCE estimator is to introduce a randomized policy π_{θ} and differentiate through the action probabilities induced by the policy. Under mild regularity conditions on π_{θ} and $c(x_k, u_k)$, the following expression holds for the policy gradient:

$$\nabla J_N(\theta) = \mathbb{E}\left[\sum_{t=0}^{N-1} \left(\sum_{k=t}^{N-1} c(x_k, u_k) \tau_{k+1}^*\right) \nabla_\theta \log \pi_\theta(u_t | x_t)\right],$$

which leads to the following policy gradient estimator:

$$\widehat{\nabla}^{\mathsf{R}} J_N(\theta; \xi_{1:N}) = \sum_{t=0}^{N-1} \left(\sum_{k=t}^{N-1} c(x_k, u_k) \tau_{k+1}^* \right) \nabla_{\theta} \log \pi_{\theta}(u_t | x_t).$$
(REINFORCE)

While being unbiased, the REINFORCE estimator is known to have a very high variance [99]. The variance arises from two sources. First, the cumulative cost $\sum_{k=t}^{N-1} c(x_k, u_k) \tau_{k+1}^*$ can be very noisy, as has been observed for queuing networks [26]. Second, as the policy converges to the optimal policy, the score function $\nabla_{\theta} \log \pi_{\theta}(u_t|x_t)$ can grow large, magnifying the variance in the cost term. Practical implementations involve many algorithmic add-ons to reduce variance, e.g., adding a 'baseline' term [99] which is usually (an estimate of) the value function $V_{\pi_{\theta}}(x_k)$,

$$\widehat{\nabla}^{\mathsf{RB}} J_N(\theta; \xi_{1:N}) = \sum_{t=0}^{N-1} \left(\sum_{k=t}^{N-1} c(x_k, u_k) \tau_{k+1}^* - V_{\pi_\theta}(x_k) \right) \nabla_\theta \log \pi_\theta(u_t | x_t).$$
(BASELINE)

These algorithmic add-ons have led to the increased complexity of existing policy gradient implementations [57] and the outsized importance of various hyperparameters [56]. It has even been observed that seemingly small implementation "tricks" can have a large impact on performance, even more so than the choice of the algorithm itself [30].

4.2 Our approach: Differentiable Discrete-Event Simulation

We can view the state trajectory as a repeated composition of the transition function $s_{k+1} = f(s_k, u_k, \xi_{k+1})$, which is affected by exogenous noise $\xi_{1:N}$, i.e., stochastic inter-arrival and service times. If the transition function were differentiable with respect to the actions u_k , then under any fixed trace $\xi_{1:N}$, one could compute a *sample-path* derivative of the cost $J(\theta; \xi_{1:N})$ using auto-differentiation frameworks such as PyTorch [77] or Jax [13]. Auto-differentiation software computes gradients efficiently using the chain rule. To illustrate, given a sample path of states, actions, and noise $(s_k, u_k, \xi_{k+1})_{k=0}^{N-1}$, we can calculate the gradient of s_3 with respect to u_1 via

$$\frac{\partial s_3}{\partial u_1} = \frac{\partial s_3}{\partial s_2} \frac{\partial s_2}{\partial u_1} = \frac{\partial f(s_2, u_2, \xi_3)}{\partial s_2} \frac{\partial f(s_1, u_1, \xi_2)}{\partial u_1}.$$

This computation is streamlined through a technique known as backpropagation, or reverse-mode auto-differentiation. The algorithm involves two steps. The first step, known as the *forward pass*, evaluates the main function or performance metric (in the example, s_i 's) and records the partial derivatives of all intermediate states relative to their inputs (e.g. $\partial s_2/\partial u_1$). This step constructs a computational graph, which outlines the dependencies among variables. The second step is a *backward pass*, which traverses the computational graph in reverse. It sequentially multiplies and accumulates partial derivatives using the chain rule, propagating these derivatives backward through the graph until the gradient concerning the initial input (in this example, u_1) is calculated. Due to this design, gradients of functions involving nested compositions can be computed in a time that is linear in the number of compositions. By systematically applying the chain rule in reverse, auto-differentiation avoids the redundancy and computational overhead typically associated with numeric differentiation methods.

However, as mentioned before, the dynamics do not have a meaningful derivative due to the non-differentiability of actions and the **argmin** operation which selects the next event based on the minimum residual event time. Yet if we can utilize suitably differentiable surrogates, it would be possible to compute meaningful approximate sample-path derivatives using auto-differentiation.

4.2.1 Capacity sharing relaxation

First, we address the non-differentiability of the action space. Recall that $u_k \in \{0, 1\}^{m \times n}$ are scheduling decisions, which assign jobs to servers. Since u_k lies in a discrete space, a small change in the policy parameters can produce a jump in the actions. To alleviate this, we consider the transportation polytope as a continuous relaxation of the original action space (4):

$$\overline{\mathcal{U}} := \left\{ u \in [0,1]^{m \times n} : u \mathbf{1}_n = \mathbf{1}_m, \mathbf{1}_m^\top u = \mathbf{1}_n, u \le M \right\}.$$
(6)

The set of extreme points of $\overline{\mathcal{U}}$ coincide with the original, integral action space \mathcal{U} . For a fractional action $u_k \in \overline{\mathcal{U}}$, we can interpret it as servers splitting their capacity among multiple job classes motivated by the fluid approximation of queues [19]. As a relaxation, it allows servers to serve multiple jobs simultaneously. The effective service rate for each job class is equal to the fraction of the capacity allocated to the job class multiplied by the corresponding service rate.

As a result, instead of considering stochastic policies over discrete actions, we approach this problem as a *continuous* control problem and consider *deterministic* policies over continuous actions, i.e., the fractional scheduling decisions. Under this relaxation, the processing times are differentiable in the (fractional) scheduling decision. Finally, it is worth mentioning that we only use this relaxation when training policies. For policy evaluation, we enforce that actions are integral scheduling decisions in \mathcal{U} . To do so, we treat the fractional action as a probability distribution and use it to sample a discrete action.

Definition 1. Under the **capacity sharing relaxation**, the service rate for queue j under the routing decision $u \in \overline{\mathcal{U}}$ is $\mu_j^\top u_j \equiv \sum_{i=1}^m \mu_{ij} u_{ij}$. Thus, given workload w_j , the processing time of the job will be

$$\tau_j^S = \frac{w_j}{\sum_{i=1}^m u_{ij} \mu_{ij}} = \frac{w_j}{\mathbf{diag}(u^\top \mu)_{j,j}}.$$
(7)

Note that this is identical to the original definition of the processing times in (5). The only difference is that we now allow fractional routing actions, under which a server can serve multiple jobs at the same time.

For a concrete example, consider a single server *i* compatible with two job classes 1 and 2 with service rates $\mu_{i1} = 9$ and $\mu_{i2} = 15$ respectively. Suppose it splits its capacity between job classes 1 and 2 according to $u_{i1} = 1/3$ and $u_{i2} = 2/3$. Then for residual workloads w_1 and w_2 , the corresponding processing times are $\tau_1^S = w_1/3$ and $\tau_2^S = w_2/10$. If $u_{i1} = 0$ and $u_{i2} = 1$ instead, then the corresponding processing times are $\tau_1^S = w_1/0 = \infty$ and $\tau_2^S = w_2/15$.

4.2.2 Differentiable event selection

To determine the next event type, the argmin operation selects the next event based on the minimum residual event time. This operation does not give a meaningful gradient.

Pitfalls of 'naive' smoothing In order to compute gradients of the sample path, we need to smooth the argmin operation. There are multiple ways to do this. A naive approach is to directly replace argmin with a differentiable surrogate. One such popular surrogate is softmin. With some inverse temperature $\beta > 0$, softmin_{β} applied to the vector of residual event times $\tau \in \mathbb{R}^{2n}_+$ returns



Figure 6. Failure modes of 'naive' smoothing. (Left) Comparison of sample paths under the original dynamics and direct smoothing with $\beta = 1$ and $\beta = 1000$. Criss-cross network under a randomized backpressure policy with identical event times in each path, for N = 200 steps. Even under high inverse temperature $\beta = 1000$, the trajectory veers off from the original trajectory after only a hundred steps. (Right) Comparison of average cosine similarity (higher is better) of gradient estimators using direct smoothing with $\beta \in \{0.2, 0.5, 1, 2, 10\}$ for the criss-cross network under a randomized MaxWeight policy for N = 1000 steps. The gradient estimators either suffer from high bias or high variance and are unable to achieve a high cosine similarity with the true gradient.

a vector in \mathbb{R}^{2n}_+ , which we use to replace the event selection operation e_{k+1} :

$$\tilde{e}_{k+1} = \operatorname{softmin}_{\beta}(\tau_k) \quad \text{where} \quad \operatorname{softmin}_{\beta}(\tau)_j = e^{-\beta\tau_j} / \sum_{l=1}^{2n} e^{-\beta\tau_l}.$$
 (Direct Smoothing)

As $\beta \to \infty$, softmin_{β} converges to argmin. Thus, one may expect that for large β , softmin_{β} would give a reliable differentiable surrogate.

However, queuing networks involve a unique challenge for this approach: one typically considers very long trajectories when evaluating performance in queuing networks, as one is often interested in long-run average or steady-state behavior. Thus, even if one sets β to be very large to closely approximate **argmin**, the smoothing nonetheless results in 'unphysical', real-valued queue lengths instead of integral ones, and small discrepancies can accumulate over these long horizons and lead to entirely different sample paths. This can be observed concretely in the left panel of Figure 6, which displays the sample paths of the total queueing length processes for a criss-cross queueing network (in Example 3) under the original dynamic and under direct smoothing, using the same inter-arrival and service times. We observe that when setting the inverse temperature $\beta = 1$, the sample path under direct smoothing is completely different from the original one, even though all of the stochastic inputs are the same. Even when setting a very high inverse temperature, i.e., $\beta = 1000$, for which **softmin**_{β} is almost identical to **argmin**, the trajectory veers off after only a hundred steps.

This can greatly affect the quality of the gradient estimation. We observe in the right panel of Figure 6 that across a range of inverse temperatures, the average cosine similarity between the surrogate gradient and the true gradient (defined in (10)) are all somewhat low. In the same plot, we also show the average cosine similarity between our proposed gradient estimator, which we will discuss shortly, and the true gradient. Our proposed approach substantially improves the gradient estimation accuracy, i.e., the average cosine similarity is close to 1, and as we will show later, it

does so across a wide range of inverse temperatures.

Our approach: 'straight-through' estimation The failure of the direct smoothing approach highlights the importance of preserving the original dynamics, as errors can quickly build up even if the differentiable surrogate is only slightly off. We propose a simple but crucial adjustment to the direct smoothing approach, which leads to huge improvements in the quality of gradient estimation.

Instead of replacing the argmin operation with softmin_{β} when generating the sample path, we preserve the original dynamics as is, and only replace the Jacobian of argmin with the Jacobian of softmin_{β} when we query gradients. In short, we introduce a gradient operator $\widehat{\nabla}$ such that

$$e_{k+1} = \operatorname{argmin}(\tau_k), \qquad \nabla e_{k+1} = \nabla \operatorname{softmin}_{\beta}(\tau_k).$$
 (8)

where ∇ is respect to the input τ . This is known as the 'straight-through' trick in the machine learning literature and is a standard approach for computing approximate gradients in discrete environments [11, 97]. To the best of our knowledge, this is the first application of this gradient estimation strategy for discrete-event dynamical systems. Using this strategy, we can use the chain rule to compute gradients of performance metrics that depend on the event selection. Consider any differentiable function g of e_{k+1} ,

$$\widehat{\nabla}g(e_{k+1}) = \frac{\partial g(e_{k+1})}{\partial e_{k+1}} \widehat{\nabla}e_{k+1} \nabla \tau_k = \frac{\partial g(e_{k+1})}{\partial e_{k+1}} \nabla \mathrm{softmin}_\beta(\tau_k) \nabla \tau_k$$

In contrast, direct smoothing involves the derivative $\partial g(\tilde{e}_{k+1})/\partial \tilde{e}_{k+1}$ where $\tilde{e}_{k+1} = \operatorname{softmin}_{\beta}(\tau_k)$. Evaluating the gradient of g at $g(\tilde{e}_{k+1})$ is a cause of additional bias.

With these relaxations, the transition function of the system is differentiable. We can now compute a gradient of the sample path cost $J_N(\theta; \xi_{1:N})$, using the chain rule on the transition functions. Given a sample path of states $s_k = (x_k, z_k)$, actions $u_k = \pi_{\theta}(x_k)$, and $\xi_k = (T_k, W_k)$, the pathwise gradient of the sample path cost is $J_N(\theta; \xi_{1:N})$ with respect to an action u_k is,

$$\widehat{\nabla}_{u_k} J_N(\theta; \xi_{1:N}) = \underbrace{\nabla_{u_k} c(x_k, u_k)}_{\text{current cost}} + \sum_{t=k+1}^N \left[\underbrace{\nabla_x c(x_t, u_t) + \nabla_u c(x_t, u_t) \nabla_{x_t} \pi_\theta(x_t)}_{\text{future costs}} \right] \nabla_{u_k} x_t$$

The gradient consists of the sensitivity of the current cost with respect to the action as well the sensitivity of future costs via the current action's impact on future states. The policy gradient with respect to θ can then be computed as

$$\widehat{\nabla}_{\theta} J_N(\theta; \xi_{1:N}) = \sum_{k=1}^N \left. \widehat{\nabla}_{u_k} J_N(\theta; \xi_{1:N}) \right|_{u_k = \pi_{\theta}(x_k)} \nabla_{\theta} \pi_{\theta}(x_k).$$
(PATHWISE)

As a result of the straight-through trick, we do not alter the event selection operation e_k 's and thus the state trajectory $\{x_k\}_{k=1}^N$ is unchanged.

We refer to the gradient estimator (PATHWISE) as the PATHWISE policy gradient estimator. Although this formula involves iterated products of several gradient expressions, these can be computed efficiently through reverse-mode auto-differentiation using libraries such as PyTorch [77] or Jax [13] with O(N) time complexity in the time horizon N. This time complexity is of the same order as the forward pass, i.e., generating the sample path itself, and is equivalent to the time complexity of REINFORCE. The policy gradient algorithm with PATHWISE gradient is summarized in Algorithm 1. In Section 5, we perform a careful empirical comparison of PATHWISE and

Algorithm 1 PATHWISE Policy Gradient (PathPG)

- 1: INPUT: POLICY π_{θ} , Number of Iterates T, horizon N, trace $\xi_{1:N}$, step-size $\alpha > 0$.
- 2: for each $t \in 1, \ldots, T$ do
- 3: Compute $J_N(\theta; \xi_{1:N}) = \sum_{k=1}^N c(x_k, \pi_\theta(x_k)) \tau_{k+1}^*$ from trace $\xi_{1:N}$.
- 4: Compute gradient $\widehat{\nabla}_{\theta} J_N(\theta; \xi_{1:N})$ via (PATHWISE)
- 5: Update policy parameters,

$$\theta_{t+1} \leftarrow \theta_t - \alpha \widehat{\nabla}_{\theta} J_N(\theta; \xi_{1:N})$$

6: **end for**

7: return π_{θ_T}

REINFORCE, and find that PATHWISE can lead to orders of magnitude improvements in sample efficiency.

It is important to re-emphasize that when computing the gradient, we evaluate the policy differently than we would for REINFORCE. Instead of drawing a random discrete action $u_k \sim \pi_{\theta}(x)$, we use the probabilities output by the policy directly as a fractional routing matrix in $\overline{\mathcal{U}}$,

REINFORCE :	$u_k \sim \pi_\theta(x_k),$	$u_k \in \mathcal{U}$
PATHWISE :	$u_k = \pi_\theta(x_k),$	$u_k \in \overline{\mathcal{U}}.$

However, this is only for gradient computation. When we evaluate the policy, we draw $u_k \sim \pi_{\theta}(x_k)$.

Bias-variance trade-off for inverse temperature: One-step analysis The inverse temperature β is a key hyperparameter that determines the fidelity of the softmin_{β} approximation to argmin. The choice of β poses a bias-variance trade-off, with a higher β leading to a smaller bias but a higher variance and a smaller β incurring a higher bias but a lower variance.

In general, it is difficult to assess the bias since we often do not know the true gradient. However, for some simple examples, we can evaluate the true gradient explicitly. We next analyze the gradient of the one-step transition of the M/M/1 queue with respect to the service rate μ ,

$$\nabla_{\mu}\mathbb{E}[x_{k+1} - x_k] = \nabla_{\mu}\mathbb{E}[De_{k+1}] = \nabla_{\mu}\frac{\lambda - \mu}{\lambda + \mu} = -2\frac{\lambda}{(\lambda + \mu)^2}.$$

This permits an exact calculation of the mean and variance of our proposed pathwise gradient estimator. Although we can derive analytical expressions for these quantities, we present the leading order asymptotics as $\beta \to \infty$ for conciseness of presentation. While it is straightforward to see that almost surely $\operatorname{softmin}_{\beta}(\tau) \to \operatorname{argmin}(\tau)$ as $\beta \to \infty$, it is much less clear whether the gradient converges, i.e., whether $\mathbb{E}[\nabla \operatorname{softmin}_{\beta}(\tau)] \to \nabla \mathbb{E}[\operatorname{argmin}(\tau)]$. Since argmin has a gradient of zero almost everywhere, the expectation and gradient operators cannot be interchanged. Instead, we analyze the expectations directly using properties of the exponential distribution.

Theorem 1. Let $\widehat{\nabla}_{\mu}(x_{k+1} - x_k) = \widehat{\nabla}_{\mu} De_{k+1}$ denote the PATHWISE gradient estimator of the one-step transition of the M/M/1 queue with respect to μ . For $x_k \ge 1$, as $\beta \to \infty$,

$$\mathbb{E}[\widehat{\nabla}_{\mu}De_{k+1}] - \nabla_{\mu}\mathbb{E}[De_{k+1}] = \beta^{-2} \cdot \frac{\pi^2\lambda(\mu^2 - \lambda^2 + 2\mu\lambda)}{6(\lambda + \mu)^2} + o\left(\beta^{-2}\right),$$



Figure 7. Bias-variance trade-off for inverse temperature β . (Left) Bias and variance of pathwise gradient estimator across a range of inverse temperatures β for the one-step change $\widehat{\nabla}_{\mu}(x_{k+1} - x_k)$ compared to true gradient $\nabla_{\mu}\mathbb{E}[x_{k+1} - x_k] = \nabla_{\mu}\frac{\lambda-\mu}{\lambda+\mu}$. Blue line specifies the theoretical values for the bias and variance. (Right) Average cosine similarity (see (10) for the definition) of PATHWISE policy gradients with the true gradient for a randomized max-weight policy (see (9)) in a 6-class reentrant network (see Figure 13) across a range of inverse temperatures computed from a single trajectory B = 1 for N = 1000 steps. True gradient is computed by averaging REINFORCE over 10^6 trajectories. PATHWISE gradient has almost a perfect cosine similarity ≈ 1 across a wide range of inverse temperatures.

$$\operatorname{Var}(\widehat{\nabla}_{\mu} D e_{k+1}) = \beta \cdot \frac{4\lambda}{\mu(\lambda+\mu)^2} + o(\beta).$$

See section B.1 for the proof. As $\beta \to \infty$, the bias is $O(1/\beta^2)$ while the variance is $O(\beta)$. This means that one can significantly reduce the bias with only a moderate size of β . The left panel of Figure 7 shows the bias-variance trade-off of $\widehat{\nabla}_{\mu}(x_{k+1} - x_k)$ for various inverse temperatures β . The blue line is based on the analytical expression for the bias-variance trade-off curve. We observe that with the inverse temperature $\beta \in [0.5, 2]$, both the bias and the variance are reasonably small.

Corollary 1. Suppose we compute the sample average of B iid samples of $\widehat{\nabla}_{\mu} De_{k+1}$, which are denoted as $\widehat{\nabla}_{\mu} De_{k+1,i}$, $i = 1, \ldots, K$. In particular, the estimator takes the form $\frac{1}{B} \sum_{i=1}^{B} \widehat{\nabla}_{\mu} De_{k+1,i}$. The choice of β that minimizes the mean-squared error (MSE) of the estimator is $\beta^* = O(B^{1/5})$ and $MSE(\beta^*) = O(B^{-4/5})$.

The PATHWISE estimator provides a more statistically efficient trade-off than other alternatives. As an example, a standard gradient estimator is the finite-difference estimator in which one evaluates the one-step transition at $\mu - h$ and $\mu + h$ for some small $h \in (0, \infty)$, and the estimator is constructed as

$$\frac{1}{B}\sum_{i=1}^{B}\frac{De_{k+1,i}(\mu+h) - De_{k+1,i}(\mu-h)}{2h}$$

where $De_{k+1,i}(\mu + h)$'s are iid samples of $De_{k+1}(\mu + h)$. If we set $h = 1/\beta$, it is well-known that the bias scales as $O(1/\beta^2)$ while the variance scales as $O(\beta^2)$. The choice of β that minimizes the MSE is $\beta^* = O(B^{1/6})$ and $MSE(\beta^*) = O(B^{-1/3})$.

While this analysis is restricted to the one-step transition of the M/M/1 queue, these insights hold for more general systems and control problems. The right panel of Figure 7 displays the average cosine similarity (defined in (10)) between the PATHWISE gradient estimator and the true gradient for a policy gradient task in a 6-class reentrant network across different congestion levels and for different inverse temperatures. We observe that for a wide range of inverse temperatures, $\beta \in \{0.2, 0.5, 1, 2\}$, the estimator has near-perfect similarity with the true gradient, while a very large inverse temperature suffers due to high variance. This indicates that while there is a bias-variance trade-off, the performance of the PATHWISE gradient estimator is not sensitive to the choice of the inverse temperature within a reasonable range. In our numerical experiments, we find that one can get good performance using the same inverse temperature across different settings without the need to tune it for each setting.

5 Empirical Evaluation of the PATHWISE Gradients

In the previous section, we introduced the PATHWISE gradient estimator for computing gradients of queuing performance metrics with respect to routing actions or routing policy parameters. In this section, we study the statistical properties of these gradient estimators and their efficacy in downstream policy optimization tasks. We use REINFORCE as the baseline gradient estimator. First, in section 5.1 we empirically study the estimation quality across a range of queuing networks, traffic intensities, and policies. After that, in section 5.2, we investigate their performance in a scheduling task: learning the $c\mu$ rule in a multi-class queuing network. Finally, we demonstrate the applicability of our framework beyond scheduling: we investigate the performance of the PATHWISE gradient estimator for admission control tasks in section 5.3.

5.1 Gradient Estimation Efficiency

In general, it is challenging to theoretically compare the statistical properties of different gradient estimators, and very few results exist for systems beyond the M/M/1 queue (see section 8 for a theoretical comparison between REINFORCE and PATHWISE for the M/M/1 queue). For this reason, we focus on numerical experiments across a range of environments and queuing policies typically considered in the queuing literature. Specifically, we will be comparing the statistical properties of PATHWISE estimator with the baseline estimator REINFORCE. While PATHWISE introduces bias into the estimation, we find in our experiments that this bias is small in practice and remains small even over long time horizons. At the same time, the PATHWISE estimator delivers dramatic reductions in variance, achieving greater accuracy with a single trajectory than REINFORCE with 10^3 trajectories.

First, recall that a policy $\pi(x)$ maps queue-lengths x to assignment between servers and queues, represented by an $m \times n$ matrix in $\overline{\mathcal{U}}$ (allowing for fractional routing matrices). We visit three classical queuing policies: priority policies [25], MaxWeight [93], and MaxPressure [27]. Each of these methods selects the routing that solves an optimization problem. This means that the routing generated by the policy is deterministic given the state and is not differentiable in the policy parameters. In order to apply either REINFORCE or the PATHWISE gradient estimator to compute a policy gradient, we require differentiable surrogates of these policies. To this end, we define softened and parameterized variants of these policies, denoted as soft priority (sPR), soft MaxWeight (sMW), and soft MaxPressure (sMP),

$$\pi_{\theta}^{\mathsf{sPR}}(x)_i = \mathsf{softmax}(\theta_j \cdot \mu_i), \quad \pi_{\theta}^{\mathsf{sMW}}(x)_i = \mathsf{softmax}(\theta_j x_j \cdot \mu_i), \quad \pi_{\theta}^{\mathsf{sMP}}(x)_i = \mathsf{softmax}((\mu \odot R(\theta x))_i)$$
(9)



Figure 8. Comparison of estimation quality between PATHWISE and REINFORCE across several settings. We perform this comparison for 3 policies (soft MaxPressure, soft MaxWeight, soft Priority in (9)), 9 network settings, and 4 levels of traffic intensity for each network. For each cell, we randomly draw 100 parameters $\theta \in \mathbb{R}^n$. For each parameter, we estimate the true gradient by averaging the REINFORCE estimator over 10^6 trajectories (with horizon N = 1000). We then compute the PATHWISE estimator with B = 1 trajectory along with the REINFORCE estimator averaged across B = 1000 trajectories. In each instance, we draw 100 samples of these estimators to estimate sim, i.e., the average cosine similarity with the true gradient. In total, we perform this comparison across 10,800 parameters. We find that across all of these diverse settings, PATHWISE delivers much higher fidelity to the true gradient, in many cases achieving an average cosine similarity close to the maximum value of 1, despite using orders of magnitude less data, whereas the cosine similarity of REINFORCE remains around 0.2-0.6 even with 1000 trajectories.

where $\theta \in \mathbb{R}^n_+$ are a vector of costs/weights for each queue, μ denotes the matrix of service rates with $\mu_i \in \mathbb{R}^n_+$ denoting the service rates associated with server *i*. The operation \odot refers to elementwise multiplication and the softmax operation maps a vector $a \in \mathbb{R}^n$ into a set of probabilities $\operatorname{softmax}(a)_i = e^{a_i} / \sum_{j=1}^n e^{a_j}$.

We are interested in identifying the parameter θ that minimizes long-run average holding cost where $c(x, u) = h^{\top}x$. We use the objective $J_N(\theta) = \mathbb{E}\left[\sum_{k=0}^{N-1} c(x_k, \pi_{\theta}(x_k))\tau_{k+1}^*\right]$ where N is a large enough number to approximate the long-run performance, and the goal of the gradient estimation is to estimate $\nabla J_N(\theta)$.

We consider the following environments, which appear throughout our computational experiments and serve as standard benchmarks for control policies in multi-class queuing networks. We describe the network structure in detail in Figure 13.

- Criss-cross: The network introduced in Example 3 (see Figure 13 (c)).
- Re-entrant 1 (*n* classes): We consider a family of multi-class re-entrant networks with a varying number of classes, which was studied in [12, 26]. The network is composed of several layers and each layer has 3 queues. Jobs processed in one layer are sent to the next

layer. Arrivals to the system come to queues 1 and 3 in the first layer while queue 2 receives re-entered jobs from the last layer (see Figure 13 (a) for a two-layer example).

• Re-entrant 2 (*n* classes): We consider another family of re-entrant network architecture that was studied in [12]. It also consists of multiple layers with 3 queues in each layer. It differs from the Re-entrant 1 environment in that only queue 1 receive external arrivals while queues 2 and 3 receive re-entered jobs from the last layer (see Figure 13 (b) for a two-layer example).

For a gradient estimator \hat{g} , the main performance metric we evaluate is $sim(\hat{g})$, which is the expected cosine similarity with the ground-truth gradient,

$$\sin(\hat{g}) \equiv \mathbb{E}[\cos\left(\hat{g}, \nabla J_N(\theta)\right)] = \mathbb{E}\left[\frac{\langle \hat{g}, \nabla J_N(\theta) \rangle}{\|\hat{g}\| \|\nabla J_N(\theta)\|}\right] \in [-1, 1]$$
(10)

where the expectation \mathbb{E} is over randomness in \hat{g} . The higher the similarity is, the more aligned \hat{g} is to the direction of $\nabla J_N(\theta)$. This metric incorporates both bias and variance of the gradient estimator. If the gradient estimator is unbiased but has a high variance, then each individual realization of \hat{g} is likely to have low correlation with the true gradient, so the average cosine similarity will be small even if $\mathbb{E}[\hat{g}] = \nabla J_N(\theta)$. At the same time, if the gradient estimator has a low variance but a high bias, then the $\sin(\hat{g})$ could still be small if $\cos(\mathbb{E}[\hat{g}], \nabla J_N(\theta))$ is small. We focus on this metric, because it directly determines how informative the gradient estimates are when applying various gradient descent algorithms. For our experiments, we evaluate (a close approximation of) the ground-truth gradient $\nabla J_N(\theta)$ by using the unbiased REINFORCE gradient estimator over exceedingly many trajectories (in our case, 10^6 trajectories).

We compare the similarity of PATHWISE with that of REINFORCE. We denote B as the number of trajectories we use to calculate each PATHWISE or REINFORCE gradient estimator.

$$\underbrace{\hat{\nabla}_{\theta} J_N(\theta; \xi_{1:N}^{(1)})}_{\text{PATHWISE with } B=1} \qquad \underbrace{\hat{\nabla}_{\theta}^{\mathsf{R}} J_{N,B}(\theta; \xi_{1:N}) := \frac{1}{B} \sum_{b=1}^{B} \hat{\nabla}_{\theta}^{\mathsf{R}} J_N(\theta; \xi_{1:N}^{(b)})}_{\text{BEINFORCE with } B \text{ trajectories}}$$
(11)

We compute the PATHWISE gradient with only B = 1 trajectory, while REINFORCE gradient is calculated using $B = 10^3$ trajectories. For each policy and setting, we compute these gradients for 100 different randomly generated values of θ , which are drawn from a Lognormal(0, 1) distribution (as the parameters must be positive in these policies). In total, we compare the gradients in 10,080 unique parameter settings, and each gradient estimator is computed 100 times to evaluate the average cosine similarity. When computing the policy gradient, we consider a time horizon of $N = 10^3$ steps.

Figure 8 compares the PATHWISE estimator with B = 1 trajectory with the REINFORCE estimator averaged over $B = 10^3$ trajectories. For the REINFORCE estimator, costs are computed with a discount factor $\gamma = 0.999$, as using a lower discount rate introduced significant bias in the estimation. For PATHWISE, we use an inverse temperature $\beta = 1$ for the softmin relaxation across all settings. Each cell in Figure 8 corresponds to a (policy, network, traffic-intensity) and the cell value is the average expected cosine similarity of the estimator averaged across the 100 randomly drawn θ values. We observe that across these diverse settings, the PATHWISE estimator consistently has a much higher average cosine similarity with the true gradient despite using only a single trajectory. In fact, for 94.5% of the 10,800 parameter settings, PATHWISE has a higher average cosine similarity with 99% confidence than REINFORCE with B = 1000 trajectories. In most cases, the cosine similarity of PATHWISE is close to 1, indicating almost perfect alignment with the true gradient even under high congestion. REINFORCE on the other hand suffers greatly from high variance. Overall, this demonstrates that PATHWISE is able to deliver greater estimation accuracy with an order of magnitude fewer samples.

5.2 Learning the $c\mu$ rule

Given the strong improvements in estimation efficiency, we turn to evaluate how these translate to a downstream optimization task. In single-server multi-class queues, it is well-known that the $c\mu$ -rule minimizes the long-run average holding cost [25]. We assess whether gradient descent with the REINFORCE or PATHWISE gradients is capable of converging to the $c\mu$ -rule, without knowing the holding costs h or μ and only using feedback from the environment. Despite its simplicity, it has been observed in prior work that this is a difficult learning task, particularly under heavy traffic [95].

We revisit the soft priority policy mentioned before, but with only the parameters $\theta \in \mathbb{R}^n$, i.e.,

$$\pi_{\theta}^{\mathsf{sPR}}(x)_i = \mathsf{softmax}(\theta_i) \tag{12}$$

We also modify the policy to ensure that it is work-conserving, i.e., not assigning the server to an empty queue (see section 6 for further discussion).

We consider a family of multi-class single-server queues with n queues. Holding costs are identically $h_j = 1$. Inter-arrival and service times are exponentially distributed, the service rates are $\mu_{1j} = 1 + \epsilon j$, for some $\epsilon > 0$, and the arrival rates are identical $\lambda_j = \lambda$ and λ are set such that the traffic intensity $\sum_{j=1}^{n} \frac{\lambda}{\mu_{1j}} = \rho$ for some $\rho \in (0, 1)$. Note that in this case, the $c\mu$ -rule prioritizes queues with higher indices j. We consider a grid of gap sizes $\epsilon \in \{1.0, 0.5, 0.1, 0.05, 0.01\}$ to adjust the difficulty of the problem; the smaller ϵ is, the harder it is to learn.

We compare PATHWISE with B = 1 trajectory and REINFORCE with B = 100 trajectories for trajectories of N = 1000 steps. In order to isolate the effect of the gradient estimator from the optimization scheme, for both estimators we use an identical stochastic gradient descent scheme with normalized gradients (as these two estimators may differ by a scale factor). That is, for gradient estimator \hat{g} , the update under step-size α is

$$\theta_{t+1} = \theta_t - \alpha \frac{\hat{g}}{\|\hat{g}\|}.$$

We run T gradient descent steps for each gradient estimator. To allow for the fact that different estimators may have different performances across different step sizes, we consider a grid of step sizes $\alpha \in \{0.01, 0.1, 0.5, 1.0\}$. Gradient normalization may prevent convergence, so we use the averaged iterate $\bar{\theta}_T$ for T. We then evaluate the long-run average holding cost under a strict priority policy determined by $\bar{\theta}_T$, i.e., $\pi_{\bar{\theta}_T}(x)_i = \operatorname{argmax} \bar{\theta}_{T,j}$.

The left panel of Figure 9 displays the values of $\bar{\theta}_T$ after T = 50 gradient iterates for PATHWISE and REINFORCE with n = 5, $\epsilon = 0.1$, and $\rho = 0.99$. We observe that while PATHWISE sorts the queues in the correct order (it should be increasing with the queue index), REINFORCE even with B = 100 trajectories fails to prioritize queues with a higher $c\mu$ index. Remarkably, we observe in the right panel of the same figure that PATHWISE with just a single trajectory achieves a lower average holding cost than REINFORCE uniformly across various step sizes and difficulty levels, whereas the performance of REINFORCE varies greatly depending on the step size. This indicates



Figure 9. Learning the $c\mu$ rule. (Left) Averaged iterate of the policy parameters after 50 gradient steps with PATHWISE (B = 1) and REINFORCE (B = 100) gradients for a 5-class queue with traffic intensity $\rho = 0.99$ and gap-size $\epsilon = 0.1$. The scores obtained by PATHWISE are increasing in the queue index, which matches the ordering of the $c\mu$ index in this instance. The scores obtained by REINFORCE do not achieve the correct ordering despite using more trajectories. (Right) Average holding cost of the average iterate after 20 steps of gradient descent in a 10-class queue with $\rho = 0.95$. This figure reports the results averaged over 50 separate runs of gradient descent, across a grid of step-sizes (denoted as α). Remarkably, the optimization performance of the PATHWISE estimator is highly similar across step-sizes, and *uniformly* outperforms REINFORCE with different step-sizes α .

that the improvements in gradient efficiency/accuracy of PATHWISE make it more robust to the step-size hyper-parameter. It is also worth mentioning that when gap size ϵ becomes smaller, it is more difficult to learn. At the same time, since μ_{1j} 's are more similar to each other, the cost difference between different priority rules also diminishes.

5.3 Admission Control

While we focus mainly on scheduling tasks in this work, our gradient estimation framework can also be applied to admission control, which is another fundamental queuing control task [75, 21, 22, 34, 62]. To manage congestion, the queuing network may reject new arrivals to the network if the queue lengths are above certain thresholds. The admission or buffer control problem is to select these thresholds to balance the trade-off between managing congestion and ensuring sufficient resource utilization.

Under fixed buffer sizes $L = \{L_j\}_{j=1}^n$, new arrivals to queue j are blocked if $x_j = L_j$. As a result, the state update is modified as follows,

$$x_{k+1} = \min\{x_k + De_{k+1}, L\}.$$
(13)

While a small L can greatly reduce congestion, it can impede the system throughput. To account for this, we introduce a cost for rejecting an arrival to the network. Let $o_k \in \{0, 1\}^n$ denote whether an arrival is *overflowed*, i.e., an arrival is blocked because the buffer is full,

$$o_{k+1} = De_{k+1} \cdot 1\{x_k + De_{k+1} > L\}.$$
(14)

Given a fixed routing policy, the control task is to choose the buffer sizes L to minimize the holding



Figure 10. Gradient descent with PATHWISE for admission control tasks. (Left) Iterates of sign gradient descent with the PATHWISE estimator for the buffer size of an M/M/1 queue with holding cost h = 1 and overflow cost b = 100. Starting from $L_0 = 1$, the iterates quickly converge to the optimal level of $L^* = 14$. (Right) buffer sizes obtained by sign gradient descent with the PATHWISE estimator for a 2 queue 1 server network with holding cost h = 1 and overflow cost b = 20. Starting from $L_0 = (1, 1)$, the iterates quickly converge to the basin of the loss surface.

and overflow costs:

$$J_N(L;\xi_{1:N}) = \sum_{k=0}^{N-1} (h^{\top} x_k) \tau_{k+1}^* + b^{\top} o_k.$$
(15)

Similar to the routing control problem, despite the fact that overflow is discrete, our gradient estimation framework is capable of computing a PATHWISE gradient of the cost with respect to the buffer sizes, which we denote as $\widehat{\nabla}_L J_N(L; \xi_{1:N})$, i.e., we can evaluate gradients at integral values of the buffer size and use this to perform updates. Since the buffer sizes must be integral, we update the buffer sizes via sign gradient descent to preserve integrality:

$$L_{t+1} = L_t - \operatorname{sign}\left(\widehat{\nabla}_L J_N(L;\xi_{1:N})\right).$$
(16)

Learning for admission control has been studied in the queuing and simulation literature [16, 17, 51, 23]. While exact gradient methods are possible in fluid models [16, 17], the standard approach for discrete queuing models is finite perturbation analysis [51], given the discrete nature of the buffer sizes. Randomized finite-differences, which is also known as Simultaneous Perturbation Stochastic Approximation (SPSA) [88, 31], is a popular optimization method for discrete search problems. This method forms a finite-differences gradient through a random perturbation. Let $\eta \sim \text{Rademacher}(n, 1/2) \in \{-1, 1\}^n$ be a random *n*-dimensional vector where each component is an independent Rademacher random variable, taking values in $\{-1, 1\}$ with equal probability. For each perturbation η , we evaluate the objective at $L \pm \eta$, i.e., $J_N(L + \eta; \xi_{1:N})$ and $J_N(L - \eta; \xi_{1:N})$, using the same sample path for both evaluations to reduce variance. For improved performance, we average the gradient across a batch of *B* perturbations, i.e., $\eta^{(b)}$ for b = 1, ..., B, drawing a new sample path $\xi_{1:N}^{(b)}$ for each perturbation. The batch SPSA gradient is

$$\widehat{\nabla}_{L}^{\mathsf{SPSA},B} J_{N}(L) = \frac{1}{B} \sum_{b=1}^{B} \frac{1}{2} \left(J_{N}(L+\eta^{(b)};\xi_{1:N}^{(b)}) - J_{N}(L-\eta^{(b)};\xi_{1:N}^{(b)}) \right) \eta^{(b)}.$$
(17)

We update the buffer sizes L according to the same sign gradient descent algorithm as in (16).

In comparison with existing works in the queuing literature (e.g. [75, 22, 34]), which derive analytical results for simple single-class or multi-class queues, we consider admission control tasks for large, re-entrant networks with multiple job classes. Each job class has its own buffer, resulting in a high-dimensional optimization problem in large networks. Moreover, the buffer size for one job class affects downstream congestion due to the re-entrant nature of the networks. For our experiments, we fix the scheduling policy to be the soft priority policy $\pi_{\theta}^{\text{sPR}}(x)$ in (12) due to its simplicity and strong performance in our environments. We emphasize however that our framework can be applied to buffer control tasks under any differentiable routing policy, including neural network policies. For each gradient estimator, we perform T = 100 iterations of sign gradient descent, and each gradient estimator is computed from trajectories of length N = 1000. For SPSA, we consider batch sizes of $B = \{10, 100, 1000\}$ whereas we compute PATHWISE with only B = 1trajectory. When evaluating the performance, we calculate the long-run average cost with the buffer size determined by the last iterate with a longer horizon $N = 10^4$ and over 100 trajectories. We also average the results across 50 runs of sign gradient descent.

The left panel of Figure 10 displays iterates of the sign gradient descent algorithm with PATHWISE for the M/M/1 queue with holding cost h = 1 and overflow cost b = 100. We observe that sign gradient descent with PATHWISE (computed over a horizon of N = 1000 steps) quickly reaches the optimal buffer size of $L^* = 14$ and remains there, oscillating between L = 14 and 15. The right panel shows the iterates for a simple 2-class queue with 1 server, h = 1, and b = 20 under a soft priority policy. We again observe that sign gradient descent with PATHWISE quickly converges to a near-optimal set of buffer sizes.

To see how the estimator performs in larger-scale problems, we consider the Re-entrant 1 and Re-entrant 2 networks introduced in Section 5.1 with varying number of job classes (i.e., varying number of layers). Figure 11 compares the last iterate performance of SPSA and PATHWISE for these two families of queuing networks with instances ranging from 6-classes to 21-classes. Holding costs are h = 1 and overflow costs are b = 1000 for all queues. We observe that PATHWISE with only a single trajectory is able to outperform SPSA with B = 1000 trajectories for larger networks. Sign gradient descent using SPSA with only B = 10 trajectories is much less stable, with several of the iterations reaching a sub-optimal set of buffer sizes that assign $L_j = 0$ to several queues. This illustrates the well-known fact that for high-dimensional control problems, zeroth-order methods like SPSA must sample many more trajectories to cover the policy space and their performance can scale sub-optimally in the dimension. Yet PATHWISE, which is an approximate first-order gradient estimator, exhibits much better scalability with dimension and is able to optimize the buffer sizes with much less data.

6 Policy Parameterization

While our gradient estimation framework offers a sample-efficient alternative for learning from the environment, there is another practical issue that degrades the performance of learning algorithms for queuing network control: instability. Standard model-free RL algorithms are based on the



Figure 11. Last iterate performance of PATHWISE ($\beta = 1$) and SPSA [88] on admission control for re-entrant networks. (Left) Average cost of the last iterate of sign gradient descent (100 iterations) for the Re-entrant 1 networks, across different numbers of job classes. While SPSA averaged across B = 100 trajectories achieves good performance, using a smaller batch B = 10 leads to instabilities that lead to much higher costs. PATHWISE with only 1 trajectory is able to consistently achieve a smaller cost, especially for larger networks. (Right) Average cost of the last iterate of sign gradient descent (100 iterations) for the Re-entrant 2 networks. Even with B = 1000 trajectories, SPSA is unable to effectively optimize the buffer sizes for larger networks, reaching the same cost as SPSA with B = 10. This illustrates that the performance of finite-difference methods such as SPSA degrade in higher-dimensional problems, whereas PATHWISE performs well in these larger instances using much less data.

'tabula rasa' principle, which aims to search over a general and unstructured policy class in order to find an optimal policy. However, it has been observed that this approach may be unsuitable for queuing network control. Due to the lack of structure, the policies visited by the algorithm often fail to stabilize the network, which prevents the algorithm from learning and improving. As a result, researchers have proposed structural modifications to ensure stability, including behavior cloning of a stabilizing policy to find a good initialization [26], switching to a stabilizing policy if the queue lengths exceed some finite thresholds [64], or modifying the costs to be stability-aware [78].

We investigate the source of instability in various queuing scheduling problems and find a possible explanation. We note that many policies obtained by model-free RL algorithms are not **work-conserving** and often allocate servers to empty queues. A scheduling policy is work-conserving if it always keeps the server(s) busy when there are compatible jobs waiting to be served. Standard policies such as the $c\mu$ -rule, MaxPressure, and MaxWeight are all work-conserving, which partly explains their success in stabilizing complex networks. We treat work conservation as an 'inductive bias' and consider a simple modification to the policy architecture that guarantees this property without sacrificing the flexibility of the policy class.

The de-facto approach for parameterizing policies in deep reinforcement learning is to consider a function $\nu_{\theta}(x)$, which belongs to a general function family, such as neural networks, and outputs real-valued scores. These scores are then fed into a softmax layer, which converts the scores to probabilities over actions. Naively, the number of possible routing actions can grow exponentially in the number of queues and servers. Nonetheless, one can efficiently sample from the action space by having the output of $\nu_{\theta}(x) \in \mathbb{R}^{m \times n}$ be a matrix where row *i*, denote as $\nu_{\theta}(x)_i$, contains the scores for matching server *i* to different queues. Then by applying the softmax for row *i*,



Figure 12. Performance of PPO with and without work-conserving softmax for the Re-entrant 1 network with 6 classes. Average holding cost of PPO without any modifications, PPO initialized from a behavior cloned policy (PPO BC), and PPO with the work-conserving softmax (PPO WC). Average cost of the $c\mu$ -rule added for reference. Without any modification, PPO is unable to stabilize the queues, resulting in an average queue-length of 10^3 and does not improve over time. Behavior cloning provides a much better initialization, and the policy improves with training. However, it fails to improve over the $c\mu$ -rule. With the work conserving softmax, even the randomly initialized policy is capable of stabilizing the network – achieving an equivalent cost as the $c\mu$ -rule – and is able to outperform the $c\mu$ over the course of training.

i.e., $\operatorname{softmax}(\nu_{\theta}(x)_i)$, we obtain the probability that server *i* is assigned to each queue. We then sample the assignment independently for each server to obtain an action in \mathcal{U} . For the purpose of computing the PATHWISE estimator, $\operatorname{softmax}(\nu_{\theta}(x)_i)$ also gives a valid fractional routing in $\overline{\mathcal{U}}$. We let $\operatorname{softmax}(\nu_{\theta}(x)) \in \overline{\mathcal{U}}$ denote the matrix formed by applying the softmax to each row in $\nu_{\theta}(x)$.

Under this 'vanilla' softmax policy, the probability $\pi_{\theta}(x)_{ij}$ that server *i* is routed to queue *j* (or alternatively, the fractional capacity server *i* allocated to *j*) is given by

$$\pi_{\theta}(x)_{ij} = \operatorname{softmax}(\nu_{\theta}(x))_{ij} = \frac{e^{\nu_{\theta}(x)_{ij}}}{\sum_{j=1}^{n} e^{\nu_{\theta}(x)_{ij}}}.$$
 (Vanilla Softmax)

Many of the policies mentioned earlier can be defined in this way, such as the soft MaxWeight policy, $\nu_{\theta}(x)_i = \{\theta_j x_j \mu_{ij}\}_{j=1}^n$. This parameterization is highly flexible and $\nu_{\theta}(x)$ can be the output of a neural network. However, for a general $\nu_{\theta}(x)$, there is no guarantee that $\pi(\theta)(x)_{i,j} = 0$ if $x_j = 0$. This means that such policies may waste service capacity by allocating capacity to empty queues even when there are non-empty queues that server *i* could work on.

We propose a simple fix, which reshapes the actions produced by the policy. We refer to this as the **work-conserving softmax**,

$$\pi_{\theta}^{\mathsf{WC}}(x)_{ij} = \mathsf{softmaxWC}(\nu_{\theta}(x))_{ij} \equiv \frac{e^{\nu_{\theta}(x)_{ij}} \mathbb{1}\{x_j > 0\} \wedge \epsilon}{\sum_{l=1}^{n} e^{\nu_{\theta}(x)_{il}} \mathbb{1}\{x_l > 0\} \wedge \epsilon},$$
(WC-Softmax)

where \wedge is the minimum and ϵ is a small number to prevent division by zero when the queue lengths are all zero.

This parameterization is fully compatible with deep reinforcement learning approaches. $\nu_{\theta}(x)$ can be a neural network and critically, the work-conserving softmax preserves the differentiability



Figure 13. Multi-class queuing networks. This figure displays the architectures for the networks considered in sections 5 and 7, and have appeared in previous works (see [26, 12]). This displays the re-entrant networks with 6 job classes but we also consider networks with a larger number of job classes/queues. For both of these network architectures, the network with n job classes has n/3 servers, and each job must be processed sequentially by all servers. Each server serves 3 job classes, and some of the jobs processed by the last server are fed back into the front of the re-entrant network.

of $\pi_{\theta}^{\mathsf{WC}}(x)$ with respect to θ . As a result, REINFORCE and PATHWISE estimators can both be computed under this parameterization, since $\nabla_{\theta} \log \pi_{\theta}^{\mathsf{WC}}(x)$ and $\nabla_{\theta} \pi_{\theta}^{\mathsf{WC}}(x)$ both exist.

This simple modification delivers substantial improvements in performance. Figure 12 compares the average holding cost across policy iterations for PPO without any modifications, PPO initialized with a policy trained to imitate MaxWeight, and PPO with the work-conserving softmax. Despite its empirical success in many other reinforcement learning problems, PPO without any modifications fails to stabilize the network and incurs an exceedingly high cost. It performs much better under an initial behavioral cloning step, which achieves stability but still underperforms the $c\mu$ -rule. On the other hand, with the work-conserving softmax, even the randomly initialized policy stabilizes the network and outperforms the $c\mu$ -rule over the course of training. This illustrates that an appropriate choice of policy architecture, motivated by queuing theory, is decisive in enabling learning-based approaches to succeed. As a result, for all of the policy optimization experiments in sections 5.2, 5.3, and 7, we equip the policy parameterization with the work-conserving softmax.

7 Scheduling for Multi-Class Queuing Networks: Benchmarks

We now benchmark the performance of the policies obtained by PATHWISE policy gradient (Algorithm 1) with standard queuing policies and policies obtained using state-of-the-art model-free reinforcement learning algorithms. We consider networks displayed in Figure 13, which were briefly described in section 5.1 and appeared in previous works [26, 12]. Dai and Gluzman [26] used Criss-cross and Re-entrant-1 networks to show that PPO can outperform standard queuing policies. Bertsimas et al. [12] consider the Re-entrant-2 network, but did not include any RL baselines. We consider networks with exponential inter-arrival times and workloads in order to compare with previous results. We also consider hyper-exponential distributions to model settings with higher coefficients of variation, as has been observed in real applications [43]. The hyper-exponential distribution $X \sim \text{HyperExp}(\lambda_1, \lambda_2, p)$ is a mixture of exponential distributions:

$$X \stackrel{d}{=} Y \cdot E_1 + (1 - Y) \cdot E_2,$$

for $Y \sim \text{Bernoulli}(p)$, $E_1 \sim \text{Exp}(\lambda_1)$, $E_2 \sim \text{Exp}(\lambda_2)$, and all are drawn independently of each other. We calibrate the parameters of the hyper-exponential distribution to have the same mean as the corresponding exponential distribution, but with a 1.5x higher variance. Our empirical validation



Figure 14. Comparison of average holding cost of queuing control policies relative to $c\mu$ -rule, including PPO-WC (ours) and PATHWISE (Algorithm 1). (Left) Average holding cost for the Reentrant-1 networks with 18, 21, 24, 27, and 30 job classes. (Right) Average holding cost for the Reentrant-1 networks with hyper-exponential inter-arrival and service times with 6, 9, 12, 14, 18, 21 job classes. PATHWISE outperforms PPO-WC for large networks in particular.

goes beyond the typical settings studied in the reinforcement learning for queuing literature, and is enabled by our discrete-event simulation framework.

Table 1: Criss Cross

Noise	$c\mu$	MaxWeight	MaxPressure	Fluid	PPO-DG [26]	PPO-WC	PATHWISE	% improve
Exp	17.9 ± 0.3	17.8 ± 0.3	19.0 ± 0.3	18.2 ± 2.7	15.4 ± 0.1	15.4 ± 0.2	15.2 ± 0.4	17.8%
HyperExp	28.4 ± 0.5	28.3 ± 0.8	28.1 ± 0.8	27.1 ± 4.7	N/A	24.0 ± 0.7	22.5 ± 0.7	20.4%

We now describe the standard queuing policies considered in this section, which can all be expressed in the form

$$\pi(x) = \operatorname*{argmax}_{u \in \mathcal{U}} \sum_{i \in [m], j \in [n]} \rho_{ij}(x) u_{ij}$$

for some index $\rho \in \mathbb{R}^{m \times n}$ that differs per method.

- $c\mu$ -rule [25]: $\rho_{ij} = h_j \mu_{ij} 1\{x_j > 0\}$. Servers prioritize queues with a higher holding cost and a larger service rate.
- MaxWeight [93, 89]: $\rho_{ij}(x) = h_j \mu_{ij} x_j$. Servers prioritize queues that are longer, with a higher holding cost, and a larger service rate.
- MaxPressure [93, 27]: $\rho_{ij} = \sum_{\ell=1}^{n} \mu_{ij} R_{j\ell} h_{\ell} x_{\ell}$. MaxPressure is a modification of MaxWeight, in the sense that it takes workload externality within the network into account through the R_{jl} terms, e.g., processing a class j job may generate a new class j' job.
- Fluid [10]: The scheduling policy is based on the optimal sequence of actions in the fluid relaxation, which approximates the average evolution of stochastic queue lengths by deterministic ordinary differential equations. We aim to solve the continuous-time problem

$$\min_{\bar{u}} \int_0^T h^\top x(t) dt$$

s.t.
$$\dot{x}(t) = \lambda - R(\mu \odot \bar{u}(t)), \quad \forall t \in [0, T]$$

 $x(t) \ge 0, \quad \forall t \in [0, T]$
 $\bar{u}(t) \in \overline{\mathcal{U}}, \quad \forall t \in [0, T].$

For tractability, we discretize the problem with time increment $\Delta t > 0$ and horizon $H = T/\Delta t$, and solve as a linear program. We then set $u_k = \overline{u}(t_k)$. The linear program is re-solved periodically to improve fidelity with the original stochastic dynamics.

We next describe the deep reinforcement learning methods considered in this section:

- PPO-DG [26]: PPO is a standard model-free policy gradient method [84, 56]. Dai and Gluzman [26] implement PPO for multi-class queuing networks and show that the policies obtained outperform several standard queuing policies. Their implementation includes an initial behavioral cloning for stability and a carefully designed variance-reducing policy gradient estimation. We report the results from their paper, although in our experiments we include several problem instances not evaluated in their work.
- PPO-WC (ours): Given the empirical success of PPO with the work-conserving softmax, we use this algorithm as the main RL benchmark. For this policy, we use the same neural network architecture, hyper-parameters, and variance reduction methods as PPO-DG [26].
- PATHWISE policy gradient (Algorithm 1): Trains a neural network policy with work-conserving softmax using the PATHWISE policy gradient estimator. We use an inverse temperature of $\beta = 10$ for all experiments in this section.

While value-based methods have also been considered for queuing network control [64, 100], our focus in this work is on benchmarking policy gradient algorithms.

For the reinforcement learning policies, we train each method over 100 episodes, each consisting of N = 50,000 environment steps. Following Dai and Gluzman [26]'s implementation, PPO-WC was trained with B = 50 actors, while PATHWISE was trained only with B = 1 actor, which means that PPO-based methods used 50x more trajectories than PATHWISE. See Appendix A for more details on the training process.

To evaluate each scheduling policy, we run 100 parallel episodes starting from empty queues $x_0 = \mathbf{0}_n$ with a long horizon N to estimate the long-run average holding cost (typically N = 200,000 steps). As in previous works (e.g. [26, 12]), we consider holding costs $h = \mathbf{1}_n$ in which case the holding cost is equivalent to the total queue length. To reiterate, for each policy π we estimate the following quantity

$$J_N(\pi) = \mathbb{E}\left[\frac{1}{t_N} \sum_{k=0}^{N-1} (\mathbf{1}_n^{\top} x_k) \tau_{k+1}^*\right] = \mathbb{E}\left[\frac{1}{t_N} \int_0^{t_N} \sum_{i=1}^n x_i(t) dt\right]$$
(18)

where t_N is the time of the Nth event. We measure the standard deviation across the 100 episodes to form 95% confidence intervals. For the reinforcement learning policies, we report the average holding cost for the best policy encountered during training.

Tables 1-5 display the results of our benchmarking across the problem instances discussed before. The column '% improve' records the relative reduction in holding cost achieved by PATHWISE over the best standard queuing policy (either $c\mu$, MaxWeight, MaxPressure, or Fluid). Our main observations on the relative performance of the standard policies and policies obtained from reinforcement learning methods are summarized as follows.

Classes	$c\mu$	MaxWeight	MaxPressure	Fluid	PPO-DG [26]	PPO-WC	PATHWISE	% improve
6	17.4 ± 0.4	17.5 ± 0.4	18.8 ± 0.5	16.8 ± 4.3	14.1 ± 0.2	13.6 ± 0.4	14.9 ± 0.5	14.3%
9	23.3 ± 0.6	26.1 ± 0.5	24.2 ± 0.6	27.7 ± 4.4	23.3 ± 0.3	22.6 ± 0.4	22.0 ± 0.6	4.3%
12	33.0 ± 0.8	34.0 ± 1.0	35.1 ± 0.9	40.6 ± 4.8	32.2 ± 0.6	29.7 ± 0.4	30.7 ± 0.7	7.0%
15	40.2 ± 1.3	43.6 ± 1.1	42.2 ± 1.3	49.8 ± 5.1	39.3 ± 0.6	38.7 ± 0.4	36.2 ± 0.8	9.9%
18	48.5 ± 1.0	51.0 ± 1.4	52.4 ± 1.6	54.5 ± 4.2	51.4 ± 1.0	47.5 ± 0.5	45.7 ± 0.7	5.7%
21	55.2 ± 1.1	59.5 ± 1.6	56.0 ± 1.6	63.7 ± 6.7	55.1 ± 1.8	56.3 ± 0.8	52.8 ± 1.2	4.3%
24	64.9 ± 1.4	69.4 ± 1.6	66.6 ± 2.2	74.0 ± 6.7	N/A	65.8 ± 0.6	60.2 ± 0.9	6.2%
27	71.1 ± 1.5	77.7 ± 1.9	72.0 ± 2.5	83.1 ± 7.1	N/A	75.8 ± 0.7	67.7 ± 1.3	3.0%
30	87.7 ± 2.5	84.5 ± 2.1	106.8 ± 2.5	93.6 ± 8.0	N/A	83.1 ± 0.7	77.8 ± 1.3	7.9%

Table 2: Re-entrant-1 (Exp)

 Table 3: Re-entrant-1 (HyperExp)

Classes	$c\mu$	MaxWeight	MaxPressure	Fluid	PPO-WC	PATHWISE	% improve
6	37.8 ± 1.3	39.2 ± 1.2	43.8 ± 1.8	43.6 ± 7.5	29.9 ± 0.7	32.2 ± 1.4	6.9%
9	50.2 ± 1.9	55.5 ± 2.2	68.7 ± 2.7	59.2 ± 8.2	47.5 ± 0.8	45.6 ± 1.6	9.2%
12	70.0 ± 2.5	72.3 ± 2.8	89.4 ± 3.6	75.6 ± 15.3	64.4 ± 1.2	61.1 ± 2.5	12.7%
15	81.7 ± 4.0	91.0 ± 3.7	112.0 ± 4.9	97.0 ± 12.9	81.8 ± 1.1	76.9 ± 2.5	5.9%
18	101.1 ± 4.7	103.9 ± 3.9	126.7 ± 6.2	111.2 ± 14.4	99.8 ± 1.5	91.5 ± 2.5	9.4%
21	116.3 ± 4.6	123.3 ± 3.9	152.3 ± 6.6	151.0 ± 21.3	118.2 ± 2.0	108.0 ± 3.2	7.2%

- PPO-WC is a strong reinforcement-learning benchmark. PPO with our proposed workconserving softmax is able to efficiently find policies that outperform standard policies across all problem instances, as well as the PPO featured in [26] (under the same policy network and hyper-parameters). This illustrates that simply ensuring work-conservation is a powerful inductive bias that delivers stability.
- PATHWISE policy gradient outperforms PPO-WC in larger networks, using 50x less data. We observe that PPO-WC and PATHWISE achieve similar performances when the number of classes is small. However, when the number of job classes gets larger, PATHWISE consistently outperforms PPO-WC, as seen in Figure 14. This is likely due to the sample efficiency gained by the PATHWISE gradient, enabling the algorithm to find better policies with less data.
- PATHWISE achieves large performance gains over PPO-WC for higher-variance problem instances, using 50x less data. We observe that for the Re-entrant-1 networks with hyper-exponential noise, the reduction of holding cost of PATHWISE relative to PPO-WC is often equivalent or even larger than the cost reduction of PPO-WC relative to the *cµ*-rule. This illustrates that even among optimized RL policies, there can be significant per-

formance differences for difficult problem instances, and the sample efficiency of PATHWISE is particularly useful in noisier environments.

• While standard queuing methods work well, RL methods meaningfully improve performance in hard instances. We observe in the '% improve' column that PATHWISE achieves a 3-20% improvement over the best standard queuing policy in each setting.

Altogether, these results illustrate that policy gradient with PATHWISE gradient estimator and work-conserving softmax policy architecture can learn effective queuing network control policies with substantially less data than model-free policy gradient algorithms for large networks with high-variance event times, mirroring real-world systems. In particular, the improved sample efficiency of the PATHWISE gradient estimator and the stability brought by work-conserving softmax policy architecture are the keys to enabling learning in large-scale systems with realistic data requirements.

	Table 4: Re-entrant-2 (Exp)							
Classes	$c\mu$	MaxWeight	MaxPressure	Fluid	PPO-WC	PATHWISE	% improve	
6	18.8 ± 0.5	17.4 ± 0.4	24.5 ± 0.7	18.6 ± 2.8	13.7 ± 0.2	14.7 ± 0.6	21.9%	
9	24.2 ± 0.6	25.8 ± 0.7	31.5 ± 1.0	26.7 ± 3.7	22.1 ± 0.3	21.6 ± 0.6	10.7%	
12	35.1 ± 0.9	34.0 ± 1.1	40.3 ± 1.4	35.7 ± 4.6	29.9 ± 0.5	29.8 ± 0.7	15.1%	
15	44.8 ± 1.3	43.8 ± 1.5	50.1 ± 1.6	44.0 ± 5.3	38.0 ± 0.5	36.2 ± 1.0	19.2%	
18	52.4 ± 1.6	48.7 ± 1.3	58.2 ± 2.1	55.2 ± 5.8	46.8 ± 0.5	45.6 ± 0.8	13.0%	
21	56.0 ± 1.6	57.5 ± 1.8	71.3 ± 2.9	62.2 ± 7.5	55.5 ± 0.6	51.4 ± 1.1	8.2%	
24	66.6 ± 2.2	69.0 ± 1.7	76.0 ± 3.0	70.8 ± 7.7	63.2 ± 0.7	59.8 ± 1.4	10.2%	
27	72.0 ± 2.5	75.9 ± 2.1	84.9 ± 3.2	82.9 ± 9.6	70.3 ± 0.9	68.4 ± 1.6	5.0%	
30	80.6 ± 2.7	83.8 ± 1.9	90.6 ± 3.2	91.2 ± 11.3	80.4 ± 0.8	75.5 ± 1.8	6.3%	
Table 5: Re-entrant-2 (HyperExp)								
Classes	$c\mu$	MaxWeight	MaxPressure	Fluid	PPO-WC	PATHWISE	% improve	
6	39.4 ± 1.4	39.1 ± 1.9	58.6 ± 2.3	39.8 ± 7.3	30.7 ± 0.8	30.9 ± 1.4	21.6%	
9	52.4 ± 2.1	58 ± 2.7	67.1 ± 3.0	55.5 ± 9.4	45.7 ± 0.8	43.4 ± 1.4	17.2%	
12	70.9 ± 2.7	77.0 ± 3.7	93.0 ± 4.0	72.1 ± 14.4	61.1 ± 1.3	58.9 ± 2.3	16.9%	
15	81.5 ± 2.7	90.5 ± 3.9	109.0 ± 5.4	84.2 ± 18.1	78.0 ± 1.7	73.8 ± 3.6	9.4%	
18	104.3 ± 5.1	103.8 ± 4.3	123.6 ± 5.3	99.5 ± 15.5	93.8 ± 1.3	92.1 ± 2.7	11.7%	
21	116.6 ± 6.0	117.9 ± 5.1	135.6 ± 6.0	118.9 ± 16.6	110.5 ± 1.6	104.2 ± 2.9	10.6%	

8 Why is REINFORCE Sample-Inefficient? A Theoretical Case Study for the M/M/1 Queue

In this section, we provide a theoretical case study explaining how PATHWISE gradients are able to learn more from a single observed trajectory compared to REINFORCE. We focus on the special case of the M/M/1 queue to explain how REINFORCE and its actor-critic variants utilizing baselines/advantages suffer from sample inefficiency. Although we are only able to analyze a substantially simpler setting than the control problems in general multi-class queueing networks we are interested in, our theoretical results illustrate the essential statistical benefits of pathwise gradient estimators. We highlight that while REINFORCE applies to virtually any setting by relying on random exploration, it fundamentally struggles to assign credit to actions, especially in noisy environments. PATHWISE, on the other hand, is much better at assigning credit to actions. This allows us to crystallize why we see such a large improvement in sample efficiency in sections 5 and 7. We also believe our result may be of broader interest in reinforcement learning, because it illustrates a practically relevant instance where REINFORCE, even with an optimal baseline, is provably sub-optimal.

We consider the M/M/1 queue with a fixed arrival rate λ under service rate control $u = \mu > \lambda$. This setting permits an analytic approach to showing that the REINFORCE estimator has a suboptimally large variance, particularly for congested systems when $\rho = \lambda/\mu \rightarrow 1$. On the other hand, the PATHWISE estimator achieves an order of magnitude improvement in estimation efficiency. In the M/M/1 queue setting, the PATHWISE approach for general queuing networks reduces to IPA based on Lindley recursion.

We consider a simple service-rate control problem where the cost is the steady-state average queue length in the M/M/1 queue

$$Q(\mu) := \mathbb{E}_{\infty}[x(t)] = \frac{\lambda}{\mu - \lambda} = \frac{\rho}{1 - \rho}.$$

Given that the service rate is continuous, it is natural to consider a policy that randomizes over $[\mu - h, \mu]$. One such option is a Beta-distributed policy $\pi_{\theta} : A = \mu - hY$, where $Y \sim \text{Beta}(\theta, 1)$ and h > 0. As $\theta \to \infty$, the policy frequently sets service rates close to μ and as $\theta \to 0$, it concentrates more probability mass on service rates close to $\mu - h$. The task is to estimate the following policy gradient

$$\nabla_{\theta} J(\theta) = \nabla \mathbb{E}_{A \sim \pi_{\theta}} [Q(A)].$$

The REINFORCE estimator of $\nabla_{\theta} J(\theta)$ involves sampling a random service rate from the policy π_{θ} , and then estimating the steady-state queue length from a trajectory under that service rate. For a trajectory with N steps, we denote the steady-state queue length estimator as $\widehat{Q}_N(A)$. Then, the REINFORCE gradient estimator takes the form

$$\widehat{\nabla}^{\mathsf{R}} J_N(\theta; \xi_{1:N}) = \widehat{Q}_N(\mu - hY) \nabla_\theta \log \pi_\theta(Y) = \widehat{Q}_N(\mu - hY) \left(\log Y + \frac{1}{\theta} \right).$$
(19)

The standard estimator for the steady-state queue length is simply the queue length averaged over a sample path: $\hat{Q}_N(a) = \frac{1}{N} \sum_{k=1}^N x_k \tau_{k+1}^*$ when A = a. As long as $\mu - h > \lambda$, it is known that as $N \to \infty$, $\hat{Q}_N(a) \to Q(a)$, which implies that $\hat{\nabla}^{\mathsf{R}} J_N(\theta; \xi_{1:N}) \to \nabla J(\theta)$.

On the other hand, the PATHWISE estimator utilizes the structure of the single server queue. First, by inverse transform sampling, $Y \stackrel{d}{=} F_{\theta}^{-1}(\omega)$, where $\omega \sim \text{Uniform}(0,1)$ and $F_{\theta}^{-1}(\omega) = \omega^{1/\theta}$. Then, we can substitute $A = \mu - h\omega^{1/\theta}$. Since $Q(\mu)$ is differentiable and the derivative is integrable, we can change the order of differentiation and integration

$$\nabla_{\theta} J(\theta) = \nabla_{\theta} \mathbb{E}_{\omega} [Q(\mu - h\omega^{1/\theta})] = -\mathbb{E}_{\omega} [\nabla Q(\mu - h\omega^{1/\theta}) \cdot h\nabla_{\theta} \omega^{1/\theta}].$$

The preceding display involves the gradient of the steady-state queue-length $Q(\mu)$ with respect to the service rate μ , i.e., $\nabla Q(\mu)$.

For the M/M/1 queue, there are consistent sample-path estimators of $\nabla Q(\mu)$. One such estimator uses the fact that by Little's law $Q(\mu) = \mathbb{E}_{\infty}[x(t)] = \lambda \mathbb{E}_{\infty}[w(t)] =: \lambda W(\mu)$ where $W(\mu)$ is the steady-state waiting time. The waiting time process W_i , which denotes the waiting time of the *i*th job arriving to the system, has the following dynamics, known as the Lindley recursion:

$$W_{i+1} = \left(W_i - T_{i+1} + \frac{S_i}{\mu}\right)^+,$$
(20)

where $T_{i+1} \stackrel{\text{iid}}{\sim} \text{Exp}(1)$ is the inter-arrival time of between the *i*th job and (i + 1)th job, and $S_i \stackrel{\text{iid}}{\sim} \text{Exp}(1)$ is the workload of the *i*th job. Crucially, this stochastic recursion specifies how the service rate affects the waiting time along the sample path, which enables one to derive a pathwise derivative via the recursion:

$$\widehat{\nabla}W_{i+1} = \left(-\frac{S_k}{\mu^2} + \widehat{\nabla}W_i\right) \mathbf{1} \left\{W_{i+1} > 0\right\},\,$$

where $\widehat{\nabla}W_i$ together with W_i form a Markov chain following the above recursion. By averaging this gradient across jobs and using Little's law, we have the following gradient estimator for ∇Q , which we denote as $\widehat{\nabla}Q_N(\mu)$:

$$\widehat{\nabla}Q_N(\mu) = \lambda \frac{1}{L_N} \sum_{i=1}^{L_N} \widehat{\nabla}W_i,$$

where L_N is the number of arrivals that occur during a sample path with N events. Using this, the PATHWISE policy gradient estimator (a.k.a. IPA estimator) is

$$\widehat{\nabla}_{\theta} J_N(\theta; \xi_{1:N}) = h \cdot \widehat{\nabla} Q_N(\mu - hY) \left(\frac{1}{\theta} Y \log Y\right).$$
(21)

As long as $\mu - h > \lambda$, it has been established that $\widehat{\nabla}Q_N(\mu)$ is asymptotically unbiased, i.e., as $N \to \infty$, $\mathbb{E}[\widehat{\nabla}Q_N(\mu)] \to \nabla Q(\mu)$, which implies $\mathbb{E}[\widehat{\nabla}_{\theta}J_N(\theta;\xi_{1:N})] \to \nabla J(\theta)$.

Since both the REINFORCE and PATHWISE gradient estimators give an asymptotically unbiased estimation of $\nabla J(\theta)$, we compare them based on their variances, which determines how many samples are needed to reliably estimate the gradient. Although the variances of the estimators are not precisely known for a finite N, \hat{Q}_N and $\hat{\nabla}Q_N$ both satisfy the central limit theorem (CLT) with explicitly characterized asymptotic variances, which we denote as $\operatorname{Var}_{\infty}(\hat{Q})$ and $\operatorname{Var}_{\infty}(\hat{\nabla}Q)$ respectively. This implies that the variance of \hat{Q}_N is approximately $\operatorname{Var}_{\infty}(\hat{Q})/N$. We define $\operatorname{Var}_{\infty}(\hat{\nabla}J_N(\theta;\xi_{1:N}))$ to be the variance of the gradient estimator when we approximate the variance of \hat{Q}_N and $\hat{\nabla}Q_N$ using $\operatorname{Var}_{\infty}(\hat{Q})/N$ and $\operatorname{Var}_{\infty}(\hat{\nabla}Q)/N$ respectively.

It is worth reiterating that the REINFORCE estimator only required an estimate of $\cot Q(\mu)$, which does not require any domain knowledge, whereas the PATHWISE gradient required an estimate of $\nabla Q(\mu)$ which requires a detailed understanding of how the service rate affected the sample path dynamics. Utilizing this structural information can greatly improve the efficiency of gradient estimation. Since $\widehat{\nabla}_{\theta} J_N(\theta; \xi_{1:N}) = O(h)$ almost surely, we have $\operatorname{Var}(\widehat{\nabla}_{\theta} J_N(\theta; \xi_{1:N})) = O(h^2)$. On the other hand, the variance of the REINFORCE estimator can be very large even if h is small. To highlight this, consider the extreme case where h = 0 for which the policy gradient $\nabla J(\theta)$ is obviously zero since the policy deterministically sets the service rate to μ regardless of θ . Strikingly, REINFORCE does not have zero variance in this case:

Observation 1. If h = 0, then the variance of the estimators are

$$\operatorname{Var}_{\infty}\left(\widehat{\nabla}J_{N}(\theta;\xi_{1:N})\right) = 0, \quad \operatorname{Var}_{\infty}\left(\widehat{\nabla}^{\mathsf{R}}J_{N}(\theta;\xi_{1:N})\right) = \Theta\left(N^{-1}(1-\rho)^{-4}\right)$$

Note that even when h = 0, the variance of the REINFORCE estimator can be quite high if the queue is congested, i.e. ρ is close to 1, while the pathwise estimator gives the correct estimate of zero with zero variance.

For non-trivial values of h, we focus on the so-called 'heavy-traffic' asymptotic regime with $\rho = \lambda/\mu \to 1$, which is of major theoretical and practical interest in the study of queues. Estimating steady-state quantities becomes harder as the queue is more congested, so $(1 - \rho)^{-1}$ emerges as a key scaling term in the variance. We set h such that $\frac{h}{\mu-\lambda} < c$ and $\frac{h}{\mu-\lambda} \to c \in (0,1)$ as $\rho \to 1$. This resembles the square-root heavy-traffic regime for capacity planning, where the service rate is set to be $\lambda + \beta \sqrt{\lambda}$ for some $\beta > 0$, and one considers the limit as $\lambda \to \infty$. In this case, if one were choosing a policy over the square-root capacity rules $A \in [\lambda + a\sqrt{\lambda}, \lambda + b\sqrt{\lambda}]$ for some b > a > 0, this is equivalent to setting $\mu = \lambda + b\sqrt{\lambda}$ and $h = (b - a)\sqrt{\lambda} = O(\sqrt{\lambda})$. Note that if c = 0, the gradient is zero (identical to Observation 1), and if $c \ge 1$, the queue with service rate $\mu - h$ is unstable.

Within this regime, we have the following comparison between the gradient estimators, which utilizes recent results concerning the asymptotic variance of $\widehat{\nabla}Q_N(\mu)$ [54].

Theorem 2. Suppose $h = c(\mu - \lambda)$ for $c \in (0, 1)$ as $\rho \to 1$. Under this scaling $\nabla J(\theta) \sim (1 - \rho)^{-1}$, and

$$\operatorname{Var}_{\infty}\left(\widehat{\nabla}J_{N}(\theta;\xi_{1:N})\right) = O\left(\underbrace{N^{-1}\left(1-\rho\right)^{-3}}_{estimation\ noise} + \underbrace{\left(1-\rho\right)^{-2}}_{policy\ randomization}\right)$$
(22)

$$\operatorname{Var}_{\infty}\left(\widehat{\nabla}^{\mathsf{R}}J_{N}(\theta;\xi_{1:N})\right) = \Theta\left(N^{-1}\left(1-\rho\right)^{-4} + (1-\rho)^{-2}\right)$$
(23)

See Appendix B.2 for the proof.

Overall, the PATHWISE estimator is much more sample efficient than the REINFORCE estimator as $\rho \to 1$, with the variance scaling as $(1 - \rho)^{-3}$ compared to $(1 - \rho)^{-4}$. The first terms in (22) and (23) represent the variance occurring from the Monte Carlo estimation and becomes smaller if one generates a longer sample path (larger N), and it scales as N^{-1} . The second terms are the variance resulting from randomness in the service rate induced by the policy.

Theorem 2 illustrates that large improvements in statistical efficiency can be achieved by leveraging the structure of the system dynamics. An existing strategy for incorporating domain knowledge in REINFORCE is to subtract a baseline b from the cost, which preserves un-biasedness:

$$\widehat{\nabla}^{\mathsf{RB}} J_N(\theta, \xi_{1:N}) = \left(\widehat{Q}_N(\mu - hY) - b\right) \left(\log Y + \frac{1}{\theta}\right).$$
(24)

In this case, one can characterize the optimal variance-reducing baseline in closed form if one has

knowledge of the true cost $Q(\mu)$. Under $h = c(\mu - \lambda)$,

$$b^* = \frac{\mathbb{E}[Q(h-hY)\nabla_{\theta}\log\pi_{\theta}(Y)^2]}{\mathbb{E}[\nabla_{\theta}\log\pi_{\theta}(Y)^2]}$$
$$= \frac{\lambda}{\mu-\lambda} \left[F_1^2(1,\theta,1+\theta,c) - 2\theta^2 \Phi(c,2,\theta) + 2c\theta^3 \Phi(c,3,1+\theta)\right]$$
$$= O((1-\rho)^{-1})$$

where F_1^2 is the hypergeometric 2F1 function and Φ is the Lerch Φ transcendental. The optimal baseline b^* is of the same order as $Q(\mu)$ as $\rho \to 1$.

Corollary 2. Consider the REINFORCE estimator with the optimal baseline b^* . As $\rho \to 1$, the variance of the estimator scales as

$$\operatorname{Var}_{\infty}\left(\widehat{\nabla}^{\mathsf{RB}}J_{N}(\theta;\xi_{1:N})\right) = \Theta\left(N^{-1}\left(1-\rho\right)^{-4} + (1-\rho)^{-2}\right)$$

The proof of Corollary 2 is provided in Appendix B.3. Simply, since the optimal baseline is a deterministic input, it is unable to improve upon the $(1-\rho)^{-4}$ dependence on ρ , which is driven by the statistical properties of \hat{Q}_N . This illustrates that the pathwise gradient estimator can offer an order of magnitude improvement in sample efficiency than the REINFORCE estimator even with an optimized baseline that requires knowledge of the true cost function (and thus precludes the need to estimate the cost in the first place).

Intuitively, the REINFORCE estimator is inefficient because it is unable to leverage the fact that $Q(\mu) \approx Q(\mu + \epsilon)$ when ϵ is small. After all, generic MDPs do not have such a structure; a slight change in the action can result in vastly different outcomes. The REINFORCE estimator cannot use the estimate of $\hat{Q}_N(\mu)$ to say anything about $Q(\mu + \epsilon)$, and must draw a new sample path to estimate $Q(\mu + \epsilon)$. Meanwhile, using a *single* sample path, the pathwise estimator can obtain an estimate for $Q(\mu + \epsilon)$ when ϵ is small enough via $\hat{Q}_N(\mu + \epsilon) \approx \hat{Q}_N(\mu) + \epsilon \hat{\nabla} Q_N(\mu)$. In this sense, the pathwise estimator can be seen as a *infinitesimal counterfactual* of the outcome under alternative—but similar—actions.

Even though we only study a single server queue here, we believe the key observations may apply more broadly.

- Higher congestion (ρ → 1) makes it more challenging to estimate the performance of queueing networks based on the sample path. This applies to both gradient estimators and baselines that could be used to reduce variance.
- It is important to reliably estimate the effects of small changes in the policy, as large changes can potentially cause instability. Pathwise gradient estimators provide a promising way to achieve this. For general networks with known dynamics, their dynamics are often not differentiable, which requires the development of the PATHWISE estimator.

9 Conclusion

In this work, we introduce a new framework for policy optimization in queuing network control. This framework uses a novel approach for gradient estimation in discrete-event dynamical systems. Our proposed PATHWISE policy gradient estimator is observed to be orders of magnitude more efficient than model-free RL alternatives such as REINFORCE across an array of carefully designed empirical

experiments. In addition, we introduce a new policy architecture, which drastically improves stability while maintaining the flexibility of neural network policies. Altogether, these illustrate how structural knowledge of queuing networks can be leveraged to accelerate reinforcement learning for queuing control problems.

We next discuss some potential extensions of our approach:

- We consider policies with preemption. Our proposed method can also handle non-preemptive policies by keeping track of the occupied servers as part of the state.
- We focus on scheduling and admission control problems in queuing network satisfying Assumptions A, but the algorithmic ideas can be extended to more general queuing networks by utilizing a larger state space that contains the residual workloads of *all* jobs in the network, rather than only the top-of-queue jobs as is done in this work. A higher dimensional state descriptor is required for more general networks as multiple jobs in the same queue can be served simultaneously.
- Beyond queuing network control, our methodology can be extended to control problems in other discrete-event dynamical systems. More explicitly, our methodology can handle systems that involve a state update of the form $x_{k+1} = g(x_k, e_{k+1})$ where g is a differentiable function and e_{k+1} is the selected event. Recall that in this work, the state update is linear in x_k and e_{k+1} : $x_{k+1} = x_k + De_{k+1}$. We also require that e_{k+1} is differentiable almost surely in the action u_k .

References

- Z. Aksin, M. Armony, and V. Mehrotra. The modern call center: A multi-disciplinary perspective on operations management research. *Production and operations management*, 16(6): 665–688, 2007.
- [2] M. Alvo, D. Russo, and Y. Kanoria. Neural inventory control in networks via hindsight differentiable policy optimization. arXiv preprint arXiv:2306.11246, 2023.
- [3] P. Andelfinger. Differentiable agent-based simulation for gradient-guided simulation-based optimization. In Proceedings of the 2021 ACM SIGSIM Conference on Principles of Advanced Discrete Simulation, pages 27–38, 2021.
- [4] P. Andelfinger. Towards differentiable agent-based simulation. ACM Transactions on Modeling and Computer Simulation, 32(4):1–26, 2023.
- [5] M. Armony and N. Bambos. Queueing dynamics and maximal throughput scheduling in switched processing systems. *Queueing systems*, 44:209–252, 2003.
- [6] M. Armony, S. Israelit, A. Mandelbaum, Y. N. Marmor, Y. Tseytlin, and G. B. Yom-Tov. On patient flow in hospitals: A data-based queueing-science perspective. *Stochastic systems*, 5(1):146–194, 2015.
- [7] F. Avram, D. Bertsimas, and M. Ricard. An optimal control approach to optimization of multiclass queueing networks. In *Proceedings of Workshop on Queueing Networks*, 1994.

- [8] L. Baird and A. Moore. Gradient descent for general reinforcement learning. Advances in neural information processing systems, 11, 1998.
- [9] S. Banerjee, D. Freund, and T. Lykouris. Pricing and optimization in shared vehicle systems: An approximation framework. *Operations Research*, 70(3):1783–1805, 2022.
- [10] N. Bäuerle. Asymptotic optimality of tracking policies in stochastic networks. The Annals of Applied Probability, 10(4):1065–1083, 2000.
- [11] Y. Bengio, N. Léonard, and A. Courville. Estimating or propagating gradients through stochastic neurons for conditional computation. arXiv preprint arXiv:1308.3432, 2013.
- [12] D. Bertsimas, E. Nasrabadi, and I. C. Paschalidis. Robust fluid processing networks. *IEEE Transactions on Automatic Control*, 60(3):715–728, 2014.
- [13] J. Bradbury, R. Frostig, P. Hawkins, M. J. Johnson, C. Leary, D. Maclaurin, G. Necula, A. Paszke, J. VanderPlas, S. Wanderman-Milne, and Q. Zhang. JAX: composable transformations of Python+NumPy programs, 2018. URL http://github.com/google/jax.
- [14] X.-R. Cao. Convergence of parameter sensitivity estimates in a stochastic experiment. IEEE Transactions on Automatic Control, 30(9):845–853, 1985.
- [15] X.-R. Cao. First-order perturbation analysis of a simple multi-class finite source queue. *Performance Evaluation*, 7(1):31–41, 1987.
- [16] C. G. Cassandras, Y. Wardi, B. Melamed, G. Sun, and C. G. Panayiotou. Perturbation analysis for online control and optimization of stochastic fluid models. *IEEE Transactions* on Automatic Control, 47(8):1234–1248, 2002.
- [17] C. G. Cassandras, G. Sun, C. G. Panayiotou, and Y. Wardi. Perturbation analysis and control of two-class stochastic fluid models for communication networks. *IEEE Transactions* on Automatic Control, 48(5):770–782, 2003.
- [18] H. Chen and A. Mandelbaum. Discrete flow networks: Bottleneck analysis and fluid approximations. *Mathematics of operations research*, 16(2):408–446, 1991.
- [19] H. Chen and A. Mandelbaum. Hierarchical modeling of stochastic networks, part i: Fluid models. In *Stochastic modeling and analysis of manufacturing systems*, pages 47–105. Springer, 1994.
- [20] H. Chen and D. D. Yao. Dynamic scheduling of a multiclass fluid network. Operations Research, 41(6):1104–1115, 1993.
- [21] N. Chr. Individual and social optimization in a multiserver queue with a general cost-benefit structure. *Econometrica: Journal of the Econometric Society*, pages 515–528, 1972.
- [22] E. B. Çil, E. L. Örmeci, and F. Karaesmen. Effects of system parameters on the optimal policy structure in a class of queueing control problems. *Queueing Systems*, 61:273–304, 2009.
- [23] A. Cohen, V. Subramanian, and Y. Zhang. Learning-based optimal admission control in a single-server queuing system. *Stochastic Systems*, 14(1):69–107, 2024.

- [24] R. Cont, S. Stoikov, and R. Talreja. A stochastic model for order book dynamics. Operations research, 58(3):549–563, 2010.
- [25] D. Cox and W. Smith. Queues. Methuen, London, 5 edition, 1961.
- [26] J. G. Dai and M. Gluzman. Queueing network controls via deep reinforcement learning. Stochastic Systems, 12(1):30–67, 2022.
- [27] J. G. Dai and W. Lin. Maximum pressure policies in stochastic processing networks. Operations Research, 53(2):197–218, 2005.
- [28] J. G. Dai and S. P. Meyn. Stability and convergence of moments for multiclass queueing networks via fluid limit models. *IEEE Transactions on Automatic Control*, 40(11):1889–1904, 1995.
- [29] D. J. Eckman and S. G. Henderson. Biased gradient estimators in simulation optimization. In 2020 Winter Simulation Conference (WSC), pages 2935–2946, 2020. doi: 10.1109/WSC48552. 2020.9383938.
- [30] L. Engstrom, A. Ilyas, S. Santurkar, D. Tsipras, F. Janoos, L. Rudolph, and A. Madry. Implementation matters in deep policy gradients: A case study on ppo and trpo. arXiv preprint arXiv:2005.12729, 2020.
- [31] M. C. FU and S. D. HILL. Optimization of discrete event systems via simultaneous perturbation stochastic approximation. *IIE transactions*, 29(3):233-243, 1997.
- [32] M. C. Fu and J.-Q. Hu. Smoothed perturbation analysis derivative estimation for markov chains. Operations Research Letters, 15(5):241–251, 1994.
- [33] M. C. Fu and J.-Q. Hu. Conditional Monte Carlo: Gradient estimation and optimization applications, volume 392. Springer Science & Business Media, 2012.
- [34] A. P. Ghosh and A. P. Weerasinghe. Optimal buffer size for a stochastic processing network in heavy traffic. *Queueing Systems*, 55(3):147–159, 2007.
- [35] P. Glasserman. *Gradient estimation via perturbation analysis*, volume 116. Springer Science & Business Media, 1990.
- [36] P. Glasserman. Derivative estimates from simulation of continuous-time markov chains. Operations Research, 40(2):292–308, 1992.
- [37] P. Glasserman. Monte Carlo methods in financial engineering, volume 53. Springer, 2004.
- [38] P. Glasserman and D. D. Yao. Some guidelines and guarantees for common random numbers. Management Science, 38(6):884–908, 1992.
- [39] P. W. Glynn. Likelilood ratio gradient estimation: an overview. In Proceedings of the 19th Winter conference on simulation, pages 366–375, 1987.
- [40] P. W. GLYNN. A gsmp formalism for discrete event systems. PROCEEDINGS OF THE IEEE, 77(1), 1989.

- [41] P. W. Glynn. Optimization of stochastic systems via simulation. In Proceedings of the 21st conference on Winter simulation, pages 90–105, 1989.
- [42] P. W. Glynn. Likelihood ratio gradient estimation for stochastic systems. Communications of the ACM, 33(10):75–84, 1990.
- [43] L. V. Green, P. J. Kolesar, and W. Whitt. Coping with time-varying demand when setting staffing requirements for a service system. *Production and Operations Management*, 16(1): 13–39, 2007.
- [44] E. Greensmith, P. L. Bartlett, and J. Baxter. Variance reduction techniques for gradient estimates in reinforcement learning. *Journal of Machine Learning Research*, 5(9), 2004.
- [45] M. Harchol-Balter. Performance modeling and design of computer systems: queueing theory in action. Cambridge University Press, 2013.
- [46] J. M. Harrison. The bigstep approach to flow management in stochastic processing networks. Stochastic Networks: Theory and Applications, 4(147-186):4, 1996.
- [47] J. M. Harrison. Heavy traffic analysis of a system with parallel servers: asymptotic optimality of discrete-review policies. The Annals of Applied Probability, 8(3):822–848, 1998.
- [48] J. M. Harrison and L. M. Wein. Scheduling networks of queues: heavy traffic analysis of a simple open network. *Queueing Systems*, 5:265–279, 1989.
- [49] J. M. Harrison and L. M. Wein. Scheduling networks of queues: Heavy traffic analysis of a two-station closed network. *Operations research*, 38(6):1052–1064, 1990.
- [50] E. Heiden, D. Millard, E. Coumans, Y. Sheng, and G. S. Sukhatme. Neuralsim: Augmenting differentiable simulators with neural networks. In 2021 IEEE International Conference on Robotics and Automation (ICRA), pages 9474–9481. IEEE, 2021.
- [51] Y. Ho, M. Eyler, and T. Chien. A new approach to determine parameter sensitivities of transfer lines. *Management science*, 29(6):700–714, 1983.
- [52] Y.-C. Ho, X. Cao, and C. Cassandras. Infinitesimal and finite perturbation analysis for queueing networks. Automatica, 19(4):439–445, 1983.
- [53] T. A. Howell, S. Le Cleac'h, J. Z. Kolter, M. Schwager, and Z. Manchester. Dojo: A differentiable simulator for robotics. arXiv preprint arXiv:2203.00806, 9(2):4, 2022.
- [54] J.-Q. Hu and T. Lian. On comparison of steady-state infinitesimal perturbation analysis and likelihood ratio derivative estimates. *Discrete Event Dynamic Systems*, 33(2):95–104, 2023.
- [55] Y. Hu, L. Anderson, T.-M. Li, Q. Sun, N. Carr, J. Ragan-Kelley, and F. Durand. Diffaichi: Differentiable programming for physical simulation. arXiv preprint arXiv:1910.00935, 2019.
- [56] S. Huang, R. F. J. Dossa, A. Raffin, A. Kanervisto, and W. Wang. The 37 implementation details of proximal policy optimization. *The ICLR Blog Track 2023*, 2022.

- [57] S. Huang, R. F. J. Dossa, C. Ye, J. Braga, D. Chakraborty, K. Mehta, and J. G. AraAšjo. Cleanrl: High-quality single-file implementations of deep reinforcement learning algorithms. *Journal of Machine Learning Research*, 23(274):1–18, 2022.
- [58] A. Ilyas, L. Engstrom, S. Santurkar, D. Tsipras, F. Janoos, L. Rudolph, and A. Madry. A closer look at deep policy gradients. arXiv preprint arXiv:1811.02553, 2018.
- [59] E. Jang, S. Gu, and B. Poole. Categorical reparameterization with gumbel-softmax. arXiv preprint arXiv:1611.01144, 2016.
- [60] M. E. Johnson and J. Jackman. Infinitesimal perturbation analysis: a tool for simulation. Journal of the Operational Research Society, 40(3):243–254, 1989.
- [61] D. P. Kingma and M. Welling. Auto-encoding variational bayes. arXiv preprint arXiv:1312.6114, 2013.
- [62] Y. L. Koçağa and A. R. Ward. Admission control for a multi-server queue with abandonment. Queueing Systems, 65:275–323, 2010.
- [63] P. L'Ecuyer and P. W. Glynn. Stochastic optimization by simulation: Convergence proofs for the gi/g/1 queue in steady-state. *Management Science*, 40(11):1562–1578, 1994.
- [64] B. Liu, Q. Xie, and E. Modiano. Rl-qn: A reinforcement learning framework for optimal control of queueing systems. ACM Transactions on Modeling and Performance Evaluation of Computing Systems, 7(1):1–35, 2022.
- [65] C. J. Maddison, A. Mnih, and Y. W. Teh. The concrete distribution: A continuous relaxation of discrete random variables. In *International Conference on Learning Representations (ICLR* 2016), 2016.
- [66] D. Madeka, K. Torkkola, C. Eisenach, A. Luo, D. P. Foster, and S. M. Kakade. Deep inventory management. arXiv preprint arXiv:2210.03137, 2022.
- [67] C. Maglaras. Discrete-review policies for scheduling stochastic networks: Trajectory tracking and fluid-scale asymptotic optimality. *The Annals of Applied Probability*, 10(3):897–929, 2000.
- [68] A. Mandelbaum and A. L. Stolyar. Scheduling flexible servers with convex delay costs: Heavy-traffic optimality of the generalized $c\mu$ -rule. Operations Research, 52(6):836–855, 2004.
- [69] N. Matloff. Introduction to discrete-event simulation and the simple language. Davis, CA. Dept of Computer Science. University of California at Davis. Retrieved on August, 2(2009): 1–33, 2008.
- [70] S. Meyn. Control techniques for complex networks. Cambridge University Press, 2008.
- [71] S. P. Meyn. Sequencing and routing in multiclass queueing networks part i: Feedback regulation. SIAM Journal on Control and Optimization, 40(3):741–776, 2001.
- [72] V. Mnih, A. P. Badia, M. Mirza, A. Graves, T. Lillicrap, T. Harley, D. Silver, and K. Kavukcuoglu. Asynchronous methods for deep reinforcement learning. In *International* conference on machine learning, pages 1928–1937. PMLR, 2016.

- [73] C. C. Moallemi, S. Kumar, and B. Van Roy. Approximate and data-driven dynamic programming for queueing networks. *working paper*, 2008.
- [74] S. Mohamed, M. Rosca, M. Figurnov, and A. Mnih. Monte carlo gradient estimation in machine learning. *The Journal of Machine Learning Research*, 21(1):5183–5244, 2020.
- [75] P. Naor. The regulation of queue size by levying tolls. Econometrica: journal of the Econometric Society, pages 15–24, 1969.
- [76] M. Neely. Stochastic network optimization with application to communication and queueing systems. Springer Nature, 2022.
- [77] A. Paszke, S. Gross, S. Chintala, G. Chanan, E. Yang, Z. DeVito, Z. Lin, A. Desmaison, L. Antiga, and A. Lerer. Automatic differentiation in pytorch. In *Neural Information Pro*cessing Systems (NIPS) Workshop on Automatic Differentiation, 2017.
- [78] B. S. Pavse, M. Zurek, Y. Chen, Q. Xie, and J. P. Hanna. Learning to stabilize online reinforcement learning in unbounded state spaces. In *Forty-first International Conference on Machine Learning*, 2024.
- [79] G. Qu, A. Wierman, and N. Li. Scalable reinforcement learning of localized policies for multiagent networked systems. In *Learning for Dynamics and Control*, pages 256–266. PMLR, 2020.
- [80] Z. Rosberg, P. Varaiya, and J. Walrand. Optimal control of service in tandem queues. *IEEE Transactions on Automatic Control*, 27(3):600–610, 1982.
- [81] S. Schoenholz and E. D. Cubuk. Jax md: a framework for differentiable physics. Advances in Neural Information Processing Systems, 33:11428–11441, 2020.
- [82] J. Schulman, S. Levine, P. Abbeel, M. Jordan, and P. Moritz. Trust region policy optimization. In *International conference on machine learning*, pages 1889–1897. PMLR, 2015.
- [83] J. Schulman, P. Moritz, S. Levine, M. Jordan, and P. Abbeel. High-dimensional continuous control using generalized advantage estimation. arXiv preprint arXiv:1506.02438, 2015.
- [84] J. Schulman, F. Wolski, P. Dhariwal, A. Radford, and O. Klimov. Proximal policy optimization algorithms. arXiv preprint arXiv:1707.06347, 2017.
- [85] D. Shah, Q. Xie, and Z. Xu. Stable reinforcement learning with unbounded state space. arXiv preprint arXiv:2006.04353, 2020.
- [86] J. G. Shanthikumar, S. Ding, and M. T. Zhang. Queueing theory for semiconductor manufacturing systems: A survey and open problems. *IEEE Transactions on Automation Science* and Engineering, 4(4):513–522, 2007.
- [87] Simio LLC. Simio simulation software. https://www.simio.com, 2024. Accessed: 2024-06-05.
- [88] J. C. Spall. Multivariate stochastic approximation using a simultaneous perturbation gradient approximation. *IEEE transactions on automatic control*, 37(3):332–341, 1992.

- [89] A. L. Stolyar. Maxweight scheduling in a generalized switch: State space collapse and workload minimization in heavy traffic. The Annals of Applied Probability, 14(1):1–53, 2004.
- [90] H. J. Suh, M. Simchowitz, K. Zhang, and R. Tedrake. Do differentiable simulators give better policy gradients? In *International Conference on Machine Learning*, pages 20668–20696. PMLR, 2022.
- [91] R. Suri. Infinitesimal perturbation analysis for general discrete event systems. Journal of the ACM (JACM), 34(3):686–717, 1987.
- [92] R. S. Sutton, D. McAllester, S. Singh, and Y. Mansour. Policy gradient methods for reinforcement learning with function approximation. Advances in neural information processing systems, 12, 1999.
- [93] L. Tassiulas and A. Ephremides. Stability properties of constrained queueing systems and scheduling policies for maximum throughput in multihop radio networks. In 29th IEEE Conference on Decision and Control, pages 2130–2132. IEEE, 1990.
- [94] The AnyLogic Company. Anylogic simulation software. https://www.anylogic.com, 2024. Accessed: 2024-06-05.
- [95] T. H. Tran, L. M. Nguyen, and K. Scheinberg. Finding optimal policy for queueing models: New parameterization. arXiv preprint arXiv:2206.10073, 2022.
- [96] G. Tucker, A. Mnih, C. J. Maddison, J. Lawson, and J. Sohl-Dickstein. Rebar: Low-variance, unbiased gradient estimates for discrete latent variable models. *Advances in Neural Information Processing Systems*, 30, 2017.
- [97] A. Van Den Oord, O. Vinyals, et al. Neural discrete representation learning. Advances in neural information processing systems, 30, 2017.
- [98] N. Walton and K. Xu. Learning and information in stochastic networks and queues. In Tutorials in Operations Research: Emerging Optimization Methods and Modeling Techniques with Applications, pages 161–198. INFORMS, 2021.
- [99] L. Weaver and N. Tao. The optimal reward baseline for gradient-based reinforcement learning. arXiv preprint arXiv:1301.2315, 2013.
- [100] H. Wei, X. Liu, W. Wang, and L. Ying. Sample efficient reinforcement learning in mixed systems through augmented samples and its applications to queueing networks. Advances in Neural Information Processing Systems, 36, 2024.
- [101] R. J. Williams. Simple statistical gradient-following algorithms for connectionist reinforcement learning. *Machine learning*, 8:229–256, 1992.

A Training Details

PPO WC was trained over 100 episodes, each consisting of 50,000 environment steps parallelized over 50 actors. We closely follow the hyper-parameters and training setup as in [26]. We used a discount factor of 0.998, a GAE [83] parameter of 0.99, and set the Kullback–Leibler divergence penalty as 0.03. For the value network, we used a batch size of 2,500, while for the policy network, we used the entire rollout buffer (batch size of 50,000) to take one gradient step. We performed 3 PPO gradient updates on the same rollout data. For all the experiments, we used the Adam optimizer with a cosine decaying warming-up learning rate scheduler. The learning rates were set to 3×10^{-4} for the value network and 9×10^{-4} for the policy network. We used 3% of the training horizon to warm up to the maximum learning rate and then cosine decayed to 1×10^{-5} for both networks. We used the same neural network architecture as those in [26], see Appendix E of [26] for more details.

For the PATHWISE policy gradient, we used the same hyperparameters across all experiments. We use an inverse temperature of $\beta = 10$ for the softmax relaxation of the event selection. We update the policy after every episode with the Adam optimizer, using constant step-size of 5×10^{-4} , momentum parameters (0.8, 0.9), and gradient clipping of 1. For the policy neural network, we used a multilayer perceptron with 3 hidden layers, each hidden layer consisting of 128 hidden units. We use the work-conserving softmax for the final output.

B Proofs

B.1 Proof of Theorem 1

We focus on the case where $x_k \ge 1$. Then,

$$\begin{aligned} x_{k+1} &= x_k + 1\{\tau_k^A < w_k/\mu\} - 1\{\tau_k^A > w_k/\mu\} \\ &= x_k + \frac{e^{-\beta\tau_k^A}}{e^{-\beta\tau_k^A} + e^{-\beta w_k/\mu}} - \frac{e^{-\beta w_k/\mu}}{e^{-\beta\tau_k^A} + e^{-\beta w_k/\mu}} \\ &= x_k + \frac{e^{-\beta\tau_k^A} - e^{-\beta w_k/\mu}}{e^{-\beta\tau_k^A} + e^{-\beta w_k/\mu}}. \end{aligned}$$

Since the inter-arrival times and workloads are exponentially distributed, by the memoryless property, we have $\tau_k^A \sim \mathsf{Exp}(\lambda)$ and $w_k \sim \mathsf{Exp}(1)$.

The true gradient is

$$\frac{d}{d\mu}\mathbb{E}[x_{k+1} - x_k] = \frac{d}{d\mu}\frac{\lambda - \mu}{\lambda + \mu} = -\frac{2\lambda}{(\lambda + \mu)^2}$$

Under our $softmin_{\beta}$ approximation for the event-selection, we have

$$\begin{split} & \mathbb{E}\left[\frac{d}{d\mu}\frac{e^{-\beta\tau_{k}^{A}}-e^{-\beta w_{k}/\mu}}{e^{-\beta\tau_{k}^{A}}+e^{-\beta w_{k}/\mu}}\right] \\ &=\mathbb{E}\left[-2\beta\frac{e^{-\beta(\tau_{k}^{A}+w_{k}/\mu)}}{(e^{-\beta\tau_{k}^{A}}+e^{-\beta w_{k}/\mu})^{2}}\frac{w_{k}}{\mu^{2}}\right] \\ &=-2\frac{\beta}{\mu}\mathbb{E}\left[\tau_{k}^{S}\left(\frac{e^{\beta(\tau_{k}^{A}-\tau_{k}^{S})}}{(e^{\beta(\tau_{k}^{A}-\tau_{k}^{S})}+1)^{2}}\mathbf{1}\{\tau_{k}^{A}<\tau_{k}^{S}\}+\frac{e^{\beta(\tau_{k}^{S}-\tau_{k}^{A})}}{(e^{\beta(\tau_{k}^{S}-\tau_{k}^{A})}+1)^{2}}\mathbf{1}\{\tau_{k}^{A}>\tau_{k}^{S}\}\right)\right], \end{split}$$

for $\tau_k^S = w_k/\mu$. Next, note that

$$\begin{split} & \mathbb{E}\left[\tau_k^S \frac{e^{\beta(\tau_k^A - \tau_k^S)}}{(e^{\beta(\tau_k^A - \tau_k^S)} + 1)^2} \mathbf{1}\{\tau_k^A < \tau_k^S\}\right] \\ &= \mathbb{E}\left[\mathbb{E}\left[\tau_k^S \frac{e^{\beta(\tau_k^A - \tau_k^S)}}{(e^{\beta(\tau_k^A - \tau_k^S)} + 1)^2} \mathbf{1}\{\tau_k^A < \tau_k^S\}\right] \middle| \tau_k^A = t\right] \\ &= \mathbb{E}\left[\mathbb{E}\left[\tau_k^S \frac{e^{\beta(t - \tau_k^S)}}{(e^{\beta(t - \tau_k^S)} + 1)^2} \middle| \tau_k^A = t, \tau_k^S > t\right] \mathbb{P}(\tau_k^S > t | \tau_k^A = t)\right] \\ &= \mathbb{E}\left[\mathbb{E}\left[(t + S') \frac{e^{-\beta S'}}{(e^{-\beta S'} + 1)^2}\right] \mathbb{P}(\tau_k^S > t | \tau_k^A = t)\right] \text{ for } S' \sim \mathsf{Exp}(\mu) \\ &= \mathbb{E}\left[(tA(\beta, \mu) + B(\beta, \mu))\mathbb{P}(\tau_k^S > t | \tau_k^A = t)\right] \\ &= \mathbb{E}\left[(\tau_k^A A(\beta, \mu) + B(\beta, \mu))e^{-\mu\tau_k^A}\right] \\ &= \frac{\lambda}{(\lambda + \mu)^2}A(\beta, \mu) + \frac{\lambda}{(\lambda + \mu)}B(\beta, \mu) \end{split}$$

where

$$\begin{split} A(\beta,\mu) &= \frac{\mu \left(\beta - \mu H\left(\frac{\mu}{2\beta}\right) + \mu H\left(\frac{\mu}{2\beta} - \frac{1}{2}\right)\right)}{2\beta^2} \\ &= \frac{\mu}{2\beta} + \frac{\mu^2}{2\beta^2} \left(\underbrace{H\left(\frac{\mu}{2\beta} - \frac{1}{2}\right) - H\left(\frac{\mu}{2\beta}\right)}_{\tilde{H}(\beta,\mu)}\right) \\ B(\beta,\mu) &= \frac{\mu}{4\beta^3} \left(2\beta H\left(\frac{\mu}{2\beta}\right) - 2\beta H\left(\frac{\mu}{2\beta} - \frac{1}{2}\right) - \mu\psi^{(1)}\left(\frac{\beta + \mu}{2\beta}\right) + \mu\psi^{(1)}\left(\frac{2\beta + \mu}{2\beta}\right)\right) \\ &= -\frac{\mu}{2\beta^2} \tilde{H}(\beta,\mu) + \frac{\mu^2}{4\beta^3} \left(\underbrace{\psi^{(1)}\left(\frac{2\beta + \mu}{2\beta}\right) - \psi^{(1)}\left(\frac{\beta + \mu}{2\beta}\right)}_{\tilde{\psi}^{(1)}(\beta,\mu)}\right) \end{split}$$

Moreover, note that

$$H\left(\frac{\mu}{2\beta}\right) = \psi^{(0)}\left(\frac{\mu}{2\beta} + 1\right) + \gamma$$
$$H\left(\frac{\mu}{2\beta} - \frac{1}{2}\right) = \psi^{(0)}\left(\frac{\mu}{2\beta} + \frac{1}{2}\right) + \gamma$$

Note that

$$H\left(\frac{\mu}{2\beta}\right) - H\left(\frac{\mu}{2\beta} - \frac{1}{2}\right) = \log\left(\frac{\mu}{2\beta} + 1\right) + \frac{1}{\frac{\mu}{2\beta} + 1} - \log\left(\frac{\mu}{2\beta} + \frac{1}{2}\right) + \frac{1}{\frac{\mu}{2\beta} + 1}.$$

Similarly,

$$\begin{split} & \mathbb{E}\left[\tau_k^S \frac{e^{\beta(\tau_k^S - \tau_k^A)}}{(e^{\beta(\tau_k^S - \tau_k^A)} + 1)^2} \mathbf{1}\{\tau_k^A > \tau_k^S\}\right] \\ &= \mathbb{E}\left[\mathbb{E}\left[s\frac{e^{-\beta(\tau_k^A - s)}}{(e^{-\beta(\tau_k^A - s)} + 1)^2} \mathbf{1}\{\tau_k^A > s\} \middle| \tau_k^S = s\right]\right] \\ &= \mathbb{E}\left[\mathbb{E}\left[s\frac{e^{-\beta(\tau_k^A - s)}}{(e^{-\beta(\tau_k^A - s)} + 1)^2} \middle| \tau_k^A > s, \tau_k^S = s\right] \mathbb{P}(\tau_k^A > s | \tau_k^S = s)\right] \\ &= \mathbb{E}\left[\mathbb{E}\left[s\frac{e^{-\beta T'}}{(e^{-\beta T'} + 1)^2}\right] \mathbb{P}(\tau_k^A > s | \tau_k^S = s)\right] \text{ for } T' \sim \mathsf{Exp}(\lambda) \\ &= \mathbb{E}\left[\tau_k^S A(\beta, \lambda) e^{-\lambda \tau_k^S}\right] \\ &= \frac{\mu}{(\lambda + \mu)^2} A(\beta, \lambda). \end{split}$$

Then,

$$\begin{split} &-2\frac{\beta}{\mu}\mathbb{E}\left[\tau_k^S\left(\frac{e^{\beta(\tau_k^A-\tau_k^S)}}{(e^{\beta(\tau_k^A-\tau_k^S)}+1)^2}\mathbf{1}\{\tau_k^A<\tau_k^S\}+\frac{e^{\beta(\tau_k^S-\tau_k^A)}}{(e^{\beta(\tau_k^S-\tau_k^A)}+1)^2}\mathbf{1}\{\tau_k^A>\tau_k^S\}\right)\right]\\ &=-2\frac{\beta}{\mu}\left(\frac{\lambda}{(\lambda+\mu)^2}A(\beta,\mu)+\frac{\lambda}{(\lambda+\mu)}B(\beta,\mu)+\frac{\mu}{(\lambda+\mu)^2}A(\beta,\lambda)\right)\\ &=-2\frac{\beta}{\mu}\left(\frac{\lambda A(\beta,\mu)+\mu A(\beta,\lambda)}{(\lambda+\mu)^2}+\frac{\lambda}{(\lambda+\mu)}B(\beta,\mu)\right)\\ &=\frac{-2\lambda}{(\lambda+\mu)^2}-\frac{2\beta}{\mu}\left(\frac{\lambda\mu^2\tilde{H}(\beta,\mu)+\mu\lambda^2\tilde{H}(\beta,\lambda)}{2\beta^2(\lambda+\mu)^2}-\frac{\lambda\mu}{2\beta^2(\lambda+\mu)}\tilde{H}(\beta,\mu)+\frac{\lambda}{(\lambda+\mu)}\frac{\mu^2}{4\beta^3}\tilde{\psi}^{(1)}(\beta,\mu)\right). \end{split}$$

Note that as $\beta \to \infty$,

$$\lim_{\beta \to \infty} \tilde{H}(\beta, \mu) = \gamma + \psi^{(0)} \left(\frac{1}{2}\right),$$
$$\lim_{\beta \to \infty} \tilde{\psi}^{(1)}(\beta, \mu) = -\frac{\pi^2}{3}.$$

This means that the leading order term is $O(1/\beta^2)$. In particular,

$$-\frac{2\beta}{\mu}\left(\frac{\lambda\mu^2\tilde{H}(\beta,\mu)+\mu\lambda^2\tilde{H}(\beta,\lambda)}{2\beta^2(\lambda+\mu)^2}-\frac{\lambda\mu\tilde{H}(\beta,\mu)}{2\beta^2(\lambda+\mu)}\right)\sim\frac{\pi^2\lambda^2(\mu-\lambda)}{6\beta^2(\lambda+\mu)^2}$$

Finally, we have the second-order term

$$-\frac{2\beta}{\mu}\frac{\lambda}{(\lambda+\mu)}\frac{\mu^2}{4\beta^3}\tilde{\psi}^{(1)}(\beta,\mu)\sim\frac{\pi^2\lambda\mu}{6\beta^2(\lambda+\mu)}.$$

Thus, we have the following characterization of the bias:

$$\mathbb{E}\left[\frac{d}{d\mu}\frac{e^{-\beta\tau_k^A} - e^{-\beta w_k/\mu}}{e^{-\beta\tau_k^A} + e^{-\beta w_k/\mu}}\right] - \left(\frac{-2\lambda}{(\lambda+\mu)^2}\right) \sim \frac{1}{\beta^2}\frac{\pi^2\lambda(\mu^2 - \lambda^2 + 2\mu\lambda)}{6(\lambda+\mu)^2} + o\left(\frac{1}{\beta^2}\right)$$

For variance, we have

$$\begin{split} & \mathbb{E}\left[\left(-2\beta \frac{e^{\beta(\tau_{k}^{A}+w_{k}/\mu)}}{(e^{\beta\tau_{k}^{A}}+e^{\beta w_{k}/\mu})^{2}}\frac{w_{k}}{\mu^{2}}\right)^{2}\right] \\ &= \mathbb{E}\left[\frac{4\beta^{2}}{\mu^{2}}\frac{e^{2\beta(\tau_{k}^{A}+\tau_{k}^{S})}}{\left(e^{\beta\tau_{k}^{A}}+e^{\beta\tau_{k}^{S}}\right)^{4}}\tau_{k}^{S,2}\right] \\ &= \frac{4\beta^{2}}{\mu^{2}}\mathbb{E}\left[\tau_{k}^{S,2}\left(\frac{e^{2\beta(\tau_{k}^{A}+\tau_{k}^{S})}}{\left(e^{\beta\tau_{k}^{A}}+e^{\beta\tau_{k}^{S}}\right)^{4}}\mathbf{1}\{\tau_{k}^{A}<\tau_{k}^{S}\}+\frac{e^{2\beta(\tau_{k}^{A}+\tau_{k}^{S})}}{\left(e^{\beta\tau_{k}^{A}}+e^{\beta\tau_{k}^{S}}\right)^{4}}\mathbf{1}\{\tau_{k}^{A}<\tau_{k}^{S}\}\right] \\ &= \frac{4\beta^{2}}{\mu^{2}}\mathbb{E}\left[\tau_{k}^{S,2}\left(\frac{e^{2\beta(\tau_{k}^{A}-\tau_{k}^{S})}}{\left(e^{\beta(\tau_{k}^{A}-\tau_{k}^{S})}+1\right)^{4}}\mathbf{1}\{\tau_{k}^{A}<\tau_{k}^{S}\}+\frac{e^{2\beta(\tau_{k}^{S}-\tau_{k}^{A})}}{\left(e^{\beta(\tau_{k}^{S}-\tau_{k}^{A})}+1\right)^{4}}\mathbf{1}\{\tau_{k}^{A}<\tau_{k}^{S}\}\right] \end{split}$$

Note that

$$\begin{split} & \mathbb{E}\left[\tau_{k}^{S,2}\frac{e^{2\beta(\tau_{k}^{A}-\tau_{k}^{S})}}{\left(e^{\beta(\tau_{k}^{A}-\tau_{k}^{S})}+1\right)^{4}}\mathbf{1}\{\tau_{k}^{A}<\tau_{k}^{S}\}\right] \\ &= \mathbb{E}\left[\mathbb{E}\left[\tau_{k}^{S,2}\frac{e^{2\beta(\tau_{k}^{A}-\tau_{k}^{S})}}{\left(e^{\beta(\tau_{k}^{A}-\tau_{k}^{S})}+1\right)^{4}}\mathbf{1}\{t<\tau_{k}^{S}\}\right]\right|\tau_{k}^{A}=t\right] \\ &= \mathbb{E}\left[\mathbb{E}\left[(t+S')^{2}\frac{e^{-2\beta S'}}{(e^{-\beta S'}+1)^{4}}\right]\mathbb{P}(\tau_{k}^{S}>t|\tau_{k}^{A}=t)\right] \text{ for } S'\sim \mathsf{Exp}(\mu) \\ &= \mathbb{E}\left[\mathbb{E}\left[(t^{2}+2S'+S'^{2})\frac{e^{-2\beta S'}}{(e^{-\beta S'}+1)^{4}}\right]\mathbb{P}(S_{i}>t|T_{i}=t)\right] \\ &= \mathbb{E}\left[\left(t^{2}\tilde{A}(\beta,\mu)+t\tilde{B}(\beta,\mu)+\tilde{C}(\beta,\mu)\right)\mathbb{P}(S_{i}>t|T_{i}=t)\right] \\ &= \mathbb{E}\left[\left(T_{i}^{2}\tilde{A}(\beta,\mu)+t_{i}\tilde{B}(\beta,\mu)+\tilde{C}(\beta,\mu)\right)e^{-\mu T_{i}}\right] \\ &= \frac{2\lambda}{(\lambda+\mu)^{3}}\tilde{A}(\beta,\mu)+\frac{\lambda}{(\lambda+\mu)^{2}}\tilde{B}(\beta,\mu)+\frac{\lambda}{(\lambda+\mu)}\tilde{C}(\beta,\mu). \end{split}$$

Similarly,

$$\mathbb{E}\left[\tau_k^{S,2} \frac{e^{2\beta(\tau_k^S - \tau_k^A)}}{\left(e^{\beta(\tau_k^S - \tau_k^A)} + 1\right)^4} \mathbf{1}\{\tau_k^A > \tau_k^S\}\right]$$
$$= \mathbb{E}\left[\mathbb{E}\left[s^2 \frac{e^{2\beta(s - \tau_k^A)}}{\left(e^{\beta(s - \tau_k^A)} + 1\right)^4} \mathbf{1}\{\tau_k^A > s\} \middle| \tau_k^S = s\right]\right]$$

$$\begin{split} &= \mathbb{E}\left[\mathbb{E}\left[s^2 \frac{e^{-2\beta(\tau_k^A - s)}}{(e^{-\beta(\tau_k^A - s)} + 1)^4} | \tau_k^A > s, \tau_k^S = s\right] \mathbb{P}(\tau_k^A > s | \tau_k^S = s)\right] \\ &= \mathbb{E}\left[\mathbb{E}\left[s^2 \frac{e^{-2\beta T'}}{(e^{-\beta T'} + 1)^4}\right] \mathbb{P}(\tau_k^A > s | \tau_k^S = s)\right] \text{ for } T' \sim \mathsf{Exp}(\lambda) \\ &= \mathbb{E}\left[\tilde{A}(\beta, \lambda) \tau_k^{S,2} e^{-\lambda \tau_k^S}\right] \\ &= \frac{2\mu}{(\lambda + \mu)^2} \tilde{A}(\beta, \lambda). \end{split}$$

Putting the above two parts together, we have

$$\begin{split} & \frac{4\beta^2}{\mu^2} \mathbb{E} \left[\tau_k^{S,2} \left(\frac{e^{2\beta(\tau_k^A - \tau_k^S)}}{\left(e^{\beta(\tau_k^A - \tau_k^S)} + 1 \right)^4} \mathbf{1}\{\tau_k^A < \tau_k^S\} + \frac{e^{2\beta(\tau_k^S - \tau_k^A)}}{\left(e^{\beta(\tau_k^S - \tau_k^A)} + 1 \right)^4} \mathbf{1}\{\tau_k^A > \tau_k^S\} \right) \right] \\ &= \frac{4\beta^2}{\mu^2} \left(\frac{2\lambda}{(\lambda + \mu)^3} \tilde{A}(\beta, \mu) + \frac{\lambda}{(\lambda + \mu)^2} \tilde{B}(\beta, \mu) + \frac{\lambda}{(\lambda + \mu)} \tilde{C}(\beta, \mu) + \frac{2\mu}{(\lambda + \mu)^2} \tilde{A}(\beta, \lambda) \right) \\ &\sim \frac{4\beta\lambda}{3\mu(\lambda + \mu)^2} \text{ as } \beta \to \infty. \end{split}$$

B.2 Proof of Theorem 2

By assumption, $h \leq c(\mu - \lambda)$ for some c < 1. First, we develop bound for $\operatorname{Var}_{\infty}(\hat{\nabla}^{\mathsf{R}}J_{N}(\theta;\xi_{1:N}))$. We can compute the variance by conditioning on the value of Y:

$$\operatorname{Var}_{\infty}(\hat{Q}_{N}(\mu - hY) \cdot (\log Y - \frac{1}{\theta})) = \mathbb{E}\left[\operatorname{Var}_{\infty}\left(\hat{Q}_{N}(\mu - hy)(\log y - \frac{1}{\theta})|Y = y\right)\right] + \operatorname{Var}\left(\mathbb{E}\left[\hat{Q}_{N}(\mu - hy)(\log y - \frac{1}{\theta})|Y = y\right)\right]\right).$$

For the first term, note that the asymptotic variance in the CLT for the ergodic estimator $\hat{Q}_N(\mu)$ is $\operatorname{Var}_{\infty}(\hat{Q}(\mu)) = \frac{2\rho(1+\rho)}{(1-\rho)^4}$. Then, we have $\operatorname{Var}_{\infty}(\hat{Q}_N(\mu)) = \frac{2\rho(1+\rho)}{N(1-\rho)^4}$. Since $h \leq \mu$, $\operatorname{Var}_{\infty}(\hat{Q}_N(\mu)) = \frac{2\rho(1+\rho)}{N(1-\rho)^4}$. Then,

$$\mathbb{E}\left[\operatorname{Var}_{\infty}\left(\hat{Q}_{N}(\mu-hy)\cdot(\log y+\frac{1}{\theta})|Y=y\right)\right]$$
$$=\mathbb{E}\left[(\log y+\frac{1}{\theta})^{2}\operatorname{Var}_{\infty}(\hat{Q}_{N}(\mu-hy)|Y=y)\right]$$
$$\geq\mathbb{E}\left[(\log Y+\frac{1}{\theta})^{2}\frac{2\rho(1+\rho)}{N(1-\rho)^{4}}\right]$$
$$=\frac{1}{\theta^{2}}\frac{2\rho(1+\rho)}{N(1-\rho)^{4}}=\Theta\left((1-\rho)^{-4}\right),$$

where the last equality uses the fact that for $Y \sim \text{Beta}(\theta, 1)$,

$$\mathbb{E}[\log Y] = \psi(\theta) - \psi(\theta + 1) = -\frac{1}{\theta}$$

and

$$\mathbb{E}\left[\left(\log Y + \frac{1}{\theta}\right)^2\right] = \operatorname{Var}(\log Y) = \psi_1(\theta) - \psi_1(\theta + 1) = \frac{1}{\theta^2}.$$

For the second term, we plug in the true estimand as the expectation of $Q_N(\mu)$, i.e. $Q(\mu) = \frac{\lambda}{\mu - \lambda}$. Then,

$$\operatorname{Var}\left(\mathbb{E}\left[\hat{Q}_{N}(\mu-hy)(\log y+\frac{1}{\theta})|Y=y)\right]\right)$$
$$=\operatorname{Var}\left(\frac{\lambda}{\mu-hY-\lambda}\left(\log Y+\frac{1}{\theta}\right)\right)$$
$$=\mathbb{E}\left[\left(\frac{\lambda}{\mu-hY-\lambda}\right)^{2}\left(\log Y+\frac{1}{\theta}\right)^{2}\right]-\mathbb{E}\left[\left(\frac{\lambda}{\mu-hY-\lambda}\right)\left(\log Y+\frac{1}{\theta}\right)\right]^{2}.$$

We proceed to evaluate these expectations analytically.

$$\mathbb{E}\left[\left(\frac{\lambda}{\mu - hY - \lambda}\right)\left(\log Y + \frac{1}{\theta}\right)\right]$$

= $\frac{\rho}{1 - \rho}\left(\frac{\Gamma(\theta)}{\Gamma(1 + \theta)}F_1^2\left(1, \theta, 1 + \theta, \frac{h}{\mu - \lambda}\right) - \theta\Phi\left(\frac{h}{\mu - \lambda}, 2, \theta\right)\right)$
= $O((1 - \rho)^{-1}),$

where F_1^2 is the Hypergoe metric 2F1 function and Φ is the Lerch transcendental function. We also have

$$\begin{split} & \mathbb{E}\left[\left(\frac{\lambda}{\mu-hY-\lambda}\right)^2 \left(\log Y + \frac{1}{\theta}\right)^2\right] \\ &= \frac{\lambda^2}{\theta^2(\mu-\lambda)^3} \left(2(\mu-\lambda) + h\theta^3 \Phi\left(\frac{h}{\mu-\lambda}, 2, \theta+1\right) + h\theta^3(\theta-1)\Phi\left(\frac{h}{\mu-\lambda}, 3, \theta+1\right) \right. \\ &+ \left(\mu-\lambda\right) \left(\theta\frac{\mu-\lambda}{\mu-h-\lambda} - (\theta-1)F_1^2(1, \theta, 1+\theta, \frac{h}{\mu-\lambda})\right) \\ &+ 2(\mu-\lambda)F_2^3\left((2, \theta, \theta), (1+\theta, 1+\theta), \frac{h}{\mu-\lambda}\right)\right) \\ &= O((1-\rho)^{-2}). \end{split}$$

Taking the difference between the above two parts under the limit as $(1 - \rho) \rightarrow 0$, we have

$$\operatorname{Var}\left(\frac{\lambda}{\mu - hY - \lambda}\left(\log Y + \frac{1}{\theta}\right)\right) = O((1 - \rho)^{-2}).$$

Next, we develop a bound for $\operatorname{Var}_{\infty}(\hat{\nabla}J_N(\theta;\xi_{1:N}))$.

$$\operatorname{Var}_{\infty}\left(h \cdot \hat{\nabla}Q_{N}(\mu - hY)\left(\frac{1}{\theta}Y\log Y\right)\right) = \mathbb{E}\left[\operatorname{Var}_{\infty}\left(h \cdot \hat{\nabla}Q_{N}(\mu - hY)\left(\frac{1}{\theta}Y\log Y\right)|Y = y\right)\right] + \operatorname{Var}\left(\mathbb{E}\left[h \cdot \hat{\nabla}Q_{N}(\mu - hY)\left(\frac{1}{\theta}Y\log Y\right)|Y = y\right]\right)$$

For the first term, we can use the fact that $|Y \log Y| \leq 1/e$ almost surely since $Y \in [0, 1]$. We also

use recent results in [54], which compute the asymptotic variance of the IPA estimator:

$$\operatorname{Var}_{\infty}(\hat{\nabla}Q_N(\mu)) = \frac{1 + 16\rho + 27\rho^2 + 2\rho^3 + 6\rho^4}{\mu^2 N(1+\rho)(1-\rho)^5} \le 52\mu^{-2}N^{-1}(1-\rho)^{-5}$$

Under $\mu - hy$, the congestion factor $1 - \frac{\lambda}{\mu - hy} = \frac{\mu - hy - \lambda}{\mu - hy} \ge \frac{\mu - cy(\mu - \lambda) - \lambda}{\mu} = (1 - cy)(1 - \rho)$ and $\mu - hy \ge \mu - h \ge (1 - c)\mu$. So we have the bound,

$$\begin{aligned} \operatorname{Var}_{\infty} \left(h \hat{\nabla} Q_N(\mu - hY) \left(\frac{1}{\theta} y \log y \right) | Y = y \right) \\ &\leq h^2 \theta^{-2} (y \log y)^2 \operatorname{Var}_{\infty} (\hat{\nabla} Q_N(\mu - hy)) \\ &\leq h^2 \mu^{-2} \theta^{-2} e^{-2} 52 N^{-1} (1 - c)^{-7} (1 - \rho)^{-5} \\ &= O(N^{-1} h^2 \mu^{-2} (1 - \rho)^{-5}) \\ &= O(N^{-1} (1 - \rho)^{-3}) \end{aligned}$$

since $h = O(1 - \rho)$.

For the second term, we plug in the true estimand as the mean of $\hat{\nabla}Q_N(\mu)$, i.e., $\nabla Q(\mu) = -\frac{\rho}{\mu(1-\rho)^2}$,

$$\mathbb{E}\left[h \cdot \hat{\nabla}Q_N(\mu - hY)\left(\frac{1}{\theta}Y\log Y\right)|Y = y\right] = -h\frac{1}{\theta}(y\log y)\frac{\lambda}{(\mu - hy - \lambda)^2}.$$

We next evaluate the variance analytically,

$$\operatorname{Var}\left(-h\frac{1}{\theta}(Y\log Y)\frac{\lambda}{(\mu-hY-\lambda)^2}\right) = \mathbb{E}\left[\left(-h\frac{1}{\theta}(Y\log Y)\frac{\lambda}{(\mu-hY-\lambda)^2}\right)^2\right] - \mathbb{E}\left[-h\frac{1}{\theta}(Y\log Y)\frac{\lambda}{(\mu-hY-\lambda)^2}\right]^2$$

Since

$$\begin{split} & \mathbb{E}\left[-h\frac{1}{\theta}(Y\log Y)\frac{\lambda}{(\mu-hY-\lambda)^2}\right] \\ &= \frac{h}{(1+\theta)^2}\frac{\lambda}{(\mu-\lambda)^2}F_2^3\left((2,1+\theta,1+\theta),(2+\theta,2+\theta),\frac{h}{\mu-\lambda}\right) \\ &= O((1-\rho)^{-1}) \end{split}$$

and

$$\begin{split} & \mathbb{E}\left[\left(-h\frac{1}{\theta}(Y\log Y)\frac{\lambda}{(\mu-hY-\lambda)^2}\right)^2\right] \\ &= 2h\theta^2\Gamma(\theta)^3\Gamma(2+\theta)^{-3}\frac{\lambda^2}{(\mu-\lambda)^2} \times \left(F_3^4\left((3,1+\theta,1+\theta,1+\theta),(2+\theta,2+\theta,2+\theta),\frac{h}{\mu-\lambda}\right)\right) \\ &-F_3^4\left((4,1+\theta,1+\theta,1+\theta),(2+\theta,2+\theta,2+\theta),\frac{h}{\mu-\lambda}\right)\right) \\ &= O((1-\rho)^{-2}), \end{split}$$

we have

$$\operatorname{Var}\left(-h\frac{1}{\theta}(Y\log Y)\frac{\lambda}{(\mu-hY-\lambda)^2}\right) = O((1-\rho)^{-2}).$$

B.3 Proof of Corollary 2

First, we can explicitly characterize the optimal baseline:

$$b^* = \frac{\mathbb{E}[Q(\mu - hY)\nabla_{\theta}\log \pi_{\theta}(Y)^2]}{\mathbb{E}[\nabla_{\theta}\log \pi_{\theta}(Y)^2]}$$

=
$$\frac{\mathbb{E}[Q(\mu - hY)(\log Y + 1/\theta)^2]}{\mathbb{E}[(\log Y + 1/\theta)^2]}$$

=
$$\frac{\lambda}{\mu - \lambda} \underbrace{\left[F_1^2\left(1, \theta, 1 + \theta, \frac{h}{\mu - \lambda}\right) - 2\theta^2 \Phi\left(\frac{h}{\mu - \lambda}, 2, \theta\right) + 2\frac{h}{\mu - \lambda}\theta^3 \Phi\left(\frac{h}{\mu - \lambda}, 3, 1 + \theta\right)\right]}_{b(\theta)}$$

=
$$O((1 - \rho)^{-1}).$$

Next, we plug this into the REINFORCE estimator.

$$\operatorname{Var}_{\infty}((\hat{Q}_{N}(\mu - hY) - b^{*}) \cdot (\log Y - \frac{1}{\theta})) = \mathbb{E}\left[\operatorname{Var}_{\infty}\left((\hat{Q}_{N}(\mu - hy) - b^{*})(\log y - \frac{1}{\theta})|Y = y\right)\right] + \operatorname{Var}\left(\mathbb{E}\left[(\hat{Q}_{N}(\mu - hy) - b^{*})(\log y - \frac{1}{\theta})|Y = y\right]\right).$$

Note that since b^* is a constant, the first term, i.e., the mean of the conditional variance given Y, has the same value as in Theorem 2. For the second term, note that since $h = c(\mu - \lambda)$,

$$\operatorname{Var}\left[\left(\frac{\lambda}{\mu - hY - \lambda} - b^*\right) \left(\log Y + \frac{1}{\theta}\right)\right]$$
$$= \left(\frac{\lambda}{\mu - \lambda}\right)^2 \operatorname{Var}\left[\left(\frac{1}{1 - cY} - b(\theta)\right) \left(\log Y + \frac{1}{\theta}\right)\right]$$

Since $\operatorname{Var}\left[\left(\frac{1}{1-cY}-b(\theta)\right)\left(\log Y+\frac{1}{\theta}\right)\right]>0$ and doesn't depend on μ or λ , this confirms that the second term is $\Theta((1-\rho)^{-2})$