VARIATIONAL INFERENCE FOR MARKOVIAN QUEUEING NETWORKS

IKER PEREZ,* *University of Nottingham*

GIULIANO CASALE,** *Imperial College London*

Abstract

Queueing networks are stochastic systems formed by interconnected resources routing and serving jobs. They induce jump processes with distinctive properties, and find widespread use in inferential tasks. Here, service rates for jobs and potential bottlenecks in the routing mechanism must be estimated from a reduced set of observations. However, this calls for the derivation of complex conditional density representations, both over the stochastic network trajectories and rates, and it is considered an intractable problem. Numerical simulation procedures designed for this purpose do not scale due to high computational costs; furthermore, *variational* approaches relying on approximating measures and full independence assumptions are unsuitable. In this paper, we offer a probabilistic interpretation of *variational* methods applied to inference tasks with queueing networks, and show that approximating measure choices routinely used with jump processes yield ill-defined optimization problems. Yet, we demonstrate that it is still possible to enable a *variational* inferential task, by considering a novel space expansion treatment over an analogue counting process for job transitions. Overall, we present and compare exemplar use cases with practical queueing networks, and show that our framework offers an efficient and improved alternative where existing variational or numerically intensive solutions fail.

Keywords: Queueing networks; variational methods; Markov jump process

* Postal address: School of Mathematical Sciences, University of Nottingham, Nottingham NG7 2RD, United Kingdom

** Postal address: Department of Computing, Imperial College London, London SW7 2RH, United Kingdom

2010 Mathematics Subject Classification: Primary 60J27

Secondary 60J28

1. Introduction

Queueing networks (QNs) are systems of theoretical and practical interest in the design of computing systems [14], as well as in the optimization of business processes arising in factories, shops, offices or hospitals [6, 15, 21]. They are formed by interconnected resources routing and serving jobs, and their behaviour often gives rise to complex families of stochastic (jump) processes. In applications, they provide the means to assess modifications, diagnose performance and evaluate robustness in multiple service infrastructures.

Formally, a QN is associated with a (Markov) jump process; the piecewise-deterministic model for job counts stationed at each resource. Here, a resource consists of a *service station* along with a *queue*, because the processing capabilities of a station are usually limited. Whenever a job is serviced, this is routed from its current resource into a new queue, and we say that the process has *jumped*. Hence, a jump adds a count unit to one resource by subtracting it to another, and we say that jumps in QNs are *coupled*. Finally, the process is parametrized by the set of *intensity* rates, which determine the pace of service across job types and the multiple stations.

Performance evaluation. In applications, the service rates in a QN may be unknown, and the various bottlenecks in the routing mechanism for jobs (across queues) are hard to anticipate. Consequently, network deployments are periodically monitored to collect *performance* measurements, which often include queue lengths, visit counts or response times. These measurements are a consequence of the stochastic evolution of the underlying Markov model; thus, they form the quantitative basis allowing to draw inferences on multiple aspects of the network. This inferential process is referred to as *performance evaluation* [8] and is not limited to estimating the service rates; but furthermore (i) quantifying any uncertainty around estimations and (ii) identifying queue length distributions under varying networks conditions.

From a probabilistic viewpoint, the primary inferential objective is to parametrize the *regular conditional probability* for stochastic network *trajectories* or *paths*, conditioned on the available performance measurements; and to derive *closed-form* expres-

sions for the induced distributions over (i) intensity rates and (ii) queue lengths over time. In this paper, we will observe that this goal calls for the evaluation of complex intractable integrals, a problem commonly encountered within Bayesian statistical settings [1]. There, popular approaches generally deviate from analytical methodology and instead construct numerical representations of the induced distributions (such as density plots and histograms), using intensive *Markov Chain Monte Carlo* procedures [27, 22]. However, numerical methods must sample the trajectories of the underlying stochastic model, conditioned on measurements. These work well in analogue problems with jump processes observed in mathematical biology [12] or genetics [11]; yet, scale poorly to complex multivariate settings common in QNs [5], which account for coupled jumps, job priorities, service types or feedback loops. Also, numerical methods are reliant on *sufficient* readily-available measurements to describe the likely behaviour of the system. In real-world applications, performance measurements are scarce [27]; for instance, in large externally managed implementations the monitoring may only be executed *end-to-end* [17], i.e., exclusively for the input and output resources where jobs enter and depart the QN.

Consequently, pragmatic statistical solutions are restricted in scope, and often ignore the temporal component of the Markov process. Instead, they construct varied inferential methodology by relying on the mass representation for network’s stationary distribution [16], and the objective is usually to only offer point estimates for the unknown service rates [see 26, for a review].

1.1. Approximate inferential methods

In order to derive *closed-form* expressions for the induced conditional distributions over service rates or network trajectories, we must resort to *approximate* inferential settings of significant probabilistic complexity. Here, the idea is to operate under some (suitably parametrized) *approximating* measure in the probability space. This approximating measure should define alternative sample trajectories for the underlying Markov jump model (for job counts across resources); and ensure analytical tractability of the various integrals that arise in the inferential problem. However, the reader will note that *closed-form* formulae derived under some (conveniently defined) approximating measure will not be necessarily representative of the real induced distributions. An

important step under an approximate inferential set-up is to minimize the dissimilarity (or *divergence*) between (i) the proposed approximating measure and (ii) the true *regular conditional probability* over events, conditioned on performance metrics (which also defines a measure).

This approximate procedure is a popular technique in Bayesian statistics, and the resulting *closed-form* representations are commonly themed *variational* approximations [4]; because their derivation relies heavily on functional optimization methods in the study of *calculus of variations*. The choices made in order to parametrize suitable approximating measures are key to success; however, the entire process is convoluted, and there exist no guidelines suited to every kind of stochastic model. The standard choice in the statistical literature relies on the strongest of all independence assumptions [cf. 18, 7], to describe each marginal resource (within the multivariate Markov jump model) as a completely independent stochastic process. The underlying parametrization is commonly referred to as the *mean-field* approximating measure, because it draws from the physical *mean-field theory*, where high-dimensional stochastic models are studied through approximations of multiple independent models combined together. Overall, the method works well in practice, and has shown to retrieve accurate conditional density representations in multiple application domains, including predator-prey models, epidemics, networks and similar processes [18, 19, 7, 28, 30, 31].

However, an important commonality across all Markov models addressed in prior work is that jumps in the marginal components take place one at a time. This ensures that the multivariate stochastic process can indeed be approximated by a collection of independent univariate models. Our intent is to approximate the stochastic dynamics for jobs counts across resources in a QN, which are subject to coupled jumps; here, the reader may note that it is probabilistically infeasible to assume full independence, for two independent jump models may not change state at the same time. Currently, there exists no viable alternative or solution that can facilitate an approximate inferential set-up tailored to multivariate systems such as QNs.

Key contributions. In this paper, we will formally describe why coupled jumps hinder the construction of approximating measures as commonly defined in the literature. We will follow by proving that it is still possible to efficiently enable the approximate inferential task, by considering a novel space expansion treatment of the underlying

jump process; in a nut-shell, we will

1. shift the scope; traditionally, a QN is formalized as a counting process of jobs *queueing* in the resources; instead, we address a process of job *transition* counts across resources,
2. augment the support space; so that job counts across queues can become negative.

Finally, we present use cases of our proposed procedure for performance evaluation tasks, applied to example inferential problems where (i) numerical methods do not scale and (ii) approximate variational procedures are unusable. The results within this paper are relevant for single or multi-class Markovian systems (with exponential or phase-type service times), with either finite or infinite processors, multiple types of service disciplines and probabilistic routings.

Structure. The rest of the paper is organised as follows. In Section 2 we offer a hierarchical formulation of a queueing system along with the problem statement. Section 3 introduces an approximating network model and offers a summary of the main results to be presented later in the paper. Sections 4 and 5 include the main contributions in our work; these discuss the treatment of the network system by means of interactions in network resources, and further present the results, proofs and technical details that contribute to later algorithmic constructions. In Section 6, we guide the reader through applications of our results within inferential and network evaluation tasks and in Section 7 we conclude the paper with a discussion.

2. Queueing systems and jump processes

In the following, we employ shorthand notation for densities, base measures and distributions whenever these are clear from the context. From here on, let $(\Omega, \mathcal{F})$ denote a measurable space with the regular conditional probability property, supporting the various rates, trajectories and observations. A general form queueing network comprises some $M \in \mathbb{N}$ service stations along with a *set of job classes* $\mathcal{C}$. The stations are connected by a *network topology* that governs the underlying routing mechanisms; when a job is serviced in one station, it can either queue for service at a different station, or depart the network. Such topology is often defined as a set of *routing probability*

matrices $\{P^c\}_{c \in \mathcal{C}}$, with elements $p_{i,j}^c$ that denote the probability for a class $c \in \mathcal{C}$ job to immediately transit to queueing station j after service completion in station i , for all $0 \leq i, j \leq M$. In open queueing systems, the index 0 is used as a *virtual* external source (and destination) of job arrivals and departures to, and from, the network. In closed systems, this index may either not exist, or instead refer to a delay server that routes departing jobs back into the network. Also, it holds that $\sum_{j=0}^M p_{i,j}^c = 1$, for all $0 \leq i \leq M, c \in \mathcal{C}$.

We address time-homogeneous Markovian systems that are parametrized by exponential inter-arrival and service times, with non-negative rates $\boldsymbol{\mu} = \{\mu_i^c \in \mathbb{R}_+ : 0 \leq i \leq M, c \in \mathcal{C}\}$, which may vary across service stations and job classes. The servers in the network stations may have finite or infinite processors, and service disciplines can differ across a range of *processor sharing* (PS) policies, *first-come first-served* (FCFS) and variations including *last-come first-served* (LCFS) or *service in random order* (SIRO). In some cases, FCFS processors may require shared processing times across the various job classes [cf. 3]. For simplicity and ease of notation, class *switching*, service *priorities* or queue-length dependent service rates are not discussed in detail, however, these follow naturally and we later present some examples of such instances. Under standard exponential service assumptions, the underlying system behaviour is described by a *Markov jump process* (MJP) $X = (X_t)_{t \geq 0}$ with values defined in a measurable space $(\mathcal{S}, \mathcal{P}(\mathcal{S}))$. Here, $\mathcal{S}$ denotes a countable set of feasible states in the network, usually infinite in open or mixed systems and finite within closed ones; $\mathcal{P}(\mathcal{S})$ denotes the power set of $\mathcal{S}$. We allow for $\mathcal{S}$ to support vectors of integers that represent job counts across the various class types and service stations, and denote by $X_t^{i,c}$ the number of class c jobs in station $i > 0$ at time $t \geq 0$. Note that here we ignore the queue lengths in the external delay ($i = 0$) within closed systems, since these are uniquely determined given the number of jobs in the remaining stations. The *infinitesimal generator matrix* Q of X is such that

$$\mathbb{P}(X_{t+dt} = x' | X_t = x) = \mathbb{I}(x = x') + Q_{x,x'} dt + o(dt)$$

for all $x, x' \in \mathcal{S}$. This can be an infinite matrix, it is generally sparse and its entries describe rates for transitions across states in $\mathcal{S}$. Rows in Q must sum to 0 so that $Q_{x,x'} \geq 0$ for all $x \neq x'$, and $Q_x := Q_{x,x} = -\sum_{x' \in \mathcal{S}: x' \neq x} Q_{x,x'}$.

Hence, jumps in the process X are caused by jobs being routed through stations in the underlying network model. We often say that a state $x' \in \mathcal{S}$ is *accessible* from $x \in \mathcal{S}$, and write $x \xrightarrow{i,j,c} x'$ for its corresponding jump, if x' may be reached from x by means of a class- c job transition between the stations i and j , in the direction $i \rightarrow j$. We further denote

$$\mathcal{T} = \{(i, j, c) \in \{0, \dots, M\}^2 \times \mathcal{C} : p_{i,j}^c > 0\} \quad (1)$$

for the finite set of all *feasible* job transitions in the system, and we remark that the generator Q of X is populated by some positive real-valued rates $\boldsymbol{\lambda} = \{\lambda_{\boldsymbol{\eta}} \in \mathbb{R}_+ : \boldsymbol{\eta} \in \mathcal{T}\}$ that define the intensities for these job routings, with $\lambda_{i,j,c} = \mu_i^c \cdot p_{i,j}^c$ for all $(i, j, c) \in \mathcal{T}$.

Example. In Figure 1 we observe diagrams that illustrate this notation in an open single-class network. On the left, we see 3 stations with different rates, disciplines and server counts. The topology P is such that $|\mathcal{T}| = 5$ and $p_{0,1} = 1 - p_{0,2} \in (0, 1)$, $p_{1,3} = p_{2,3} = p_{3,0} = 1$ ($p_{i,j} = 0$ otherwise). On the right, we find the corresponding job transition rates across the 4 pairs of connected stations. Now, let $\vee$ and $\wedge$ reflect shorthand notation for maxima and minima, respectively. In this single-class example, X monitors counts across the stations s.t. $X_t = (X_t^1, X_t^2, X_t^3) \in \mathcal{S}$ for all $t \geq 0$; also, the generator Q is an infinite matrix with $Q_{x,x'} = \lambda_{i,j} \cdot (K_i \wedge x_i)$ for all pairs $x, x' \in \mathcal{S}$ with associated transition $x \xrightarrow{i,j} x'$, where $K_i, x_i \in \mathbb{N}_0$ denote the number of processors and the queue-length within station $i \geq 0$. We finally have $K_1 = 1, K_2 = \infty$ and $K_3 = 2$; at the virtual source, it holds $K_0 \wedge x_0 = 1$ always. Thus, note that transition rates in X further depend on the queue lengths, and resemble kinetic laws within chemical reaction models [11].

2.1. A hierarchical formulation of queueing systems

Within a hierarchical multilevel formulation, rates in $\boldsymbol{\lambda}$ have a distribution (or *image*) $\mathbb{P}_{\boldsymbol{\lambda}} \equiv \boldsymbol{\lambda}_* \mathbb{P}$ under a reference measure $\mathbb{P}$ on $(\Omega, \mathcal{F})$. We assume this to admit a density $f_{\boldsymbol{\lambda}}$ w.r.t. a base measure that will further induce (by properties of exponential transitions) distributions over the service rates $\boldsymbol{\mu}$ and routing topology. Next, note that a network trajectory over a finite interval is a piecewise deterministic jump process, such that $X \equiv (\boldsymbol{t}, \boldsymbol{x})$ is represented by a sequence of transition times $\boldsymbol{t}$ along with states

FIGURE 1: Left, open bottleneck network with 3 service stations. Shaded circles indicate servers, queueing areas are pictured as empty rectangles. The box is a probabilistic junction for the routing of arrivals. Right, job transition intensities across network stations.

$\mathbf{x}$. Each pair $(\mathbf{t}, \mathbf{x})$ is furthermore a random variable on a measurable space $(\mathcal{X}, \Sigma_{\mathcal{X}})$ supporting finite $\mathcal{S}$ -valued trajectories, and a conditional density $f_{X|\boldsymbol{\lambda}}$ may be defined w.r.t a dominating base measure $\mu_{\mathcal{X}}$, s.t. the regular conditional probability $\mathbb{P}(A|\boldsymbol{\lambda})$, $A \in \mathcal{F}$ satisfies

$$\mathbb{P}(X^{-1}(B)|\boldsymbol{\lambda}) = \int_B f_{X|\boldsymbol{\lambda}}(\mathbf{t}, \mathbf{x}) \mu_{\mathcal{X}}(d\mathbf{t}, d\mathbf{x})$$

for all $B \in \Sigma_{\mathcal{X}}$ (see Appendix A for details). In this case,

$$f_{X|\boldsymbol{\lambda}}(\mathbf{t}, \mathbf{x}) = \pi(x_0) e^{Q_{x_I}(T-t_I)} \prod_{i=1}^I Q_{x_{i-1}, x_i} e^{Q_{x_{i-1}}(t_i - t_{i-1})}, \quad (2)$$

for every pair of ordered times $\mathbf{t} = \{0, t_1, \dots, t_I\}$ in $[0, T]$ and states $\mathbf{x} = \{x_0, \dots, x_I\}$. Here, $\pi(\cdot)$ denotes an arbitrary distribution over initial states, and $Q \equiv Q(\boldsymbol{\lambda})$ is the matrix of infinitesimal rates associated with fixed values in $\boldsymbol{\lambda}$. The queueing network model is thus fully parametrized by a collection of hyper-parameters, and analogue modelling choices for *continuous-time Markov chains* (CTMCs) or MJPs can be found in [13, 2] or [32], to name a few.

2.2. Network evaluation and problem statement

Let $T > 0$ denote some arbitrary terminal time and $\mathbf{x}_0 \in \mathcal{S}$ an initial state in X . For simplicity, this is assumed to be a 0-valued vector, where no jobs populate the system. Now, let $0 \leq t_1 < \dots < t_K \leq T$ denote some fixed network monitoring times along with *observation* variables $\{O_k \in \mathcal{O}, k = 1, \dots, K\}$, for some arbitrary support set $\mathcal{O}$, such that

$$\mathbb{P}\left(\bigcap_{k=1}^K O_k^{-1}(\mathbf{o}_k) | X\right) = \prod_{k=1}^K \mathbb{P}(O_k^{-1}(\mathbf{o}_k) | X) = \prod_{k=1}^K f_{O|X_{t_k}}(\mathbf{o}_k) \quad (3)$$

for any sequence of elements $\mathbf{o}_1, \dots, \mathbf{o}_K$ where $\mathbf{o}_k$ denotes the time- t_k network observation across all stations. Hence, observations are conditionally mutually independent given the network states. The term $f_{O|X_{t_k}}$ stands for a conditional mass function assigned to measurements across the M stations; defined w.r.t a counting measure μ_O . In this paper it is assumed that $f_{O|x} > 0$ (everywhere) for all $x \in \mathcal{S}$, however, deterministic observations such as queue lengths can be easily approximated by means of *regularised* indicator functions. Extensions to continuous settings are straightforward.

Example. Consider a vanilla closed network as pictured in Figure 2. It includes a single FCFS service station, $K = 1$ processing units and a delay routing a population of N jobs in a closed loop. The evolution of $X = (X_t)_{t \geq 0}$ monitors the total number of jobs within the service station, with $X_0 = 0$. Variables $O_k = (O_{0,k}, O_{1,k})$ at times t_k approximate a system with deterministic observations of queue lengths (in both station and delay) and are supported on $\mathcal{O} = \{0, \dots, N\}^2$. The observation model may thus factor across the network components s.t. for example $f_{O|x}(\mathbf{o}) = \tilde{f}_{O|N-x}(o_0) \cdot \tilde{f}_{O|x}(o_1)$ with $\tilde{f}_{O|x}(o) = \frac{\epsilon}{N} + \mathbb{I}(o = x) \cdot (1 - \frac{N+1}{N}\epsilon)$ and

$$\mathbb{P}(O_k^{-1}(\mathbf{o})|X) = \begin{cases} (1 - \epsilon)^2 & o_0 = N - X_{t_k}, o_1 = X_{t_k}, \\ (\epsilon/N)^2 & o_0 \neq N - X_{t_k}, o_1 \neq X_{t_k}, \\ (1 - \epsilon) \cdot \epsilon/N & \text{otherwise.} \end{cases} \quad (4)$$

for a *slack* regularising parameter $\epsilon > 0$ and all $k = 1, \dots, K$. This will account for some $\%100 \cdot \epsilon$ faulty measurements. Finally, note that the framework can generalise in order to accommodate different variants of observation types. For simplicity in the presentation, in this paper we focus on queue lengths.

FIGURE 2: Closed queueing network with a single FCFS service station and a delay.

Now, let $\mathbb{P}(A|\mathbf{o}_1, \dots, \mathbf{o}_K)$, $A \in \mathcal{F}$, $(\mathbf{o}_1, \dots, \mathbf{o}_K) \in \mathcal{O}^K$ denote the regular conditional

probability across global events, conditioned on observations. Our primary interest lies in its induced distribution over the intensity rates (which we denote $\mathbb{P}_{\lambda|\mathbf{o}_1, \dots, \mathbf{o}_K}$) and trajectories. Within (Bayesian) inferential settings, these induced distributions are referred to as a *posteriors*. For simplicity, we will restrict the problem statement to intensity rate posteriors; however, the reader will note that an analogue presentation is easily deduced for network trajectories. The posterior distribution exists and admits a density carried by its corresponding *prior* $\mathbb{P}_\lambda$ (see Appendix A), moreover, the transformation is proportional to a weighted product of network paths, and defined by the Radon-Nikodym derivative

$$\frac{d\mathbb{P}_{\lambda|\mathbf{o}_1, \dots, \mathbf{o}_K}}{d\mathbb{P}_\lambda} = \frac{\int_{\mathcal{X}} [\prod_{k=1}^K \mathbb{P}(O_k^{-1}(\mathbf{o}_k)|\mathbf{t}, \mathbf{x})] f_{X|\lambda}(\mathbf{t}, \mathbf{x}) \mu_{\mathcal{X}}(d\mathbf{t}, d\mathbf{x})}{\mathbb{P}(O_1^{-1}(\mathbf{o}_1) \cap \dots \cap O_K^{-1}(\mathbf{o}_K))}, \quad (5)$$

which corresponds to Bayes' equation. There, the denominator denotes a normalising constant that integrates over trajectories and rates. This transformation will often induce a density representation $f_{\lambda|\mathbf{o}_1, \dots, \mathbf{o}_K}$ for the posterior distribution w.r.t a suitable (Lebesgue) base. In these cases, we may think of the above derivative as a Likelihood-ratio. However, this ratio poses a tractability problem, that is, the integral over trajectories cannot be computed analytically and must be approximated. This is a common problem in inferential tasks with jump processes [cf. 12, 24, 23], and proposed solutions often rely on intensive MCMC procedures that iterate between trajectories and parameters; including direct sampling, rejection sampling or uniformization-based methods. Yet, algorithms are hard to implement, computationally demanding or only applicable to reduced classes of problems. In the case of queueing networks, strong temporal dependencies in the stochastic trajectories X impose hard coupling properties amongst rates and paths [27], which limits the applicability of numerical state of the art solutions to the simplest types of network evaluation problems [22].

Alternatively, complex integrals in (5) may be approximated through a variational approach, where we suitably parametrize an alternative approximating measure to $\mathbb{P}$, that will decompose the integrand term into multiple independent and analytically tractable bits. However, there exists yet no approximating measure design to achieve this goal in light of the complexity of jump processes induced by networks of queues. In the following, we present theoretical results leading to a new variational inferential design for use with jump processes, which significantly deviates from and overcomes

multiple limitations found in standard methods that resort to *full-independence* assumptions [18, 19, 7, 28, 30, 31].

3. Overview of results

Note that under the natural measure $\mathbb{P}$ tied to the infinitesimal generator Q , an underlying MJP X as introduced in Section 2 is supported in a set $\mathcal{S}$ of feasible vectors of integers, which is often just $\mathcal{S} = \mathbb{N}_0^{|C| \times M}$. Further recall that $\mathcal{T}$ in (1) corresponds to possible transition *routes* in the queueing system, that is, triplets $(i, j, c) \in \{0, \dots, M\}$ s.t. a class- c job currently in station i has a non-negative probability $p_{i,j}^c > 0$ of transitioning to station j , under the measure $\mathbb{P}$.

Space expansion. Assume the existence of an approximating measure $\tilde{\mathbb{P}}$ on an augmented space of network paths $\tilde{\mathcal{X}}$, such that we further assign a mass to states with *negative* queue lengths. Under $\tilde{\mathbb{P}}$, rates for transitions across states are induced by a generator $\tilde{Q} = \tilde{Q}(\delta, \lambda)$ with

$$\tilde{Q}_{x,x'} = \delta + Q_{x,x'}, \quad \delta > 0 \quad (6)$$

whenever $x \xrightarrow{i,j,c} x'$ and $(i, j, c) \in \mathcal{T}$; that is, whenever a transition from state x to x' is such that

- a class- c job is serviced at station i and routed to j , and
- $p_{i,j}^c > 0$ within the original probabilistic *routing topology* under $\mathbb{P}$.

In any other case, it holds that $\tilde{Q}_{x,x'} = Q_{x,x'} = 0$. Note that the original connectivity structure of the queueing network is preserved under $\tilde{\mathbb{P}}$; however, jobs may be serviced within queueing stations even if *they do not exist*. Specifically,

- following (6), intensities for job transitions between stations i and j are strictly positive whenever $p_{i,j}^c > 0$. If no class- c job exists within station i , then this intensity corresponds to $\delta > 0$,
- in the event of a t -time class- c job departure from station i when $\lim_{s \rightarrow t} X_s^{i,c} \leq 0$ (the job does not exist), then we assume this job to be *virtually generated* and

$$X_t^{i,c} = \lim_{s \rightarrow t} X_s^{i,c} - 1,$$

i.e., a unit is subtracted from the state vector at the corresponding index,

- this preserves the global job population count, and forces network states to accommodate negative queue lengths.

For values of δ small enough, the $\tilde{\mathbb{P}}$ -density assigned to trajectories outside of $\mathcal{X}$ is negligible. Note that a density $\tilde{f}$ in (2) with generator $\tilde{Q}$ in (6) is such that, for any network path $(\mathbf{t}, \mathbf{x}) \in \tilde{\mathcal{X}} \setminus \mathcal{X}$, it holds

$$\tilde{f}_{X|\lambda}(\mathbf{t}, \mathbf{x}) \leq \prod_{i=1}^I \tilde{Q}_{x_{i-1}, x_i} = O(\delta^r) \quad \text{as } \delta \rightarrow 0$$

for some $r \in \{1, \dots, I\}$. Thus, $X_* \tilde{\mathbb{P}}(\mathcal{X}) = 1 - \int_{\tilde{\mathcal{X}} \setminus \mathcal{X}} \tilde{f}_{X|\lambda} d\tilde{\mu}_{\mathcal{X}} \xrightarrow{\delta \rightarrow 0} 1$, where $\tilde{\mu}$ denotes an appropriately augmented base measure, and the limiting system dynamics under $\tilde{\mathbb{P}}$ will offer a perfect approximation to the original network model. Within the rest of the paper,

- In Section 4, we present a counting process Y over job transitions in the augmented network with generator $\tilde{Q}$ in (6), and parametrize an alternative absolutely continuous approximating measures $\mathbb{Q}$. In Lemma 1, we derive a lower bound to the equivalent log-likelihood for the network measurements.
- Propositions 5.1 and 5.2 within Section 5 inspect the structure of $\mathbb{Q}$ that best approximates the target regular conditional probability. Corollaries 5.1 and 5.2 focus on the rate density $d\mathbb{P}_{\lambda|\sigma_1, \dots, \sigma_K}$ by looking at conjugacy properties and limiting behaviour as $\delta \rightarrow 0$.
- Finally, Section 6 describes applications of our results within inferential tasks, allowing to approximate (image) measures across the various service rates μ , routing probabilities in $\mathcal{P}$ and trajectories X , conditioned on network measurements. This includes comparisons with existing alternative methods.

4. A counting process over job transitions

A network system as introduced in Section 2 gives raise to a multivariate Markov counting process $Y = (Y_t)_{t \geq 0}$ on $(\Omega, \mathcal{F})$, where each indexed $Y_t = (Y_t^\eta)_{\eta \in \mathcal{T}} \in \mathcal{S}^Y$ accounts for job transitions across all classes in $\mathcal{T}$, up to a time $t \geq 0$. That is, each Y_t^η denotes the cumulative count in Y of transitions $x \xrightarrow{\eta} x'$ in X , with $x, x' \in \mathcal{S}$, and $Y_0^\eta = 0$ for all $\eta \in \mathcal{T}$. At a basic level, these are simply non-decreasing counting

processes for job transitions in the directions defined within $\mathcal{T}$. We further note that $|\mathcal{T}|$ is often small, as underlying network topologies impose strict routing mechanisms. The support set $\mathcal{S}^Y$ for the counting process is determined by the connectivity structure amongst the stations. Under the approximating measure $\tilde{\mathbb{P}}$, it holds $\mathcal{S}^Y = \mathbb{N}_0^{|\mathcal{T}|}$, since job transitions may occur even if origin stations have no jobs queueing. Now, let

$$\mathcal{T}_{i,c}^{\leftarrow} = \{\boldsymbol{\eta} \in \mathcal{T} : \eta_2 = i, \eta_3 = c\} \quad \text{and} \quad \mathcal{T}_{i,c}^{\rightarrow} = \{\boldsymbol{\eta} \in \mathcal{T} : \eta_1 = i, \eta_3 = c\}$$

denote the subsets of $\mathcal{T}$ that include class $c \in \mathcal{C}$ job transitions to, and from, the network station $i \in \{0, \dots, M\}$, respectively. Also, recall that $X_t^{i,c}$ denotes the number of class c jobs in station $i > 0$ at time $t \geq 0$, then

$$X_t^{i,c} = \sum_{\boldsymbol{\eta} \in \mathcal{T}_{i,c}^{\leftarrow}} Y_t^{\boldsymbol{\eta}} - \sum_{\boldsymbol{\eta} \in \mathcal{T}_{i,c}^{\rightarrow}} Y_t^{\boldsymbol{\eta}} \quad (7)$$

for all $t \geq 0$, assuming initially empty networked systems. We note that for all $\boldsymbol{\eta} = (i, j, c) \in \mathcal{T}$ it holds $\boldsymbol{\eta} \in \mathcal{T}_j^{\leftarrow}$ and $\boldsymbol{\eta} \in \mathcal{T}_i^{\rightarrow}$. Thus, paths in X and Y differ in that the former is *coupled*, i.e., a job transition in the direction $\boldsymbol{\eta} = (i, j, c)$ is relevant to (and thus is synchronized across) a pair of marginal processes $(X_t^{i,c})_{t \geq 0}, (X_t^{j,c})_{t \geq 0}$; in the latter, this is only relevant to the indexed process $(Y_t^{\boldsymbol{\eta}})_{t \geq 0}$.

In view of (7), we further denote $x_{i,c} = \sum_{\boldsymbol{\eta} \in \mathcal{T}_{i,c}^{\leftarrow}} y_{\boldsymbol{\eta}} - \sum_{\boldsymbol{\eta} \in \mathcal{T}_{i,c}^{\rightarrow}} y_{\boldsymbol{\eta}}$ to the class- c queue-length in station $i > 0$ for any $y \in \mathcal{S}^Y$. Then, the $\tilde{\mathbb{P}}$ -associated infinitesimal generator matrix Ξ of Y is such that $\Xi_{y,y'} \equiv \Xi_{y,\boldsymbol{\eta}} = \delta + \lambda_{\boldsymbol{\eta}} \cdot [\Upsilon(y, \eta_1, \eta_3) \vee 0]$ with a station *load*

$$\Upsilon(y, i, c) = x_{i,c} \cdot \left(\frac{K_i}{\sum_{c' \in \mathcal{C}} x_{i,c'}} \wedge 1 \right) \quad (8)$$

for all jumps $y \xrightarrow{\boldsymbol{\eta}} y'$, $\boldsymbol{\eta} = (i, j, c)$, where the origin station $i > 0$ has PS discipline (here we have set $0/0 = 0$), and

$$\Upsilon(y, i, c) = K_i \wedge x_{i,c} \quad (9)$$

in stations $i > 0$ with FCFS policy within single-class networks. The station *load* $\Upsilon(y, i, c)$ differs from the queue length $x_{i,c}$ in that it accounts for the weighting derived from the service discipline. We further have $\Xi_{y,y'} = \delta + \lambda_{0,j,c}$ for arrivals from virtual stations (in open networks) and $\Xi_{y,y'} = \delta + \lambda_{0,j,c} \cdot (N + \sum_{\boldsymbol{\eta} \in \mathcal{T}_{0,c}^{\leftarrow}} y_{\boldsymbol{\eta}} - \sum_{\boldsymbol{\eta} \in \mathcal{T}_{0,c}^{\rightarrow}} y_{\boldsymbol{\eta}})$ for arrivals from delays, where N denotes the job population in a closed system. Finally, $\Xi_y := \Xi_{y,y} = -\sum_{y' \in \mathcal{S}^Y : y \neq y'} \Xi_{y,y'}$.

4.1. A variational decomposition

The likelihood for observation events in (3) readily transfers to counts Y by means of (7), we thus may write $f_{O|Y_{t_k}}(\mathbf{o}_k) \equiv f_{O|X_{t_k}}(\mathbf{o}_k)$. Under the measure $\tilde{\mathbb{P}}$ network states can have negative values, the likelihood is undefined in such instances. Now, note that piecewise $\mathcal{S}^Y$ -valued trajectories also represent elements $(\mathbf{t}, \mathbf{y})$ in a space $(\mathcal{Y}, \Sigma_Y)$, similar to network paths in X . Let $f_{Y|\boldsymbol{\lambda}, \mathbf{o}_1, \dots, \mathbf{o}_K}$ be a density function, w.r.t. some base measure μ_Y , where for all $B \in \Sigma_Y$ it holds

$$\mathbb{P}(Y^{-1}(B)|\boldsymbol{\lambda}, \mathbf{o}_1, \dots, \mathbf{o}_K) = \int_B f_{Y|\boldsymbol{\lambda}, \mathbf{o}_1, \dots, \mathbf{o}_K} d\mu_Y.$$

It may be shown by properties of conditional distributions that, conditioned on observations, Y is a non-homogeneous semi-Markov process with *hazard* functions

$$\Lambda_{y, y'}(t) = \Xi_{y, y'} \cdot \frac{\mathbb{P}(\bigcap_{k:t_k > t} O_k^{-1}(\mathbf{o}_k) | Y_t = y')}{\mathbb{P}(\bigcap_{k:t_k > t} O_k^{-1}(\mathbf{o}_k) | Y_t = y)} \quad (10)$$

for $y' \neq y$, and $\Lambda_y(t) = -\sum_{y' \neq y} \Lambda_{y, y'}(t)$, s.t.

$$f_{Y|\boldsymbol{\lambda}, \mathbf{o}_1, \dots, \mathbf{o}_K}(\mathbf{t}, \mathbf{y}) = \pi(y_0) e^{\int_{t_I}^T \Lambda_{y_I}(u) du} \prod_{i=1}^I \Lambda_{y_{i-1}, y_i}(t_i) e^{\int_{t_{i-1}}^{t_i} \Lambda_{y_{i-1}}(u) du}.$$

Here, $\Xi \equiv \Xi(\delta, \boldsymbol{\lambda})$ denotes the generator matrix associated with fixed values in $\boldsymbol{\lambda}$. For a deeper look at conditional jump processes we refer the reader to [25, 9]. This conditional counting process is of key importance, however, the structure of rates in (10) poses a trivial analytical impediment. In our approximating effort, we assume the existence of an alternative measure $\mathbb{Q}$ on $(\Omega, \mathcal{F})$. Under this measure, count trajectories in Y are subject to a *full* decomposition, that is, the $\mathbb{Q}$ -law of Y is that of a family of $|\mathcal{T}|$ independent non-homogeneous Poisson counting processes with state-dependent intensity functions $\nu^\eta = (\nu_t^\eta(\cdot))_{t \geq 0}$, for all $\eta \in \mathcal{T}$. Here,

- Intensity rates for jumps $y \xrightarrow{t, \eta} y'$, $t \geq 0$, are independent of $\boldsymbol{\lambda}$, change over time, and are given by $\nu_t^\eta(y_\eta)$.
- Holding rates in Y evolve according to $|\nu_t(Y_t)|$, with $\nu_t(Y_t) = -\sum_{\eta \in \mathcal{T}} \nu_t^\eta(Y_t^\eta)$.
- The state probability of the multivariate process Y factors across the job transition directions, s.t.

$$\mathbb{Q}(Y_t = y) = \prod_{\eta \in \mathcal{T}} \mathbb{Q}(Y_t^\eta = y_\eta)$$

for every $y \in \mathcal{S}^Y$.

Similar full-independence decompositions are often referred to as a variational *mean-field* approximations within Bayesian inferential settings [7]. In order to ensure computational tractability within forthcoming procedures, the intensity functions ν must be bounded from above by some arbitrary constant, s.t. $\nu_t^\eta(y_\eta) \leq \bar{\nu}$ for all $t > 0, \eta \in \mathcal{T}$ and $y \in \mathcal{S}^Y$. We finally note that $\mathbb{Q}$ and $\tilde{\mathbb{P}}$ are equivalent on $\mathcal{F}$, as both assign a positive measure to every marginally-increasing sequence of $\mathbb{N}_0^{|\mathcal{T}|}$ -valued counts.

4.2. Recapitulation and the variational lower bound

It is important to observe that the intensity functions ν^η completely define our alternative measure. In the rest of the paper, we inspect how to fine-tune these functions to ensure that $\mathbb{Q}$ offers a good approximation to the measure $\mathbb{P}$, conditioned on observations in $\mathbf{O}$. First, within the diagram in Figure 3 we offer an outline of the various measures, stochastic processes, probabilistic dependencies and infinitesimal generators introduced to this point. There, we also summarize the primary steps that help construct a variational framework to enable the inferential task. In short:

- The reference measure $\mathbb{P}$ along with the generator $Q = Q(\lambda)$ describes the real *unknown* stochastic network trajectories in X .
- $\tilde{\mathbb{P}}$ is a minor departure from $\mathbb{P}$ and expands the support of X s.t. queue lengths can become negative. It is associated with the generators $\tilde{Q} = \tilde{Q}(\delta, \lambda)$ and $\Xi = \Xi(\delta, \lambda)$, for trajectories in X and job transition counts in Y , respectively.
- $\mathbb{Q}$ is an alternative measure under which Y is a set of independent non-homogeneous counting processes. It is associated with intensity functions $\nu^\eta = (\nu_t^\eta(\cdot))_{t \geq 0}$ for each count process in the directions $\eta \in \mathcal{T}$.

In our next result, a *variational* lower bound will interrelate the transition rates λ for jumps under $\mathbb{P}$ and $\tilde{\mathbb{P}}$ with the intensity functions ν^η , in light of the available observations/measurements in $\mathbf{O}$. From now on, transition rates λ are assumed mutually independent under $\mathbb{Q}$, and admit undetermined densities $d\mathbb{Q}_\lambda = \{d\mathbb{Q}_{\lambda, \eta}, \eta \in \mathcal{T}\}$, that must integrate to 1 on $(\mathbb{R}_+, \mathcal{B}(\mathbb{R}_+))$.

Lemma 1. (Variational lower bound.) *Denote $\mathbf{O} = \bigcap_{k=1}^K O_k^{-1}(\mathbf{o}_k)$ and let $\tilde{\mathbb{P}}$ and $\mathbb{Q}$ be the probability measures on $(\Omega, \mathcal{F})$, as defined above. Recall notation $\Xi_{y, y'} \equiv \Xi_{y, \eta}$*

FIGURE 3: Overview of the approximating variational framework presented in this paper, summarizing the various measures used and primary goals that each step will accomplish.

for jumps $y \xrightarrow{\eta} y'$ with direction η , then

$$\begin{aligned} \log \tilde{\mathbb{P}}(O) &\geq \sum_{k=1}^K \mathbb{E}_{Y_{t_k}}^{\mathbb{Q}} [\log f_{O|Y_{t_k}}(o_k)] - \mathbb{E}_{\lambda}^{\mathbb{Q}} \left[\log \frac{d\mathbb{Q}_{\lambda}}{d\mathbb{P}_{\lambda}} \right] \\ &\quad - \int_0^T \mathbb{E}_{Y_t, \lambda}^{\mathbb{Q}} \left[\sum_{\eta \in \mathcal{T}} \nu_t^{\eta}(Y_t^{\eta}) \log \frac{\nu_t^{\eta}(Y_t^{\eta})}{\Xi_{Y_t, \eta}} - \Xi_{Y_t} + \nu_t(Y_t) \right] dt \end{aligned} \quad (11)$$

offers a lower bound on the $\tilde{\mathbb{P}}$ -probability of retrieved observation events.

Proof. Note that

$$\begin{aligned} \log \tilde{\mathbb{P}}(O) &= \log \int_{\mathcal{Y} \times \mathbb{R}_+^{|\mathcal{T}|}} \mathbb{P}(O|Y) d(Y, \lambda)_* \tilde{\mathbb{P}} = \log \mathbb{E}_{Y, \lambda}^{\mathbb{Q}} \left[\mathbb{P}(O|Y) \frac{d(Y, \lambda)_* \tilde{\mathbb{P}}}{d(Y, \lambda)_* \mathbb{Q}} \right] \\ &\geq \mathbb{E}_Y^{\mathbb{Q}} [\log \mathbb{P}(O|Y)] - \mathbb{E}_{Y, \lambda}^{\mathbb{Q}} \left[\log \frac{d(Y, \lambda)_* \mathbb{Q}}{d(Y, \lambda)_* \tilde{\mathbb{P}}} \right] \end{aligned} \quad (12)$$

where we use Jensen's inequality for finite measures. This is known as a variational lower bound on the log-likelihood, and

$$\mathbb{E}_Y^{\mathbb{Q}} [\log \mathbb{P}(O|Y)] = \mathbb{E}_Y^{\mathbb{Q}} [\log \prod_{k=1}^K f_{O|Y_{t_k}}(o_k)] = \sum_{k=1}^K \mathbb{E}_{Y_{t_k}}^{\mathbb{Q}} [\log f_{O|Y_{t_k}}(o_k)]$$

follows directly from (3). The negative part in (12) is the *Kullback-Leibler* (KL) divergence between image measures of $\mathbb{Q}$ and $\tilde{\mathbb{P}}$. By noting that these share base measures, and Y, λ are independent under $\mathbb{Q}$, it holds

$$\mathbb{E}_{Y, \lambda}^{\mathbb{Q}} \left[\log \frac{d(Y, \lambda)_* \mathbb{Q}}{d(Y, \lambda)_* \tilde{\mathbb{P}}} \right] = \mathbb{E}_{\lambda}^{\mathbb{Q}} \left[\log \frac{d\mathbb{Q}_{\lambda}}{d\mathbb{P}_{\lambda}} \right] + \mathbb{E}_{\lambda}^{\mathbb{Q}} \left[\mathbb{E}_Y^{\mathbb{Q}} \left[\log \frac{g_Y}{f_{Y|\lambda}} \right] \right], \quad (13)$$

where g_Y and $f_{Y|\lambda}$ denote the Y -trajectory densities associated with rates ν^η and $\Xi(\lambda)$, respectively. The last term in (13) is a $\mathbb{Q}$ -average of the KL divergence on Y , where the mean is taken across the infinitesimal transition rates. For a fixed starting $Y_0 \in \mathcal{S}^Y$, the inner expectation is shown in [18] to take the equivalent form

$$\mathbb{E}_Y^\mathbb{Q} \left[\log \frac{g_Y}{f_{Y|\lambda}} \right] = \lim_{R \rightarrow \infty} \sum_{r=0}^{R-1} \mathbb{E}_{Y_{\frac{T_r}{R}}}^\mathbb{Q} \left[\sum_{y \in \mathcal{S}^Y} \mathbb{Q}(Y_{\frac{T(r+1)}{R}} = y | Y_{\frac{T_r}{R}}) \log \frac{\mathbb{Q}(Y_{\frac{T(r+1)}{R}} = y | Y_{\frac{T_r}{R}})}{\tilde{\mathbb{P}}(Y_{\frac{T(r+1)}{R}} = y | Y_{\frac{T_r}{R}}, \lambda)} \right].$$

Note that within an infinitesimal time interval a jump in Y may only happen in one direction within $\mathcal{T}$. With this in mind, we retrieve the limit of a Riemann sum in the interval $[0, T]$, i.e.

$$\begin{aligned} \mathbb{E}_Y^\mathbb{Q} \left[\log \frac{g_Y}{f_{Y|\lambda}} \right] &= \lim_{R \rightarrow \infty} \frac{T}{R} \sum_{r=0}^{R-1} \mathbb{E}_{Y_{\frac{T_r}{R}}}^\mathbb{Q} \left[\sum_{\eta \in \mathcal{T}} \nu_{\frac{T_r}{R}}^\eta(Y_{\frac{T_r}{R}}) \log \frac{\nu_{\frac{T_r}{R}}^\eta(Y_{\frac{T_r}{R}})}{\Xi_{Y_{\frac{T_r}{R}}, \eta}} + \frac{R}{T} \log \frac{1 + \frac{T}{R} \nu_{\frac{T_r}{R}}^\eta(Y_{\frac{T_r}{R}})}{1 + \frac{T}{R} \Xi_{Y_{\frac{T_r}{R}}, \eta}} \right] \\ &= \int_0^T \mathbb{E}_{Y_t}^\mathbb{Q} \left[\sum_{\eta \in \mathcal{T}} \nu_t^\eta(Y_t) \log \frac{\nu_t^\eta(Y_t)}{\Xi_{Y_t, \eta}} - \Xi_{Y_t} + \nu_t(Y_t) \right] dt, \end{aligned}$$

and

$$\mathbb{E}_{Y, \lambda}^\mathbb{Q} \left[\log \frac{d(Y, \lambda)_* \mathbb{Q}}{d(Y, \lambda)_* \tilde{\mathbb{P}}} \right] = \mathbb{E}_\lambda^\mathbb{Q} \left[\log \frac{d\mathbb{Q}_\lambda}{d\mathbb{P}_\lambda} \right] + \int_0^T \mathbb{E}_{Y_t, \lambda}^\mathbb{Q} \left[\sum_{\eta \in \mathcal{T}} \nu_t^\eta(Y_t) \log \frac{\nu_t^\eta(Y_t)}{\tilde{\Xi}_{Y_t, \eta}} - \tilde{\Xi}_{Y_t} + \nu_t(Y_t) \right] dt$$

completes the proof. $\square$

Thus, the lower bound in (11) depends on both the *latent* variables Y and λ , accounting for the various counts and rates. On a basic level, this is built by 3 distinguishable components, that is, the expected log-observations, the *Kullback-Leibler* divergence across service rate densities, and a $\mathbb{Q}$ -weighted divergence across hazard functions and rates, further integrated along the entire network trajectory.

5. A functional representation

The above bound includes the prior rates density $d\mathbb{P}_\lambda$ along with the $\tilde{\mathbb{P}}$ -generator Ξ for the network with negative queues. In addition, we find the unknown $\mathbb{Q}$ -distribution for the count random variables Y_t , along with the hazard functions ν . Hence, by maximising this bound wrt ν , we may derive properties on $\mathbb{Q}$ that allow for the construction of approximating distributions to $d\mathbb{P}_{X|o_1, \dots, o_K}$ and $d\mathbb{P}_{\lambda|o_1, \dots, o_K}$. In this Section, we begin by generalizing work in [18] and present results that (i) accommodate

parameter uncertainty in transition rates and (ii) impose computational restrictions in the resulting iterative system of equations. Later, we move on to inspect posterior rate densities and conjugacy properties as $\delta \rightarrow 0$.

Proposition 5.1. *Let $d\mathbb{Q}_\lambda$ be some valid joint density assigned to the instantaneous rates λ under the approximating variational measure $\mathbb{Q}$. Also, define $Y_t^{\setminus \boldsymbol{\eta}} = \{Y_t^{\boldsymbol{\eta}'} : \boldsymbol{\eta}' \in \mathcal{T} \setminus \{\boldsymbol{\eta}\}\}$. Then, the $\mathbb{Q}$ -dynamics of Y that optimize the lower bound (11) may be parametrized by a system of equations, so that the intensity functions $\nu_t^{\boldsymbol{\eta}}(y) \leq \bar{\nu}$ are given by*

$$\nu_t^{\boldsymbol{\eta}}(y) = \frac{r_t^{\boldsymbol{\eta}}(y+1)}{r_t^{\boldsymbol{\eta}}(y)} e^{\mathbb{E}_{Y_t^{\setminus \boldsymbol{\eta}}, \lambda}^{\mathbb{Q}} [\log \Xi_{Y_t, \boldsymbol{\eta}} | Y_t^{\boldsymbol{\eta}} = y] - k_t^{\boldsymbol{\eta}}(y) / \mathbb{Q}(Y_t^{\boldsymbol{\eta}} = y)} \quad (14)$$

for all $t \in [0, T]$, $\boldsymbol{\eta} \in \mathcal{T}$ and $y \in \mathbb{N}_0$, with $\kappa_t^{\boldsymbol{\eta}}(y) \geq 0$ and

$$\frac{dr_t^{\boldsymbol{\eta}}(y)}{dt} = r_t^{\boldsymbol{\eta}}(y) \mathbb{E}_{Y_t^{\setminus \boldsymbol{\eta}}, \lambda}^{\mathbb{Q}} [\Xi_{Y_t, \boldsymbol{\eta}} | Y_t^{\boldsymbol{\eta}} = y] - \left(1 + \frac{k_t^{\boldsymbol{\eta}}(y)}{\mathbb{Q}(Y_t^{\boldsymbol{\eta}} = y)}\right) r_t^{\boldsymbol{\eta}}(y+1) \frac{e^{\mathbb{E}_{Y_t^{\setminus \boldsymbol{\eta}}, \lambda}^{\mathbb{Q}} [\log \Xi_{Y_t, \boldsymbol{\eta}} | Y_t^{\boldsymbol{\eta}} = y]}}{e^{k_t^{\boldsymbol{\eta}}(y) / \mathbb{Q}(Y_t^{\boldsymbol{\eta}} = y)}} \quad (15)$$

whenever $t \neq t_k$, $k = 1, \dots, K$, and

$$\lim_{t \rightarrow t_k^-} r_t^{\boldsymbol{\eta}}(y) = r_{t_k}^{\boldsymbol{\eta}}(y) \exp \left(\mathbb{E}_{Y_{t_k}^{\setminus \boldsymbol{\eta}}}^{\mathbb{Q}} [\log f_{O|Y_{t_k}}(\mathbf{o}_k) | Y_{t_k}^{\boldsymbol{\eta}} = y] \right) \quad (16)$$

at network observation times. In addition, $\kappa_t^{\boldsymbol{\eta}}(y)(\nu_t^{\boldsymbol{\eta}}(y) - \bar{\nu}) = 0$.

Proof. We identify a stationary point to the Lagrangian associated with this constrained optimization problem, where optimization is w.r.t. the jump rates and the finite dimensional distributions of Y , subject to $\nu_t^{\boldsymbol{\eta}}(y) \leq \bar{\nu}$ and the master equation

$$\frac{d\mathbb{Q}(Y_t^{\boldsymbol{\eta}} = y)}{dt} = \nu_t^{\boldsymbol{\eta}}(y-1) \cdot \mathbb{Q}(Y_t^{\boldsymbol{\eta}} = y-1) - \nu_t^{\boldsymbol{\eta}}(y) \cdot \mathbb{Q}(Y_t^{\boldsymbol{\eta}} = y) \quad (17)$$

for $y \geq 1$, with $d\mathbb{Q}(Y_t^{\boldsymbol{\eta}} = 0) = -\nu_t^{\boldsymbol{\eta}}(0)\mathbb{Q}(Y_t^{\boldsymbol{\eta}} = 0)dt$. Denote by $\phi_t^{\boldsymbol{\eta}}(y) = \mathbb{Q}(Y_t^{\boldsymbol{\eta}} = y)$ the functional representing the marginal $\mathbb{Q}$ -probability of the state Y_t in the direction of $\boldsymbol{\eta}$, for all $y \in \mathbb{N}_0$. In view of (11), the object function may be expressed as the functional

$$\Phi[\phi, \nu, l] = C + \sum_{k=1}^K \mathbb{E}_{Y_{t_k}}^{\mathbb{Q}} [\log f_{O|Y_{t_k}}(\mathbf{o}_k)] - \int_0^T \sum_{\boldsymbol{\eta} \in \mathcal{T}} \mathbb{E}_{Y_t^{\boldsymbol{\eta}}}^{\mathbb{Q}} [\Psi[Y_t^{\boldsymbol{\eta}}, \phi_t^{\boldsymbol{\eta}}(Y_t^{\boldsymbol{\eta}}), \nu_t^{\boldsymbol{\eta}}(Y_t^{\boldsymbol{\eta}}), l_t^{\boldsymbol{\eta}}(Y_t^{\boldsymbol{\eta}})]] dt$$

with

$$\begin{aligned} \Psi[Y_t^\eta, \phi_t^\eta(Y_t^\eta), \nu_t^\eta(Y_t^\eta), l_t^\eta(Y_t^\eta)] &= \nu_t^\eta(Y_t^\eta) \left(\log \nu_t^\eta(Y_t^\eta) - \mathbb{E}_{Y_t^\eta, \lambda}^\mathbb{Q} [\log \Xi_{Y_t, \eta} | Y_t^\eta] - 1 \right) \\ &\quad + \mathbb{E}_{Y_t^\eta, \lambda}^\mathbb{Q} [\Xi_{Y_t, \eta} | Y_t^\eta] - l_t^\eta(Y_t^\eta) \left(\nu_t^\eta(Y_t^\eta) + \frac{d \log \phi_t^\eta(Y_t^\eta)}{dt} \right) \\ &\quad + l_t^\eta(Y_t^\eta + 1) \nu_t^\eta(Y_t^\eta) - \frac{k_t^\eta(Y_t^\eta)}{\phi_t^\eta(Y_t^\eta)} (\bar{\nu} - \nu_t^\eta(Y_t^\eta)), \end{aligned}$$

where $l^\eta = (l_t^\eta(\cdot))_{t \geq 0}$ and $k^\eta = (k_t^\eta(\cdot))_{t \geq 0}$ are multiplier functions that ensure (17) and the complementary inequality on rates are satisfied. Above, the term C includes the remainder bits in the lower bound in (11) that are independent of the finite dimensional distributions of Y under $\mathbb{Q}$. Hence, we obtain the following functional derivatives

$$\begin{aligned} \frac{\delta \Phi}{\delta \phi_t^\eta(y)} &= \sum_{k=1}^K \delta(t - t_k) \mathbb{E}_{Y_t^\eta}^\mathbb{Q} [\log f_{O|Y_t}(\mathbf{o}_k) | Y_t^\eta = y] - \mathbb{E}_{Y_t^\eta, \lambda}^\mathbb{Q} [\Xi_{Y_t, \eta} | Y_t^\eta = y] - \frac{dl_t^\eta(y)}{dt} \\ &\quad - \nu_t^\eta(y) \left(\log \nu_t^\eta(y) - \mathbb{E}_{Y_t^\eta, \lambda}^\mathbb{Q} [\log \Xi_{Y_t, \eta} | Y_t^\eta = y] - 1 - l_t^\eta(y) + l_t^\eta(y + 1) \right), \end{aligned}$$

and

$$\frac{\delta \Phi}{\delta \nu_t^\eta(y)} = -\phi_t^\eta(y) \left(\log \nu_t^\eta(y) - \mathbb{E}_{Y_t^\eta, \lambda}^\mathbb{Q} [\log \Xi_{Y_t, \eta} | Y_t^\eta = y] + l_t^\eta(y + 1) - l_t^\eta(y) \right) - k_t^\eta(y),$$

for all $t \in [0, T]$, $\eta \in \mathcal{T}$ and $y \in \mathbb{N}_0$, to be complemented by the slackness conditions $\kappa_t^\eta(y)(\nu_t^\eta(y) - \bar{\nu}) = 0$, $\kappa_t^\eta(y) \geq 0$ and $\nu_t^\eta(y) \leq \bar{\nu}$. By letting $l_t^\eta(y) = -\log r_t^\eta(y)$ and setting the above expressions to 0, we obtain

$$\begin{aligned} \frac{dr_t^\eta(y)}{dt} &= r_t^\eta(y) \cdot \left(\mathbb{E}_{Y_t^\eta, \lambda}^\mathbb{Q} [\Xi_{Y_t, \eta} | Y_t^\eta = y] - \sum_{k=1}^K \delta(t - t_k) \mathbb{E}_{Y_t^\eta}^\mathbb{Q} [\log f_{O|Y_t}(\mathbf{o}_k) | Y_t^\eta = y] \right) \\ &\quad - \left(1 + \frac{k_t^\eta(y)}{\phi_t^\eta(y)} \right) \cdot r_t^\eta(y + 1) \cdot e^{\mathbb{E}_{Y_t^\eta, \lambda}^\mathbb{Q} [\log \Xi_{Y_t, \eta} | Y_t^\eta = y] - k_t^\eta(y) / \phi_t^\eta(y)} \end{aligned}$$

and

$$\nu_t^\eta(y) = \frac{r_t^\eta(y + 1)}{r_t^\eta(y)} \cdot e^{\mathbb{E}_{Y_t^\eta, \lambda}^\mathbb{Q} [\log \Xi_{Y_t, \eta} | Y_t^\eta = y] - k_t^\eta(y) / \phi_t^\eta(y)}.$$

Observe above that, for fixed values of ϕ and r , if

$$\frac{r_t^\eta(y + 1)}{r_t^\eta(y)} \cdot e^{\mathbb{E}_{Y_t^\eta, \lambda}^\mathbb{Q} [\log \Xi_{Y_t, \eta} | Y_t^\eta = y]} < \bar{\nu}$$

then the complementary slackness conditions imply $k_t^\eta(y) = 0$; otherwise, $\nu_t^\eta(y) = \bar{\nu}$

and

$$k_t^\eta(y) = \phi_t^\eta(y) \cdot \left[\mathbb{E}_{Y_t^\eta, \lambda}^\mathbb{Q} [\log \Xi_{Y_t, \eta} | Y_t^\eta = y] - \log \frac{\bar{\nu} \cdot r_t^\eta(y)}{r_t^\eta(y + 1)} \right] \geq 0,$$

yielding a valid system of equations, leading to (14)-(16) and concluding the proof. $\square$

Corollary 5.1. (Distributed network monitoring.) *Assume that network observations are distributed and independent across the stations, so that*

$$f_{O|Y_{t_k}}(\mathbf{o}_k) = \prod_{i=1}^M f_{O|\{Y_{t_k}^\eta : \eta \in \mathcal{T}_i\}}(\mathbf{o}_k^i)$$

for some conditional mass function $f_{O|\{Y_{t_k}^\eta : \eta \in \mathcal{T}_i\}}$, where $\mathcal{T}_i = (\cup_{c \in \mathcal{C}} \mathcal{T}_{i,c}^{\leftarrow}) \cup (\cup_{c \in \mathcal{C}} \mathcal{T}_{i,c}^{\rightarrow})$ is the set of job transitions relevant to network activity in station $i > 0$, and $\mathbf{o}_k^i$ denotes the time t_k observations across classes in the station. Further assume that $\bar{\nu} = \infty$, so that there exists no bound on intensity rates $\nu_t^\eta(y)$ under $\mathbb{Q}$. Then, the system of equations in Proposition 5.1 reduces to

$$\nu_t^\eta(y) = \frac{r_t^\eta(y+1)}{r_t^\eta(y)} e^{\mathbb{E}_{Y_t \setminus \eta, \lambda}^{\mathbb{Q}} [\log \Xi_{Y_t, \eta} | Y_t^\eta = y]}$$

for all $t \in [0, T]$, $\eta \in \mathcal{T}$ and $y \in \mathbb{N}_0$, with

$$\frac{dr_t^\eta(y)}{dt} = r_t^\eta(y) \mathbb{E}_{Y_t \setminus \eta, \lambda}^{\mathbb{Q}} [\Xi_{Y_t, \eta} | Y_t^\eta = y] - r_t^\eta(y+1) e^{\mathbb{E}_{Y_t \setminus \eta, \lambda}^{\mathbb{Q}} [\log \Xi_{Y_t, \eta} | Y_t^\eta = y]}$$

whenever $t \neq t_k$, $k = 1, \dots, K$, and

$$\lim_{t \rightarrow t_k^-} r_t^\eta(y) = r_{t_k}^\eta(y) e^{\mathbb{E}_{Y_t \setminus \eta}^{\mathbb{Q}} [\log f_{O|\{Y_{t_k}^{\eta'} : \eta' \in \mathcal{T}_{\eta_1}\}}(\mathbf{o}_k^{\eta_1}) | Y_t^\eta = y]} e^{\mathbb{E}_{Y_t \setminus \eta}^{\mathbb{Q}} [\log f_{O|\{Y_{t_k}^{\eta'} : \eta' \in \mathcal{T}_{\eta_2}\}}(\mathbf{o}_k^{\eta_2}) | Y_t^\eta = y]}$$

accounting for observations at origin and departure stations in $\eta \in \mathcal{T}$.

In Proposition 5.1, we have derived a system of equations with iterated dependencies, given a distribution $\mathbb{Q}_\lambda$ for the service rates. These equations can reconstruct the hazard rates $\nu = \nu^\eta = (\nu_t^\eta(\cdot))_{t \geq 0}$ (for each counting process) that best approximate the trajectories of Y under $\tilde{\mathbb{P}}$, given the observations $\mathbf{O}$. Probabilities over the states of Y may ultimately be derived by means of the master equation (17). In Corollary 5.1, we further notice that by simplifying the network observation model, and easing restrictions on rates under the approximating measure $\mathbb{Q}$ we retrieve an analogue result to that previously presented in [18, 7]. However, this is reportedly problematic and can cause a computational bottleneck when reconstructing the jump rates $\nu_t^\eta(y)$, as these may approach infinity at observation times. When presenting our experimental results at the end of the paper, we will rely on the formulae within Proposition 5.1.

Next, we wish to determine the optimal form for $\mathbb{Q}_\lambda$, given a family of hazard rates ν ; s.t. we best approximate the conditional densities $d\mathbb{P}_{\lambda|\mathbf{o}_1, \dots, \mathbf{o}_K}$.

Proposition 5.2. *Let densities for the infinitesimal rates λ be defined w.r.t to the (Lebesgue) product base measure μ_λ , so that $d\mathbb{Q}_\lambda = g_\lambda d\mu_\lambda$ with $g_\lambda = \prod_{\eta \in \mathcal{T}} g_\lambda^\eta$ and marginal densities $g_\lambda^\eta = d\mathbb{Q}_{\lambda, \eta} / d\mu_\lambda$. Also, let $\nu_t^\eta(y)$, $\eta \in \mathcal{T}$, be some (independent) intensity functions assigned to Y under the approximating variational measure $\mathbb{Q}$. Finally, define $\lambda_{\setminus \eta} = \{\lambda_{\eta'} : \eta' \in \mathcal{T} \setminus \{\eta\}\}$ and recall definitions for network station loads Υ in (8) and (9). Then, as $\delta \rightarrow 0$ in (6), the distribution $\mathbb{Q}_\lambda$ that optimizes the lower bound (11) is such that*

$$g_\lambda^\eta(z) \propto e^{\mathbb{E}_{\lambda_{\setminus \eta}}^\mathbb{Q}[\log f_\lambda(z)] - z \cdot \int_0^T \mathbb{E}_{Y_t}^\mathbb{Q}[\Upsilon(Y_t, \eta_1, \eta_3) \vee 0] dt} \cdot z^{\int_0^T \mathbb{E}_{Y_t}^\mathbb{Q}[\nu_t^\eta(Y_t^\eta)] dt}$$

up to a normalizing constant, for $z \in \mathbb{R}_+$ and every $\eta \in \mathcal{T}$.

Proof. We again identify a stationary point to the Lagrangian associated with a constrained optimization problem, w.r.t to arbitrary (positive) densities g_λ^η with $\int_{\mathbb{R}_+} g_\lambda^\eta d\mu_\lambda = 1$. Since $\mathbb{P}_\lambda$ and $\mathbb{Q}_\lambda$ share base measures, the object function can be written as

$$\Phi[g] = C - \sum_{\eta \in \mathcal{T}} \mathbb{E}_{\lambda_\eta}^\mathbb{Q}[\Psi[\lambda_\eta, g_\lambda^\eta]] - \sum_{\eta \in \mathcal{T}} l^\eta \left[\int_{\mathbb{R}_+} g_\lambda^\eta d\mu_\lambda - 1 \right]$$

where $\{l^\eta\}_{\eta \in \mathcal{T}}$ are non-functional Lagrange multipliers, and

$$\Psi[\lambda_\eta, g_\lambda^\eta] = \log g_\lambda^\eta - \frac{1}{|\mathcal{T}|} \mathbb{E}_{\lambda_{\setminus \eta}}^\mathbb{Q}[\log f_\lambda] + \int_0^T \mathbb{E}_{Y_t}^\mathbb{Q} \left[\nu_t^\eta(Y_t^\eta) \log \frac{\nu_t^\eta(Y_t^\eta)}{\Xi_{Y_t, \eta}} + \Xi_{Y_t, \eta} - \nu_t^\eta(Y_t) \right] dt.$$

The term C includes the remainder bits in the lower bound in (11) that are independent of the rates λ . It follows that

$$\frac{\delta \Phi}{\delta g_\lambda^\eta} = \mathbb{E}_{\lambda_{\setminus \eta}}^\mathbb{Q}[\log f_\lambda] - \log(g_\lambda^\eta) - 1 - l^\eta - \int_0^T \mathbb{E}_{Y_t}^\mathbb{Q} \left[\nu_t^\eta(Y_t^\eta) \log \frac{\nu_t^\eta(Y_t^\eta)}{\Xi_{Y_t, \eta}} + \Xi_{Y_t, \eta} - \nu_t^\eta(Y_t) \right] dt$$

in its support set $\mathbb{R}_+$, for all $\eta \in \mathcal{T}$. By equating the above to 0, considering constraints and analysing the relevant terms up to proportionality, we note that

$$g_\lambda^\eta \propto \exp \left(\mathbb{E}_{\lambda_{\setminus \eta}}^\mathbb{Q}[\log f_\lambda] + \int_0^T \mathbb{E}_{Y_t}^\mathbb{Q} \left[\nu_t^\eta(Y_t^\eta) \log \Xi_{Y_t, \eta} - \Xi_{Y_t, \eta} \right] dt \right),$$

so that

$$g_\lambda^\eta(z) \propto e^{\mathbb{E}_{\lambda_{\setminus \eta}}^\mathbb{Q}[\log f_\lambda(z)] - \int_0^T z \cdot \mathbb{E}_{Y_t}^\mathbb{Q}[\Upsilon(Y_t, \eta_1, \eta_3) \vee 0] dt + \int_0^T \mathbb{E}_{Y_t}^\mathbb{Q}[\nu_t^\eta(Y_t^\eta) \log(\delta + z \cdot \Upsilon(Y_t, \eta_1, \eta_3))] dt}$$

and

$$g_{\lambda}^{\eta}(z) \propto e^{\mathbb{E}_{\lambda \setminus \eta}^{\mathbb{Q}}[\log f_{\lambda}(z)] - z \cdot \int_0^T \mathbb{E}_{Y_t}^{\mathbb{Q}}[\Upsilon(Y_t, \eta_1, \eta_3) \vee 0] dt} \cdot z^{\int_0^T \mathbb{E}_{Y_t}^{\mathbb{Q}}[\nu_t^{\eta}(Y_t^{\eta})] dt}$$

as $\delta \rightarrow 0$, for $z \in \mathbb{R}_+$ and every $\eta \in \mathcal{T}$. $\square$

Corollary 5.2. (Conjugate prior.) *Assume that the prior density on λ also factors across the individual rates, s.t $d\mathbb{P}_{\lambda} = \prod_{\eta \in \mathcal{T}} f_{\lambda}^{\eta} d\mu_{\lambda}$, where f_{λ}^{η} for $\eta \in \mathcal{T}$ denote Gamma density functions with shape α_{η} and rate β_{η} . Then, as $\delta \rightarrow 0$ in (6), these are conjugate priors and*

$$\lambda_{\eta} \stackrel{\mathbb{Q}}{\sim} \Gamma\left(\alpha_{\eta} + \int_0^T \mathbb{E}_{Y_t}^{\mathbb{Q}}[\nu_t^{\eta}(Y_t^{\eta})] dt, \beta_{\eta} + \int_0^T \mathbb{E}_{Y_t}^{\mathbb{Q}}[\Upsilon(Y_t, \eta_1, \eta_3) \vee 0] dt\right),$$

for all $\eta \in \mathcal{T}$.

Hence, as the network model with negative queues under $\tilde{\mathbb{P}}$ offers a better approximation of its original counterpart ($\delta \rightarrow 0$), we may numerically approximate posterior distributions across the infinitesimal rates in λ , under the variational measure $\mathbb{Q}$. In the special case with independent Gamma prior densities, this is an easily interpretable posterior where the shape and rate parameters depend, respectively, on the integrated expected jump intensities and the integrated expected station *loads* in (8) and (9).

6. Applications

The results in this paper suggest an iterative approximation procedure to (5) by means of coordinate ascent. Here, we iteratively update the values of the various rates, functions and Lagrange multipliers while evaluating, and assessing convergence, in the lower bound (11) to the log-likelihood. This is a standard approach in variational inference when looking for a (local) maxima [4], and the problem is known to be convex.

Maximising the bound by calibrating the measure $\mathbb{Q}$ will yield an approximation to the regular conditional probability of events in $\mathcal{F}$ under $\tilde{\mathbb{P}}$, conditioned on the observations. This may then be projected over the rates λ or the trajectories in X and Y . This projected densities are valid in order to approximate the conditional distributions of the service rates μ and routing probabilities $\mathcal{P}$ in the original queueing network system, given observations. The final iterative procedure is described here.

1. Input network observations in (3) and assign a (conjugate) Gamma (image) density $d\mathbb{P}_\lambda$ across job transition intensities λ , with shape parameters α_η and a (shared) rate parameter $\beta_\eta = \beta$, $\eta \in \mathcal{T}$.
2. Define a discretization grid of the time interval $[0, T]$, and operate through interpolation within points in the grid.
3. Set an arbitrary density $d\mathbb{Q}_\lambda$. Fix $\kappa_t^\eta(y) = 0, r_t^\eta(y) = 1$ and input (valid) arbitrary starting values $\mathbb{Q}(Y_t^\eta = y)$, for all $t \in [0, T]$, $\eta \in \mathcal{T}$ and $y \in \mathbb{N}_0$.
4. Iterate until convergence:

- 4.1. In each direction $\eta \in \mathcal{T}$, numerically compute $r_t^\eta(y)$ for every $t \in [0, T]$ and $y \in \mathbb{N}_0$, by means of (15)-(16). Then, update intensity and slack functions $\nu_t^\eta(y), \kappa_t^\eta(y)$ with (14), so that $k_t^\eta(y) = 0$ if

$$\frac{r_t^\eta(y+1)}{r_t^\eta(y)} \cdot e^{\mathbb{E}_{Y_t^\eta, \lambda}^{\mathbb{Q}} [\log \Xi_{Y_t, \eta} | Y_t^\eta = y]} < \bar{\nu},$$

and

$$\nu_t^\eta(y) = \bar{\nu}, \quad k_t^\eta(y) = \mathbb{Q}(Y_t^\eta = y) \cdot [\mathbb{E}_{Y_t^\eta, \lambda}^{\mathbb{Q}} [\log \Xi_{Y_t, \eta} | Y_t^\eta = y] - \log \frac{\bar{\nu} \cdot r_t^\eta(y)}{r_t^\eta(y+1)}]$$

otherwise. Renew transient state probabilities in Y by means of the master equation (17), for all $t \in [0, T]$.

- 4.2. Derive expected intensities $\mathbb{E}_{Y_t^\eta}^{\mathbb{Q}} [\nu_t^\eta(Y_t^\eta)]$ and station loads $\mathbb{E}_{Y_t}^{\mathbb{Q}} [\Upsilon(Y_t, \eta_1, \eta_3) \vee 0]$, for all directions $\eta \in \mathcal{T}$ and times $t \in [0, T]$. Update $\mathbb{Q}$ -densities s.t.

$$\lambda_\eta \stackrel{\mathbb{Q}}{\sim} \Gamma\left(\alpha_\eta + \int_0^T \mathbb{E}_{Y_t^\eta}^{\mathbb{Q}} [\nu_t^\eta(Y_t^\eta)] dt, \beta + \int_0^T \mathbb{E}_{Y_t}^{\mathbb{Q}} [\Upsilon(Y_t, \eta_1, \eta_3) \vee 0] dt\right),$$

for all $\eta \in \mathcal{T}$. The likelihood ratio in (5) can be computed at this stage.

- 4.3. Evaluate the lower bound (11), given the current densities and infinitesimal rates under the approximating measure $\mathbb{Q}$. Assess variation in the bound across iterations and establish convergence.
5. Finally, infer the structure of the various service rates and routing probabilities in the queueing network system.

- 5.1. Note that $\mathbb{E}_{Y_t}^{\mathbb{Q}} [\Upsilon(Y_t, \eta_1, \eta_3) \vee 0]$ remains the same across directions $\eta \in \mathcal{T}$ with shared origin station. Since $\mu_i^c = \sum_{\eta \in \mathcal{T}_{i,c}} \lambda_\eta$ it holds

$$\mu_i^c \stackrel{\mathbb{Q}}{\sim} \Gamma\left(|\mathcal{T}_{i,c}| \cdot \alpha_\eta + \sum_{\eta \in \mathcal{T}_{i,c}} \int_0^T \mathbb{E}_{Y_t^\eta}^{\mathbb{Q}} [\nu_t^\eta(Y_t^\eta)] dt, \beta + \int_0^T \mathbb{E}_{Y_t}^{\mathbb{Q}} [\Upsilon(Y_t, \eta_1, \eta_3) \vee 0] dt\right),$$

for all $0 \leq i \leq M, c \in \mathcal{C}$.

- 5.2. Retrieve distributions for routing probabilities by noting that $p_{i,j}^c = \lambda_{i,j,c} / \mu_i^c$ for all $0 \leq i, j \leq M, c \in \mathcal{C}$. This suggests a Dirichlet distribution.

The following examples treat open and closed network models; source code can be found at github.com/IkerPerez/variationalQueues.

6.1. Single class closed network

We begin by analysing the small closed network example previously shown in Figure 2. We recall that this includes one FCFS service station, with $K_1 = 1$ processing unit, along with a delay, together processing a population of N jobs cyclically in a closed loop. All jobs belong to the same class and have equal service rates, we denote by μ_1 the job processing rate within the service station. On completion, a job proceeds to the delay station where it awaits for an exponentially distributed time before being routed back to the queue. We use μ_0 to denote the delay rate; and note that the arrival rate to the queue is directly proportional to the number of jobs at the delay.

Both stations are independent and μ_0 is fixed in order to ensure model identifiability within the service station. In this instance, the network topology is deterministic and trivial, and the evolution of $X = (X_t)_{t \geq 0}$ monitors the total number of jobs within the service station, with $X_0 = 0$. The generator Q of X is finite and s.t.

$$Q = \begin{array}{c} \begin{array}{ccccccc} & 0 & 1 & 2 & \dots & N-2 & N-1 & N \\ \begin{array}{c} 0 \\ 1 \\ \vdots \\ N-1 \\ N \end{array} & \left[\begin{array}{ccccccc} -N\mu_0 & N\mu_0 & 0 & \dots & 0 & 0 & 0 \\ \mu_1 & -N\mu_0 + \mu_0 - \mu_1 & (N-1)\mu_0 & \dots & 0 & 0 & 0 \\ \vdots & \vdots & \vdots & \ddots & \vdots & \vdots & \vdots \\ 0 & 0 & 0 & \dots & \mu_1 & -(\mu_0 + \mu_1) & \mu_0 \\ 0 & 0 & 0 & \dots & 0 & \mu_1 & -\mu_1 \end{array} \right] \end{array},$$

where row and column labels denote the number of jobs in the queueing station. Since μ_0 is fixed, our interest lies in $\lambda \equiv \lambda_{1,0} = \mu_1 \cdot p_{1,0} = \mu_1$. We assign to this rate a distribution $\mathbb{P}_\lambda \equiv \lambda_\star \mathbb{P}$ with (Gamma) density f_λ such that its hyperparameters fix some reasonably uninformative *prior* knowledge on the system. We monitor the delay and FCFS service station at fixed and equally spaced times $t_1 < \dots < t_K$ in an interval $[0, T]$. Here, variables O_k are supported on $\mathcal{O} = \{0, \dots, N\}^2$ and the observation

model factors across the network components s.t. $f_{O|x}(\mathbf{o}) = \tilde{f}_{O|N-x}(o_0) \cdot \tilde{f}_{O|x}(o_1)$ with $\tilde{f}_{O|x}(o) = \frac{\epsilon}{N} + \mathbb{I}(o = x) \cdot (1 - \frac{N+1}{N}\epsilon)$ and

$$\mathbb{P}(O_k^{-1}(\mathbf{o})|X) = \begin{cases} (1 - \epsilon)^2 & o_0 = N - X_{t_k}, o_1 = X_{t_k}, \\ (\epsilon/N)^2 & o_0 \neq N - X_{t_k}, o_1 \neq X_{t_k}, \\ (1 - \epsilon) \cdot \epsilon/N & \text{otherwise.} \end{cases} \quad (18)$$

for $\epsilon > 0$ and all $k = 1, \dots, K$. This accounts for some $\%100 \cdot \epsilon$ faulty measurements, also, we note that a system with discrete observations is approximated as $\epsilon \rightarrow 0$. Now, assume there exist some sample observations $\mathbf{o}_1, \dots, \mathbf{o}_K$ from a model realization in this closed network. These can be easily produced from (18) given a trajectory $(X_t)_{t \in [0, T]}$. In order to produce the trajectory from (2), given the service rates, we may employ Gillespie's algorithm [10] or faster *uniformization* based alternatives [24].

Remark. The transformation $d\mathbb{P}_{\lambda|\mathbf{o}_1, \dots, \mathbf{o}_K}/d\mathbb{P}_{\lambda}$ is such that, conditioned on $\mathbf{o}_1, \dots, \mathbf{o}_K$, the distribution over λ admits a density carried by the Lebesgue measure μ_{λ} , so that $d\mathbb{P}_{\lambda|\mathbf{o}_1, \dots, \mathbf{o}_K} = f_{\lambda|\mathbf{o}_1, \dots, \mathbf{o}_K} d\mu_{\lambda}$. In this simple example, numerical MCMC procedures [22] or basic *generator-matrix exponentiations* combined with a forward-backward algorithm [32] can offer such density approximations; however, this is reportedly *very inefficient* when N is large. Moreover, in involved networks/processes for complex applications (see next example), such alternatives are simply *unusable* (i.e. they do not scale).

In the following, we analyse simulated data ($N = 50, \mu_0 = 0.1, \epsilon = 0.2, T = 100, K = 50$) by assigning a conjugate Gamma density to $\lambda \equiv \mu_1$, so that $\lambda \sim \Gamma(\alpha, \beta)$ with $\alpha = 5$ and $\beta = 2$ under the reference measure $\mathbb{P}$. Recall that μ_0 is fixed to ensure model identifiability, and $X_t \in \{0, \dots, N\}$ denotes the number of jobs in the service station at any time $t \geq 0$. For later reference, the stationary distribution of the system is given by

$$\pi_{\mathbb{P}}(x|\lambda) = \lim_{t \rightarrow \infty} \mathbb{P}(X_t = x|\lambda) = \left(\frac{\mu_0}{\lambda}\right)^x \frac{1}{(N-x)!} \Bigg/ \sum_{x=0}^N \left(\frac{\mu_0}{\lambda}\right)^x \frac{1}{(N-x)!},$$

so that, assuming the observations are sufficiently spaced, and that the system has

reached stationarity, it holds

$$\mathbb{P}\left(\bigcap_{k=1}^K O_k^{-1}(\mathbf{o}_k) \mid \lambda\right) \approx \prod_{k=1}^K \sum_{x=0}^N f_{O|x}(\mathbf{o}_k) \pi_{\mathbb{P}}(x \mid \lambda) = \frac{\prod_{k=1}^K \sum_{x=0}^N f_{O|x}(\mathbf{o}_k) \left(\frac{\mu_0}{\lambda}\right)^x \frac{1}{(N-x)!}}{\left(\sum_{x=0}^N \left(\frac{\mu_0}{\lambda}\right)^x \frac{1}{(N-x)!}\right)^K}, \quad (19)$$

where $f_{O|x}(\mathbf{o}_k)$ is as defined in (18). Note that here $|\mathcal{T}| = 2$, and the process $Y = (Y_t^{0,1}, Y_t^{1,0})_{t \geq 0}$ monitors transitions between the delay and service station, in both the directions $0 \rightarrow 1$ and $1 \rightarrow 0$. The lower bound to the log-likelihood in (11) reduces to

$$\log \tilde{\mathbb{P}}(\mathbf{O}) \geq \sum_{k=1}^K \mathbb{E}_{Y_{t_k}}^{\mathbb{Q}} [\log \mathbb{P}(O_k^{-1}(\mathbf{o}_k) \mid Y_{t_k}^{0,1} - Y_{t_k}^{1,0})] - \mathbb{E}_{\lambda}^{\mathbb{Q}} \left[\log \frac{g_{\lambda}^{1,0}}{d\mathbb{P}_{\lambda}} \right] + \int_0^T \mathbb{E}_{Y_t, \lambda}^{\mathbb{Q}} [\Psi[Y_t, \nu_t, \lambda]] dt \quad (20)$$

with

$$\begin{aligned} \Psi[Y_t, \nu_t, \lambda] &= \nu_t^{1,0}(Y_t^{1,0}) + \nu_t^{0,1}(Y_t^{0,1}) - 2\delta - \lambda \cdot \mathbb{I}(Y_{t_k}^{0,1} - Y_{t_k}^{1,0} > 0) - \mu_0 \cdot (N + Y_{t_k}^{1,0} - Y_{t_k}^{0,1}) \\ &\quad - \nu_t^{1,0}(Y_t^{1,0}) \log \frac{\nu_t^{1,0}(Y_t^{1,0})}{\delta + \lambda \cdot \mathbb{I}(Y_{t_k}^{0,1} - Y_{t_k}^{1,0} > 0)} - \nu_t^{0,1}(Y_t^{0,1}) \log \frac{\nu_t^{0,1}(Y_t^{0,1})}{\delta + \mu_0 \cdot (N + Y_{t_k}^{1,0} - Y_{t_k}^{0,1})} \end{aligned}$$

s.t. it contains only two hazard functions in the approximating measure $\mathbb{Q}$, namely $\nu^{0,1}$ and $\nu^{1,0}$. In (20), we again notice that the lower bound is dominated by 3 distinguishable components, i.e. (i) the expected log-observations, (ii) the *Kullback-Leibler* divergence across the service rate density, and (iii) a weighted $\mathbb{P}$ -to- $\mathbb{Q}$ divergence in the expected path likelihood, further integrated along the entire network trajectory.

The differential equations for functionals in (16) reduce to

$$\begin{aligned} \frac{dr_t^{0,1}(y)}{dt} &= r_t^{0,1}(y) (\delta + \mu_0 \cdot \mathbb{E}_{Y_t^{1,0}}^{\mathbb{Q}} [(Y_t^{1,0} - y) \vee 0]) \\ &\quad - \frac{1 + k_t^{0,1}(y)/\mathbb{Q}(Y_t^{0,1} = y)}{e^{k_t^{0,1}(y)/\mathbb{Q}(Y_t^{0,1} = y)}} r_t^{0,1}(y+1) e^{\mathbb{E}_{Y_t^{1,0}}^{\mathbb{Q}} [\log(\delta + \mu_0 \cdot [(Y_t^{1,0} - y) \vee 0])]}, \end{aligned}$$

and

$$\begin{aligned} \frac{dr_t^{1,0}(y)}{dt} &= r_t^{1,0}(y) (\delta + \mathbb{E}_{\lambda}^{\mathbb{Q}} [\lambda] \cdot \mathbb{Q}(Y_t^{0,1} > y)) \\ &\quad - \frac{1 + k_t^{1,0}(y)/\mathbb{Q}(Y_t^{1,0} = y)}{e^{k_t^{1,0}(y)/\mathbb{Q}(Y_t^{1,0} = y)}} r_t^{1,0}(y+1) e^{\mathbb{E}_{Y_t^{0,1}, \lambda}^{\mathbb{Q}} [\log(\delta + \lambda \cdot \mathbb{I}(Y_t^{0,1} > y))]} \end{aligned}$$

In Figure 4 (left) we observe the evolution of the lower bound (20) during the iterative inferential procedure, for a sufficiently small and negligible value of δ . There,

FIGURE 4: Left, evolution of lower bound to the log-likelihood during the inferential procedure. Right, 95% credible intervals and point estimates for the service rate λ under $\mathbb{Q}$; grey corresponds to the proposed method, blue is for an (adapted) traditional variational procedure.

we notice that the procedure has converged to a (local) optima within approximately 13 iterations. On the right hand side of the Figure, we further observe credible intervals and point estimates across iterations for λ , under the approximating measure $\mathbb{Q}$. Displayed in black/gray are the estimations obtained using the proposed method in this paper; in blue, we find approximations extracted by adapting variational procedures in [18], to allow for prior knowledge and conjugacy properties. Next, in Figure 5 (left) we find the $\mathbb{P}$ -prior density for λ with the (final iteration) $\mathbb{Q}$ -posterior superimposed (in gray); along with them, the red/blue densities represent

- posterior density through *Metropolis-Hastings* Markov chain Monte Carlo, by means of strong stationarity assumptions leading to the likelihood function shown in (19),
- the approximate posterior by adapting variational procedures in [18].

On the right hand side we observe the network observations on both the service station and the delay. Delay observations are displayed by subtracting their value from the job population N (thus representing a second measurement on the service station). Whenever both observations match, these are displayed with a large-sized dot. Along with it, we find:

- In gray, a mean-average network trajectory and 95% credible interval for job counts on the service station $X_t = (Y_t^{0,1} - Y_t^{1,0})_{t \geq 0}$, under the approximating

FIGURE 5: Left, prior (light grey) and posterior (dark grey) densities for λ ; along with MCMC (red) and traditional variational density (blue) estimates. The black dot on the horizontal axis represents the original value in the network simulation. Right, network observations along with mean-average network trajectory and 95% credible interval for job counts in the service station; in gray, our proposed method, in blue, existing variational alternative method.

measure $\mathbb{Q}$ and with methods introduced in this paper.

- In blue, a similar confidence interval and mean-average path obtained using benchmark methods in [18].

Noticeably, the average trajectory under our proposed variational technique flows through the most informative observations (thick dots), and the credible interval widens-up to account for some faulty measurements within either network station. On the other hand, traditional variational approaches quickly converge to a local optima, and restrict mean-average dynamics further compressing confidence intervals in regions with noisy data. In the next example, we notice how this poses a problem for traditional methods; that is, within complex and *synchronized* stochastic processes we will fail to obtain sensible estimates for network parameters. Moreover, inference by MCMC/forward-backwards methods in our next example is virtually intractable [cf. 22].

6.2. Multi-class parallel tandems with bottleneck and service priorities

We analyse an open multi-class queueing network as pictured in Figure 6. In this network, there exists two classes ($c = 1, 2$) of jobs that simultaneously transit the system. The first class consists of high priority jobs with low arrival and service

intensity rates. The second class includes low priority jobs with high arrival and service rates. Once a job enters the system, a probabilistic routing junction (pictured as a square within the Figure) sends this job through either a *PS* or *priority-FCFS* tandem; later, it will be serviced within an *infinite* station before leaving the network. In the top processor-sharing tandem, each station has 5 processing units; these will fraction their working capacity as seen in (8), in order simultaneously service all jobs regardless of their class and priority level, however, service rates will differ depending on the job class. On the contrary, the bottom tandem includes two *FCFS* stations with a single processing unit and priority scheduling. Here, low priority jobs are *only* serviced if each station is *fully empty* of any high priority jobs; consequently, station loads in (9) are rewritten s.t.

$$\Upsilon(y, i, 1) = 1 \wedge x_{i,1} \quad \text{and} \quad \Upsilon(y, i, 2) = (1 \wedge x_{i,2}) \cdot \mathbb{I}(x_{i,1} < 1),$$

at stations $i \in \{2, 4\}$ and for any $y \in \mathcal{S}^Y$, where we recall

$$x_{i,c} = \sum_{\boldsymbol{\eta} \in \mathcal{T}_{i,c}^{\leftarrow}} y_{\boldsymbol{\eta}} - \sum_{\boldsymbol{\eta} \in \mathcal{T}_{i,c}^{\rightarrow}} y_{\boldsymbol{\eta}}$$

and thus $x_{2,c} = y_{0,2,c} - y_{2,4,c}$, $x_{4,c} = y_{2,4,c} - y_{4,5,c}$, for $c = 1, 2$. Due to the presence of service priorities, the ordering of jobs within the queue is irrelevant (this is also the case with *random order* disciplines); hence, our inferential framework allows for different service rates assigned to jobs in each class. Finally, the last service station includes an infinite amount of processing units, and processing rates also differ depending on the job class.

FIGURE 6: Open queueing network with one routing junction (pictured as a square) and 5 service stations with varied disciplines and processing rates.

We analyse synthetic data created during a time interval $[0, T]$ ($T = 100$), with

arrival intensities $\mu_0^1 = 0.5, \mu_0^2 = 3$, routing probabilities $p_{0,i}^c = 0.5$, $i, c \in \{1, 2\}$ and service rates as shown in Table 1. We collect a reduced set of *noiseless* and equally spaced observations with $K = 50$; these are essentially snapshots of the full system state across its service stations and job classes, s.t. $\mathcal{O} = \mathbb{N}^{10}$ and the observation density in (3) is defined with

$$f_{O|x}(\mathbf{o}) = \prod_{i=1}^5 \prod_{c=1}^2 \mathbb{I}(x_{i,c} = o_{i,c})$$

for $x \in \mathcal{S}$, where $o_{i,c}$ is an indexed observation in the element $\mathbf{o}$ denoting the class- c queue length at station $i > 0$. Within the inferential procedure, this observation likelihood must be approximated with some regularized variant similar to (18), while taking $\epsilon \rightarrow 0$. Next, we assign conjugate Gamma priors to the various service intensities; in order to ensure identifiability in the problem, arrival rates and routing probabilities are fixed and we focus this inferential task on the various service stations. Hence $\boldsymbol{\lambda} \equiv \{\mu_i^c : c = 1, 2 \text{ and } i = 1, \dots, 5\}$, and we set $\lambda_{\boldsymbol{\eta}} \sim \Gamma(1, 0.3)$ under the reference measure $\mathbb{P}$, for all $\boldsymbol{\eta} \in \mathcal{T}$.

In the following, we omit the cumbersome mathematical details related to this complex model formulation, and we focus on discussing prior choices, calibration of the algorithm, results and method comparisons following the inferential procedure.

Remarks on using MCMC data-augmentation for inference. Transient inference in a stochastic system with priorities is specially challenging, due to the strong dependencies this generates on the queue lengths across the stations and classes. Specifically,

- data-augmentation methods relying on MCMC techniques do not scale [cf. 27, 22], as dependences yield very autocorrelated output chains,
- there exist no analytic product-form distributions to enable approximate inferential methods under assumptions of system stationarity, as discussed in the previous example,
- *generator-matrix* exponentiations with a forward-backward algorithm [32] are simply unscalable to such large multivariate systems.

Source code for the data simulation process may be found at github.com/IkerPerez/variationalQueues.

Remarks on using benchmark variational methods for inference. Note that traditional variational methods [cf. 18, 7] are centred around populations or *lengths* in the individual queues. In this example, *populations* **may not** be factorized under an approximating measure $\mathbb{Q}$, since system jumps are *synchronized*; i.e., a jump *down* in one queue corresponds to a jump *up* in another. As a consequence, pairs of approximating rates under $\mathbb{Q}$ will be interlinked with the same *real* transition rate under $\mathbb{P}$, and derivations such as the lower bound in (11), or equations (14)-(16) are unattainable. For the sake of completeness and comparisons, we adapt existing variational algorithms to the current task; however, we must

- allow a factorization $\mathbb{Q}(X_t = \mathbf{x}) = \prod_i \mathbb{Q}(X_{i,t} = x_i)$, s.t. jobs may be virtually created and removed in any queue; i.e., jobs do not transition a network, they reach and depart servers individually. The full population of jobs in the network is not preserved,
- duplicate intensities for transitions in the real model; that is, we have a rate for (i) a job departing a queue, and (ii) the job arriving at another. Technically, a job could arrive at a new server **before** it departs the previous one; synchronization is lost; point estimates for parameters are averaged across pairs and weighted for network load.

We will see, this leads to drastic performance issues that deem the method unusable.

6.2.1. Algorithm calibration Within our method, prior choices in the system state Y must initially accommodate a strictly positive, albeit not necessarily large, likelihood for low-priority jobs to be serviced at any point in time. Here, we achieve this by means of assigning Poisson process priors to task transition counts in Y ; that is, we first run the master equation (17) with some user-specified constant intensity rates. This creates monotone mean average queue lengths in the service stations, and we ensure they flow aligned to the network observations in every instance.

Also, the presence of strong temporal dependencies will often trigger the approximating rates ν in (14) to become unreasonably large, ultimately deeming the algorithm computationally *unfeasible*. This is a phenomena also observed in [18] or [7], within the context of simpler stochastic dynamics. To ensure computational tractability, we

exploit the capping functionals k as in introduced in Proposition 5.1, and set a global rate cap of $\bar{\nu} = 50$. Furthermore, we run the differential equations for r in (16) in log-form. Specific details can be found within the aforementioned source code.

6.2.2. Results Within the plots in Figure 7, with the exception of the bottom right one, we observe 95% credible intervals for the queue length processes $X_t^{i,c}$ over time, across the various service stations and job classes. There, intervals in dark gray colour relate to high priority jobs, and their corresponding queue length observations are represented by black circles. This information is superimposed over its analogue for low priority jobs, where intervals are coloured in light gray and observations represented by small diamonds. These interval approximations ignore small positive densities that are sometimes assigned to negative queue lengths. Note that this is a consequence of employing counts across job transitions in Y as a basis for inference on X , however, we recall this is a necessity in order to overcome the coupling challenges described in Sections 2 and 4. Overall, we note that the variational flow captures well the collected observations, with some few exceptions in the stations with priority scheduling; hence, it offers a good basis to build approximate estimates for parameters and the likelihood ratio (5).

The bottom right plot in Figure 7 shows an overview of the expected jump intensity $\mathbb{E}_{Y_t^\eta}^\mathbb{Q}[\nu_t^\eta(Y_t^\eta)]$ and station load $\mathbb{E}_{Y_t}^\mathbb{Q}[\Upsilon(Y_t, \eta_1, \eta_3) \vee 0]$ in the direction $\eta = (0, 1, 1)$ at times $t \in [0, T]$. The sharp peaks in the intensities come at observations times, and ensure the process density transits through the observations. Finally, we notice that the expected station load differs from the estimate of the high-priority queue-length in station 1, as this process combines and weights the queue-length across the two priorities according to (8).

Next, we find in Table 1 summary statistics for the posterior service rates under the approximating variational measure $\mathbb{Q}$, along with point estimates by adapting benchmark variational techniques in [18]. There, we observe how the proposed framework allows for us to gain a good overview of the system properties and variability in the processing speed across the various stations; while existing methods are far from offering reasonable approximations to system behaviour (they instead seem to construct an *averaged* estimation of network flow). Noticeably, there exists a few

FIGURE 7: 95% credible intervals for queue lengths across the service stations. Dark (light) gray corresponds to high (low) priority jobs. Also, expected jump intensity and station load in the direction $\eta = (0, 1, 1)$.

significant deviations from real values, within the posterior estimates for high priority service rates in *PS* stations. This is likely due to a combination of sampling variance, high model complexity and the limitations of such approximate variational procedures for transient analyses of stochastic processes.

7. Discussion

In this paper, we have enabled the variational evaluation of approximating measures for partially-observed *coupled* systems of jump stochastic processes, with a focus on mixed systems of queueing networks. We furthermore have presented a flexible approx-

TABLE 1: Summary statistics for posterior service rates in the queueing network in Figure 6.

	Real	O/S	Summary		Quantiles				
			Mean	StDev	2.5%	25%	50%	75%	97.5%
μ_1^1	0.25	0.364	0.307	0.043	0.228	0.276	0.304	0.335	0.397
μ_2^1	1.5	1.242	1.387	0.188	1.043	1.256	1.378	1.508	1.778
μ_3^1	0.25	0.421	0.339	0.049	0.250	0.305	0.337	0.371	0.442
μ_4^1	1.5	1.482	1.635	0.219	1.233	1.482	1.625	1.777	2.093
μ_5^1	0.5	0.837	0.761	0.075	0.622	0.709	0.758	0.810	0.915
μ_1^2	0.5	0.551	0.501	0.041	0.424	0.473	0.499	0.528	0.584
μ_2^2	4.0	3.496	3.740	0.298	3.177	3.534	3.731	3.935	4.346
μ_3^2	0.5	0.568	0.504	0.041	0.425	0.475	0.502	0.531	0.588
μ_4^2	4.0	3.265	3.670	0.296	3.112	3.465	3.661	3.863	4.270
μ_5^2	1.0	0.976	0.984	0.056	0.877	0.946	0.983	1.021	1.097

imate Bayesian framework, capable of overcoming the challenges posed by coupling properties, and applicable in scenarios where existing MCMC or variational solutions are unusable. To achieve this goal, we have built on existing variational theory [see 18, 7], and discussed an alternate optimization procedure with slack variables and inequality constraints that can address computational limitations within existing techniques. Notably, results within this paper contribute to existing Bayesian statistical literature in [27, 29, 22], and first allow for the study of the latent stochastic behaviour across complex mixed network models, by means of an augmented process for interactions in the resources.

Even though the proposed framework relies on an approximated network model as a basis for inference (which ensures the absolute continuity across base measures), and while it further analyses queue-lengths by means of augmented job transitions in the resources, we have shown that we can reliably capture the finite-dimensional posterior distributions of the various marginal stochastic processes, and offer a good overview of the network structure and likely flow of workload. This is important as it can enable the evaluation and uncertainty quantification tasks in several networked systems found

in many application domains, where full data observations may be hard to retrieve. Currently, existing state-of-the-art alternatives rely on strong assumptions leading to stationary analyses of such systems, or use alternate MCMC procedures that reportedly find limitations due to existing computational constraints [27, 22].

Appendix A. Construction

Let $(\Omega, \mathcal{F})$ be a measurable space with the regular conditional probability property; also, let $0 \leq t_1 < \dots < t_K \leq T$ be some fixed observation times, with $T > 0$. In a standard queueing network with M stations, Ω may denote a product set supporting instantaneous rates, trajectories and observations, and $\mathcal{F}$ the corresponding product σ -algebra. The space of rates and observations will consist of trivial Borel algebras and power sets, so that λ is an $(\mathbb{R}_+^n, \mathcal{B}(\mathbb{R}_+^n))$ -valued random variable of rates in the infinitesimal generator matrix Q of X , where $n \in \mathbb{N}$ denotes an arbitrary number determined by the network topology. In addition, $\{O_k : k = 1, \dots, K\}$ corresponds to random measurement variables for the network monitoring activity, each defined on $(\mathcal{O}, \mathcal{P}(\mathcal{O}))$, where $\mathcal{O}$ denotes an arbitrary countable support set for observations in every service station. A network trajectory $X = (X_t)_{0 \leq t \leq T}$ is an $(\mathcal{S}, \mathcal{P}(\mathcal{S}))$ -valued stochastic process with a countably infinite support set $\mathcal{S}$. Note that this is a piecewise deterministic jump-process, so that $X = (\mathbf{t}, \mathbf{x})$ is formed by a sequence of transition times $\mathbf{t}$ along with states $\mathbf{x}$. Every pair $(\mathbf{t}, \mathbf{x})$ can be further defined as a random variable on a measurable space $(\mathcal{X}, \Sigma_{\mathcal{X}})$, with $\mathcal{X} = \cup_{i=0}^{\infty} ([0, T] \times \mathcal{S})^i$ and the corresponding union σ -algebra $\Sigma_{\mathcal{X}}$. This space can support all finite $\mathcal{S}$ -valued trajectories and allows the assignment of a dominating base measure $\mu_{\mathcal{X}}$ w.r.t which define a trajectory density. For details, we refer the reader to [9].

Let $\mathbb{P}$ be a reference probability measure on $(\Omega, \mathcal{F})$. For all $A \in \mathcal{B}(\mathbb{R}_+^n)$, we write

$$\mathbb{P}(\lambda^{-1}(A)) = \mathbb{P}_{\lambda}(A) = \int_A f_{\lambda}(\mathbf{a}) \mu_{\mathbb{R}_+^n}(d\mathbf{a}),$$

where f_{λ} denotes the joint density function of n independent Gamma distributed variables. Hence, we assume that the distribution of instantaneous rates under $\mathbb{P}$ admits a density carried by the (Lebesgue) measure $\mu_{\mathbb{R}_+^n}$. Next, let $\kappa_1 : \mathcal{F} \times \mathbb{R}_+^n \rightarrow [0, 1]$ be a regular conditional probability; i.e., a Markov kernel that defines a probability

measure on $\mathcal{F}$ for all $\boldsymbol{\lambda} \in \mathbb{R}_+^n$, with

$$\mathbb{P}(B \cap \boldsymbol{\lambda}^{-1}(A)) = \int_A \kappa_1(B, \mathbf{a}) f_{\boldsymbol{\lambda}}(\mathbf{a}) \mu_{\mathbb{R}_+^n}(d\mathbf{a})$$

for $A \in \mathcal{B}(\mathbb{R}_+^n)$ and $B \in \mathcal{F}$. By definition, $\kappa_1(B, \mathbf{a}) = \mathbb{P}(B | \boldsymbol{\lambda} = \mathbf{a})$ and most importantly

$$\kappa_1(X^{-1}(C), \mathbf{a}) = \int_C f_{X|\boldsymbol{\lambda}=\mathbf{a}}(\mathbf{t}, \mathbf{x}) \mu_{\mathcal{X}}(d\mathbf{t}, d\mathbf{x})$$

for all $C \in \Sigma_{\mathcal{X}}$ (note this often poses an intractable integral). The conditional density $f_{X|\boldsymbol{\lambda}=\mathbf{a}}$ is such that for every $I \in \mathbb{N}$ and pair of ordered times $\mathbf{t} = \{0, t_1, \dots, t_I\}$ in $[0, T]$ and states $\mathbf{x} = \{x_0, \dots, x_I\}$ in $\mathcal{S}$ we have

$$f_{X|\boldsymbol{\lambda}=\mathbf{a}}(\mathbf{t}, \mathbf{x}) = \pi(x_0) e^{Q_{x_I}(T-t_I)} \prod_{i=1}^I Q_{x_{i-1}, x_i} e^{Q_{x_{i-1}}(t_i-t_{i-1})},$$

where $Q \equiv Q(\mathbf{a})$ is the matrix of infinitesimal transition rates in X associated to values in $\mathbf{a}$. Finally, network observations are assumed to be discrete events, independent of transition rates given a trajectory. Thus, there exists a kernel $\kappa_2 : \mathcal{F} \times (\mathcal{X} \times \mathbb{R}_+^n) \rightarrow [0, 1]$ s.t.

$$\mathbb{P}(O_k \in D | X = (\mathbf{t}, \mathbf{x}), \boldsymbol{\lambda} = \mathbf{a}) = \kappa_2(O_k^{-1}(D), (\mathbf{t}, \mathbf{x}), \mathbf{a}) = \sum_{d \in D} f_{O_k|(\mathbf{t}, \mathbf{x})}(d) \mu_{\mathcal{O}}(d)$$

for all $k = 1, \dots, K$ and $D \in \mathcal{P}(\mathcal{O})$. Here, $f_{O_k|(\mathbf{t}, \mathbf{x})}$ defines an arbitrary probability mass function on $\mathcal{O}$ carried by a counting measure; in our applications, each observation only depends on the state of the system at the observation time, so the above expression could be further simplified.

Under the above model construction, the support over infinitesimal rates is a standard Borel space and the existence of a posterior distribution is guaranteed [cf. 20]. Also, measures induced by the kernel κ_2 are σ -finite and such that $\kappa_2(\cdot, (\mathbf{t}, \mathbf{x}), \mathbf{a}) < \mu_{\mathcal{O}}$, for every $((\mathbf{t}, \mathbf{x}), \mathbf{a}) \in \mathcal{X} \times \mathbb{R}_+^n$. The posterior is thus carried by its corresponding prior and defined by means of the Radon-Nikodym derivative

$$\frac{d\mathbb{P}_{\boldsymbol{\lambda}|O_1=o_1, \dots, O_K=o_K}}{d\mathbb{P}_{\boldsymbol{\lambda}}}(\mathbf{a}) = \frac{\int_{\mathcal{X}} \prod_{k=1}^K f_{O_k|(\mathbf{t}, \mathbf{x})}(o_k) f_{X|\boldsymbol{\lambda}=\mathbf{a}}(\mathbf{t}, \mathbf{x}) \mu_{\mathcal{X}}(d\mathbf{t}, d\mathbf{x})}{\int_{\mathbb{R}_+^n} \int_{\mathcal{X}} \prod_{k=1}^K f_{O_k|(\mathbf{t}, \mathbf{x})}(o_k) f_{X|\boldsymbol{\lambda}=\mathbf{a}}(\mathbf{t}, \mathbf{x}) \mu_{\mathcal{X}}(d\mathbf{t}, d\mathbf{x}) f_{\boldsymbol{\lambda}}(\mathbf{a}) \mu_{\mathbb{R}_+^n}(d\mathbf{a})},$$

where we employ the shorthand notation $d\mathbb{P}_{\boldsymbol{\lambda}|\cdot}(\mathbf{a}) = \mathbb{P}_{\boldsymbol{\lambda}}(d\mathbf{a}|\cdot)$.

Acknowledgements

The authors would like to thank the anonymous reviewers who greatly improved the quality of this paper. Work supported by RCUK through the Horizon Digital Economy Research grants (EP/G065802/1, EP/M000877/1).

References

- [1] ARMERO, C. AND BAYARRI, M. J. (1994). Prior assessments for prediction in queues. *The Statistician* **43**, 139–153.
- [2] BAELE, G., VAN DE PEER, Y. AND VANSTEELENDT, S. (2010). Using non-reversible context-dependent evolutionary models to study substitution patterns in primate non-coding sequences. *Journal of molecular evolution* **71**, 34–50.
- [3] BASKETT, F., CHANDY, K. M., MUNTZ, R. R. AND PALACIOS, F. G. (1975). Open, closed, and mixed networks of queues with different classes of customers. *Journal of the ACM* **22**, 248–260.
- [4] BLEI, D. M., KUCUKELBIR, A. AND MCAULIFFE, J. D. (2017). Variational inference: A review for statisticians. *Journal of the American Statistical Association* **112**, 859–877.
- [5] BOBBIO, A., GRIBAUDO, M. AND TELEK, M. (2008). Analysis of large scale interacting systems by mean field method. In *Quantitative Evaluation of Systems, 2008. QEST'08. Fifth International Conference on*. IEEE. pp. 215–224.
- [6] BUZACOTT, J. A. AND SHANTHIKUMAR, J. G. (1993). *Stochastic models of manufacturing systems* vol. 4. Prentice Hall, New Jersey.
- [7] COHN, I., EL-HAY, T., FRIEDMAN, N. AND KUPFERMAN, R. (2010). Mean field variational approximation for continuous-time bayesian networks. *Journal of Machine Learning Research* **11**, 2745–2783.
- [8] COOPER, R. B. (1981). Queueing theory. In *Proceedings of the ACM'81 conference*. pp. 119–122.
- [9] DALEY, D. J. AND VERE-JONES, D. (2007). *An introduction to the theory of point processes: volume II: general theory and structure*. Springer Science & Business Media.
- [10] GILLESPIE, D. T. (1977). Exact stochastic simulation of coupled chemical reactions. *The journal of physical chemistry* **81**, 2340–2361.
- [11] GOLIGHTLY, A. AND WILKINSON, D. J. (2015). Bayesian inference for markov jump processes with informative observations. *Statistical applications in genetics and molecular biology* **14**, 169–188.
- [12] HOBOLTH, A. AND STONE, E. A. (2009). Simulation from endpoint-conditioned, continuous-time markov chains on a finite state space, with applications to molecular evolution. *The Annals of Applied Statistics* **3**, 1204–1231.

- [13] HUELSENBECK, J. P., BOLLBACK, J. P. AND LEVINE, A. M. (2002). Inferring the root of a phylogenetic tree. *Systematic biology* **51**, 32–43.
- [14] KLEINROCK, L. (1976). *Queueing Systems Vol II: Computer Applications*. Wiley, New York.
- [15] KOOLE, G. AND MANDELBAUM, A. (2002). Queueing models of call centers: An introduction. *Annals of Operations Research* **113**, 41–59.
- [16] KRAFT, S., PACHECO-SANCHEZ, S., CASALE, G. AND DAWSON, S. (2009). Estimating service resource consumption from response time measurements. In *Proceedings of the Fourth International ICST Conference on Performance Evaluation Methodologies and Tools*. p. 48.
- [17] LIU, Z., WYNTER, L., XIA, C. H. AND ZHANG, F. (2006). Parameter inference of queueing models for IT systems using end-to-end measurements. *Performance Evaluation* **63**, 36 – 60.
- [18] OPPER, M. AND SANGUINETTI, G. (2008). Variational inference for markov jump processes. In *Advances in Neural Information Processing Systems*. pp. 1105–1112.
- [19] OPPER, M. AND SANGUINETTI, G. (2010). Learning combinatorial transcriptional dynamics from gene expression data. *Bioinformatics* **26**, 1623–1629.
- [20] ORBANZ, P. AND TEH, Y. W. (2010). *Bayesian Nonparametric Models*. Springer US, Boston, MA. pp. 81–89.
- [21] OSORIO, C. AND BIERLAIRE, M. (2009). An analytic finite capacity queueing network model capturing the propagation of congestion and blocking. *European Journal of Operational Research* **196**, 996 – 1007.
- [22] PEREZ, I., HODGE, D. AND KYPRAIOS, T. (2017). Auxiliary variables for bayesian inference in multi-class queueing networks. *Statistics and Computing*.
- [23] PEREZ, I. AND KYPRAIOS, T. (2019). Scalable bayesian inference for population markov jump processes. *arXiv preprint arXiv:1904.08356*.
- [24] RAO, V. A. AND TEH, Y. W. (2013). Fast MCMC sampling for Markov jump processes and extensions. *Journal of Machine Learning Research* **14**, 3295–3320.
- [25] SERFOZO, R. F. (1972). Conditional poisson processes. *Journal of Applied Probability* **9**, 288–302.
- [26] SPINNER, S., CASALE, G., BROSIG, F. AND KOUNEV, S. (2015). Evaluating approaches to resource demand estimation. *Performance Evaluation* **92**, 51 – 71.
- [27] SUTTON, C. AND JORDAN, M. I. (2011). Bayesian inference for queueing networks and modeling of internet services. *The Annals of Applied Statistics* **5**, 254–282.
- [28] WANG, C., BLEI, D. AND HECKERMAN, D. (2012). Continuous time dynamic topic models. *arXiv preprint arXiv:1206.3298*.

- [29] WANG, W., CASALE, G. AND SUTTON, C. (2016). A bayesian approach to parameter inference in queueing networks. *ACM Transactions on Modeling and Computer Simulation* **27**, 2:1–2:26.
- [30] ZECHNER, C., UNGER, M., PELET, S., PETER, M. AND KOEPPL, H. (2014). Scalable inference of heterogeneous reaction kinetics from pooled single-cell recordings. *Nature methods* **11**, 197–202.
- [31] ZHANG, B., PAN, J. AND RAO, V. A. (2017). Collapsed variational bayes for markov jump processes. In *Advances in Neural Information Processing Systems*. pp. 3749–3757.
- [32] ZHAO, T., WANG, Z., CUMBERWORTH, A., GSPONER, J., DE FREITAS, N., BOUCHARD-CÔTÉ, A. ET AL. (2016). Bayesian analysis of continuous time markov chains with application to phylogenetic modelling. *Bayesian Analysis* **11**, 1203–1237.