

Auxiliary Variables for Bayesian Inference in Multi-Class Queueing Networks *

Iker Perez ^{†1}, David Hodge², and Theodore Kypraios²

¹*Horizon Digital Economy Research, University of Nottingham, Nottingham, UK*

²*School of Mathematical Sciences, University of Nottingham, Nottingham, UK*

This is a post-peer-review, pre-copy/edit version of an article published in Statistics and Computing. The final authenticated version is available online at: <http://dx.doi.org/10.1007/s11222-017-9787-x>.

Abstract

Queueing networks describe complex stochastic systems of both theoretical and practical interest. They provide the means to assess alterations, diagnose poor performance and evaluate robustness across sets of interconnected resources. In the present paper, we focus on the underlying continuous-time Markov chains induced by these networks, and we present a flexible method for drawing parameter inference in multi-class Markovian cases with switching and different service disciplines. The approach is directed towards the inferential problem with missing data, where transition paths of individual tasks among the queues are often unknown. The paper introduces a slice sampling technique with mappings to the measurable space of task transitions between the service stations. This can address time and tractability issues in computational procedures, handle prior system knowledge and overcome common restrictions on service rates across existing inferential frameworks. Finally, the proposed algorithm is validated on synthetic data and applied to a real data set, obtained from a service delivery tasking tool implemented in two university hospitals.

Keywords— Queueing networks, Continuous-time Markov Chains, Uniformization, Markov chain Monte Carlo, Slice Sampler

1 Introduction

Recent literature addressing *queueing networks* (QNs) has aimed to study inferential methods for the estimation of service requirements. These networks offer the means to describe complex stochastic systems through sets of interacting resources, and have found applications in the design

of engineering and computing systems (Kleinrock, 1976), or within call centres (Kooze and Mandelbaum, 2002), factories (Buzacott and Shanthikumar, 1993) and hospitals (Osorio and Bierlaire, 2009). Enabling the understanding of service performance is very important, since it provides quantitative input for the optimal design of interconnected service stations. Yet, drawing inference on parameters is a challenging errand, since in most applications successive network states are never fully observed. Hence, proposed approaches often rely on reduced summaries such as queue lengths, visit counts or response times, and perform inference in different ways, including regression-based estimation procedures, non-linear numerical optimization or maximum likelihood methods. For a recent review on the matter we refer the reader to Spinner et al. (2015) and references therein.

In this paper, we focus on the underlying *continuous-time Markov chains* (CTMCs) induced by general-form open QNs, and we develop a flexible framework for drawing Bayesian inference on parameters that govern these models; in the presence of general patterns of missing data currently only discussed in (Sutton and Jordan, 2011). Statistical computation is very difficult within this family of models, as it involves working with often countably infinite state spaces where observations provide little indirect information. Here, we target multi-class Markovian cases with possible class switching and different service disciplines, where few or no individual job departure times are observed at specific servers. Hence, knowledge is mostly restricted to task arrival and departures times to, and from, the network. A *task* is a collection of *jobs* undertaken at different service stations, and high loads make it virtually impossible to determine the state of the network at any point in time, including the ordering of jobs across multiple queues. We propose an inferential framework that allows the imposition of prior system knowledge and overcomes common restrictions on service rates across popular service *disciplines* in traditional modelling approaches. A key contribution is that we introduce a slice-sampling approach relying on mappings to the measurable space of task transitions across the service stations; this enables studying systems where the transition paths of tasks among the queues is unknown, and leads to an efficient sampler. The approach draws motivation

*Work supported by RCUK through the Horizon Digital Economy Research grants (EP/G065802/1, EP/M000877/1) and The Health Foundation through the Insight 2014 project “Informatics to identify and inform best practice in out of hours secondary care” (7382).

[†]Corresponding author, e-mail: iker.perez@nottingham.ac.uk

from techniques aimed to explore countably infinite state spaces within *Dirichlet mixture models* or *infinite-state hidden Markov models* (Walker, 2007; Van Gael et al., 2008; Kalli et al., 2011), and sits well within a uniformization oriented MCMC scheme for jump processes as presented in Rao and Teh (2013).

Currently, common assumptions in inferential frameworks include the existence of complete data, product-form equilibrium distributions or unique classes with shared *service rates*. However, we often encounter systems where the completion of jobs at individual stations is only occasionally registered. In addition, inference on the basis of balance may in cases be inaccurate; for instance, the existence of equilibrium in service delivery systems with human workers is a strong assumption, since workload is usually externally controlled and arrivals hardly constitute a Poisson process. In addition, there exist concerns regarding the use of steady-state metrics whenever prior knowledge and constraints are imposed on parameters (Armero and Bayarri, 1994); and the use of product-form solutions within popular *BCMP networks* (Baskett et al., 1975) restricts *first come first served* (FCFS) queues to share service distributions over different task classes.

Aiming for flexible inferential methods, Bayesian procedures relying on *Markov Chain Monte Carlo* techniques were first explored in Sutton and Jordan (2011). There, the authors discussed a latent variable model targeting networks where only subsets of transition times are observed; the method was applicable to open QNs and defined through deterministic transformations between the data and independent *service times* across different disciplines. Later, Wang et al. (2016) proposed the use of a Gibbs sampler relying on product-form distributions and queue length observations, and it advanced the study of closed BCMP networks, offering an approximation method for the normalizing constant within the network’s equilibrium distribution. To the best of our knowledge, no further advances exist in the study of exact Monte Carlo inferential frameworks overcoming known restrictions in the study of QNs. Yet, significant progress has been made with sampling techniques and approximate inference methods for continuous-time dynamic systems often modelled as CTMCs or *continuous-time Bayesian networks* (CTBN) (Nodelman et al., 2002; Fan and Shelton, 2008). However, simulating system dynamics conditioned on scarce observations remains a complex task; a review on the efficiency of various methods for this purpose (including direct sampling, rejection sampling and uniformization methods) can be found in Hobolth and Stone (2009).

Recently, authors Rao and Teh (2013) have presented a noteworthy contribution based on the principles of uniformization (Lippman, 1975; Jensen, 1953). Their work explores a class of auxiliary variable MCMC methods allowing for the efficient and exact computation of state evolutions in systems with discrete support (such as Markov jump processes). The framework relies on producing highly depen-

dent time discretizations within subsequent blocked steps in a Gibbs sampler, and is hypothetically applicable to the study of system evolutions within QNs. However, such systems exhibit strong and characteristic temporal dependencies (cf. Sutton and Jordan (2011)), transitions over an infinite set of states, varying specifications of service disciplines and Markovian regimes often subject to switching. Hence, we face major impediments which require elaborate implementations of slice sampling techniques (Neal, 2003). In this work, we describe a method that controls the computational complexity within simulation procedures; for that matter, we employ families of auxiliary variables across steps in a Gibbs sampler targeting network paths. The result is a method that imposes strong restrictions within the vast space of permissible network transitions at each iteration; however, each subsequent step in the sampler allows for significant timing and routing deviations in limited numbers of tasks routed through the network, ensuring convergence to (i) the distribution of network path evolutions across its full space, given the evidence (ii) the posterior distribution of the arrival and service rates. Finally, we present results on both synthetic and real data, obtained from a service delivery tasking tool implemented in two jointly coordinated university hospitals in the United Kingdom.

The rest of the paper is organised as follows. Section 2 describes CTMCs induced by general form QNs, introduces notions of compatibility with observations, and states the problem addressed in the work. In section 3 the principle of uniformization and its application to networks is briefly revised, mappings to task transitions and auxiliary variables are introduced and the proposed sampler is described. Section 4 introduces results for three example networks of varying complexity with both synthetic and real data. Finally, Section 5 offers a brief closing discussion.

2 Queue networks and continuous-time Markov processes

Consider an open Markovian network with M single service stations, a *population set* \mathcal{C} of different task classes and a non-deterministic *network topology* defined by a family of *routing probability matrices* $\mathcal{P} = \{P^c : c \in \mathcal{C}\}$, such that

- $P_{i,j}^c$ denotes the probability of a class $c \in \mathcal{C}$ task immediately moving to station j after completing a job service in station i , for all $1 \leq i, j \leq M$.
- $P_{i,0}^c$ denotes the probability of a class $c \in \mathcal{C}$ task immediately exiting the network after completing a job service in station i , for all $1 \leq i \leq M$.
- $\sum_{j=0}^M P_{i,j}^c = 1$, for all $1 \leq i \leq M, c \in \mathcal{C}$.

Furthermore, let $\lambda_c > 0$ denote external arrival rates for each task class $c \in \mathcal{C}$; and $p_{0,i}^c$ the corresponding probabilities for its first job to enter station i , $1 \leq i \leq M$. Servers in the network are assumed independent and may differ in their queueing discipline. Service times are non-negative,

have constant rates, and vary over servers and classes; we denote them μ_i^c for all $1 \leq i \leq M, c \in \mathcal{C}$. *Switching* is allowed and thus classes are not permanent categorizations; state-dependent service rates are not considered but follow naturally.

In Figure 1 we observe two example networks further examined within Section 4 in this paper. There, shaded circles indicate servers with exponential service rates μ_i^c , all accompanied by corresponding job *queueing* areas pictured as empty rectangles. Together, such server and queue pairs each represent a service station i , $1 \leq i \leq M$. The shaded boxes are probabilistic routing junctions, where task destinations after a job service (or arrival) are determined according to P^c (or p^c). Finally, λ_c show rates for exponential task arrivals from outside the network.




Figure 1: On top, a bottleneck network with 3 servers; bottom, 2 networks routed in a loop with a single entry and exit server.

Under exponential and independence assumptions, there exists an underlying continuous-time Markov process $X = (X_t)_{t \geq 0}$ that describes the system behaviour. Formally, denoting by \mathcal{S} the countably infinite set of possible states in the network, X is a right-continuous stochastic process such that time-indexed variables X_t are defined within a measurable space $(\mathcal{S}, \Sigma_{\mathcal{S}})$, where $\Sigma_{\mathcal{S}}$ stands for the power set of \mathcal{S} . On a basic level, X holds the ordering of jobs in each queue and server, along with their classes and task identifiers; and \mathcal{S} is the multidimensional product of all possible congruent states at every station. The *infinitesimal generator matrix* Q of X is infinite and such that

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

for all $x, x' \in \mathcal{S}$. Elements in the generator describe rates for transitions within states in the chain, in addition $Q_{x,x'} \geq 0$ for all $x \neq x'$, and $Q_x := Q_{x,x} = -\sum_{x' \in \mathcal{S}: x \neq x'} Q_{x,x'}$. Hence, rows in Q sum to 0, and the full rate for a state departure is given by $|Q_x|$, for all $x \in \mathcal{S}$. Note that transition rates are the product between routing probabilities and exponential rates above; for instance,

- $\lambda_c p_{0,i}^c$ is the transition rate among states in \mathcal{S} accounting for a class- c arrival to service station i ,
- $\mu_i^c P_{i,j}^c$ is the transition rate among states in \mathcal{S} accounting for a job of class c serviced at station i immediately transitioning to station j .

2.1 Observations

Let $\Gamma = \{0, \dots, M\}^2 \times \mathbb{N}$ define a *task transition space*. A triplet $\gamma = (i, j, k) \in \Gamma$ denotes a transition for a uniquely identifiable task k , with i and j specifying the departure and entry stations respectively. Note that it is possible to augment Γ in order to include task classes, yet given unique identifiers this information is redundant. In this work, a transition triplet is never fully observed; instead, we define a set of partial observations $\mathcal{O} = \mathcal{O}_1 \cup \mathcal{O}_2 \subset \Sigma_{\Gamma}$, with

$$\begin{aligned} \mathcal{O}_1 &= \{\sigma \in \Sigma_{\Gamma} : \gamma_1 = \gamma'_1, \gamma_3 = \gamma'_3 \text{ and} \\ &\quad \gamma_2, \gamma'_2 > 0, \text{ for all } \gamma, \gamma' \in \sigma\}, \\ \mathcal{O}_2 &= \{\sigma \in \Sigma_{\Gamma} : \gamma_2 = \gamma'_2 = 0, \gamma_3 = \gamma'_3 \text{ and} \\ &\quad \gamma_1, \gamma'_1 > 0, \text{ for all } \gamma, \gamma' \in \sigma\}, \end{aligned}$$

where Σ_{Γ} stands for the power set of Γ .

Definition 2.1. An *observation* in \mathcal{O} is a subset of Γ that contains all permitted task transitions in the network at some specified time $t > 0$, given external information on an arrival, departure or job service.




Figure 2: Sample observations generated by a single task transitioning a bottleneck network with 3 servers. Observations are represented by rectangles. Dots inform us of transition times. Information below the dots specifies the actual task transitions at each step.

In Figure 2 we observe a bottleneck network produce four partial observations as it evolves over time. The network corresponds to that in Figure 1 (top), and observations include a single task arrival, two job services for the task, and a departure immediately after the final service. There, each task transition $(i, j, k) \in \Gamma$ is marked as $k_{i \rightarrow j}$ at its corresponding time point; note that indexes i, j take the value 0 in order to specify an external arrival or a departure. The observations take the form of elements of \mathcal{O} , i.e.

$$1_{i \rightarrow \cdot} = \{\gamma \in \Gamma : \gamma_1 = i, \gamma_2 > 0, \gamma_3 = 1\} \in \mathcal{O}_1,$$

for $i \in \{1, 2, 3\}$, and

$$1_{\cdot \rightarrow 0} = \{\gamma \in \Gamma : \gamma_1 > 0, \gamma_2 = 0, \gamma_3 = 1\} \in \mathcal{O}_2.$$

In this toy example, it is possible to deduce the original path X in the network when considering the available observations along with the topology in Figure 1; including task orderings across all queues and servers at every point in time. However, in real world applications job service observations are often missing or do not exist at all. In this work, only arrivals and departures are assumed to always be available.

2.2 Compatibility

Let $\mathcal{T} : \mathcal{S}^2 \rightarrow \Gamma \cup \emptyset$ define a measurable function, equipped with the corresponding products of discrete algebras, which maps a pair of states $x, x' \in \mathcal{S}$ to its task transition triplet in Γ . For instance,

$$\mathcal{T}(x, x') = (2, 3, 12)$$

should x' be reachable from x by servicing a job for task 12 in server 2 and immediately routing it to queue 3. Note that for this to be possible, a job for task 12 must be in server 2 within x , and the remaining tasks in the system must be distributed and ordered across stations so that there will exist full agreement with x' . If a state x' is not directly reachable from x , then $\mathcal{T}(x, x') = \emptyset$. We note that the pre-image of a triplet in \mathcal{T} is given by a countably infinite set of pairs of network states in \mathcal{S} , unless bounds on the task population are imposed.

Definition 2.2. Fix some terminal time $T > 0$ and let $\{O_{t_r} \in \mathcal{O} : r = 1, \dots, R\}$ be a sequence of observations at times $0 \leq t_1 < \dots < t_R \leq T$. Also, let $Y_{t_r} = \mathcal{T}^{-1}(O_{t_r}) \in \Sigma_{\mathcal{S}}^2$ for all $r = 1, \dots, R$. Then, we say that a process X is *compatible* with an observation O_{t_r} , and we write $X \perp O_{t_r}$ if

$$\lim_{s \nearrow t_r} X_s = y \quad \text{and} \quad X_{t_r} = y',$$

for some pair of network configurations $(y, y') \in Y_{t_r}$. Furthermore, we say that a process X is *fully compatible* with the observations if $X \perp O_{t_r}$ for all $r = 1, \dots, R$.




Figure 3: Example network paths, all compatible with arrival and departure observations for a single task entering and leaving a bottleneck network with three servers.

In Figure 3 we observe task transitions for sample paths X which are compatible with the arrival and departure information as shown in Figure 2. There, notice that the first sequence corresponds to the original path forming the observations. This time, no job services have been retained and there exist infinitely many paths X that could have produced the same output, with varying transition times and task orderings across the different stations. In large networks with multiple tasks and all simultaneously transitioning the system, it is hard to picture the infinite amount of fully compatible paths X , unless large proportions of job services are retrieved.

2.3 Latent network and problem statement

Denote by $x_0 \in \mathcal{S}$ the initial state in X . In this paper, this is assumed to be an *empty* state, where no jobs populate the network. It is however possible to define an initial distribution π over states, s.t. $\pi(x) := \mathbb{P}(X_0 = x)$ for all $x \in \mathcal{S}$. Now, assume we retrieve $K \in \mathbb{N}$ observation sequences $\tilde{\mathbf{O}} = \{\mathbf{O}_k\}_{k=1, \dots, K}$ collected during different realizations $\mathbf{X} = \{X^k\}_{k=1, \dots, K}$ in the network; with

$$\mathbf{O}_k = \{O_{t_r} \in \mathcal{O} : r = 1, \dots, R_k\}$$

at times $0 \leq t_1 < \dots < t_{R_k} \leq T_k$, for $k = 1, \dots, K$.

The likelihood function is proportional to the product of network path densities fully compatible with $\tilde{\mathbf{O}}$, and is thus intractable. A Gibbs sampling approach centred around latent network evolutions is appropriate, iterating between paths and parameters. For that, note that every X is a piecewise-constant process and may be fully characterized by a set of transition times $\mathbf{t} = \{t_1, \dots, t_n\}$ along with network states $\mathbf{x} = \{x_1, \dots, x_n\}$, so that $X \equiv (\mathbf{t}, \mathbf{x})$ with $X_0 = x_0$. To ease notation, denote $\boldsymbol{\theta} \equiv (\mathcal{P}, \mathbf{p}, \boldsymbol{\lambda}, \boldsymbol{\mu})$, where \mathbf{p} is the vector of arrival routing probabilities. Now, let $\delta_{\mathbf{x}}$ be the number of transitions in \mathbf{x} excluding task arrivals and departures. For each $k = 1, \dots, K$, the density of (\mathbf{t}, \mathbf{x}) given \mathbf{O}_k is (up to proportionality) such that

$$\begin{aligned} f_X((\mathbf{t}, \mathbf{x}) | \boldsymbol{\theta}, \mathbf{O}_k, x_0) & \propto f_{\mathcal{O}}(\mathbf{O}_k | (\mathbf{t}, \mathbf{x}), x_0) f_X((\mathbf{t}, \mathbf{x}) | \boldsymbol{\theta}, x_0) \\ & \propto (1 - q)^{\delta_{\mathbf{x}} - \delta_{\mathbf{o}}} \times \mathbb{I}((\mathbf{t}, \mathbf{x}) \perp \mathbf{O} : \mathbf{O} \in \mathbf{O}_k) \\ & \times e^{Q_{x_n}(T_k - t_n)} \prod_{i=1}^n Q_{x_{i-1}, x_i} e^{Q_{x_{i-1}}(t_i - t_{i-1})}, \quad (1) \end{aligned}$$

where q denotes the probability that a job service in X is observed, and $\delta_{\mathbf{o}} \leq \delta_{\mathbf{x}}$ is the corresponding amount of service observations in \mathbf{O}_k . This density is supported in a suitably defined space of finite network evolutions and the term on top is proportional to Bernoulli trials penalizing network paths with unobserved job services. The term below follows from the properties of the minimum of exponentially distributed random variables.

Hence, drawing parameter inference entails the complex task of simulating network configurations from (1), over increasingly large state spaces and with strong conditional

dependencies. In the following, we revise the notion of uniformization and sampling methods for jump processes introduced in Rao and Teh (2013), and we present an auxiliary observation-variable sampler fit for inference in QN models.

3 Uniformization and auxiliary observations

A generative process for sampling X requires alternating between exponentially distributed times and transitions in proportion to rates. Instead, a *uniformization*-based (Lippman, 1975; Jensen, 1953) sampling scheme employs a *dominating rate* Ω and introduces the notion of *virtual transitions*, so that all times are sampled in a blocked step. In Algorithm 1 we observe a uniformization procedure to produce network paths.

Algorithm 1 Uniformization procedure for process X

- 1: Fix a dominating rate $\Omega \geq \max_{x \in \mathcal{S}} |Q_x|$.
- 2: Sample transition times $0 \leq t_1 < \dots < t_m \leq T$ from a homogeneous Poisson process with rate Ω .
- 3: Set initial state $X_0 = x_0$.
- 4: **for** $i \in \{1, \dots, m\}$ **do**
- 5: Sample x_i from $\{x_{i-1}\} \cup \mathcal{X}_{x_{i-1}}$, with

$$\mathcal{X}_{x_{i-1}} = \{x \in \mathcal{S} \setminus \{x_{i-1}\} : \mathcal{T}(x_{i-1}, x) \neq \emptyset\},$$

and probabilities $\pi_{x_{i-1}}$ given by

$$\pi_{x_{i-1}} = \{1 + Q_{x_{i-1}}/\Omega\} \cup \{Q_{x_{i-1},x}/\Omega : x \in \mathcal{X}_{x_{i-1}}\}.$$

- 6: **end for**
 - 7: Return (\mathbf{t}, \mathbf{x}) .
-

A proof of probabilistic equivalence between a generative and uniformized sampling scheme involves comparing the marginal distribution across states at any time $t \geq 0$, and can be found in Hobolth and Stone (2009). A uniformization procedure yields an augmented set of times $\mathbf{t}' = \{t'_1, \dots, t'_m\}$ and states $\mathbf{x}' = \{x'_1, \dots, x'_m\}$ that accounts for both real and virtual transitions in X . Whenever $x_i = x_{i-1}$ we refer to transition i as *virtual* and note that the number of such transitions is dependent on the choice of Ω . Finally, the density function in (1) may be rewritten to include virtual jumps, so that

$$\begin{aligned} f_X((\mathbf{t}', \mathbf{x}') | \Omega, \boldsymbol{\theta}, \mathbf{O}_k, x_0) &\propto (1-q)^{\delta_{\mathbf{x}'} - \delta_{\mathbf{o}}} \times \mathbb{I}((\mathbf{t}', \mathbf{x}') \perp \mathbf{O} : \mathbf{O} \in \mathbf{O}_k) \\ &\times \prod_{i=1}^m Q_{x'_{i-1}, x'_i}^{\mathbb{I}(x'_i \neq x'_{i-1})} (\Omega + Q_{x'_{i-1}})^{\mathbb{I}(x'_i = x'_{i-1})}, \end{aligned}$$

where terms not proportional to $(\mathbf{t}', \mathbf{x}')$ are omitted.

In practice, simulating X only requires considering a limited number of candidate states in each transition; in close relation to the number of service stations. In Figure 4 we

observe a graphical representation of times, states and transition probabilities for a uniformization-based procedure in the single-class bottleneck network in Figure 1 (top). There, we observe only one task from entry to departure, and we notice \mathbf{x}' is unaltered after virtual transitions. Vertical rectangles are divided in proportion to rates for services and arrivals, and infeasible services are hashed in grey (the additional hashed area in the bottom accounts for a strictly positive dominating rate Ω). This determines the probabilities leading to new states at subsequent times, with virtual jumps associated to the sum of all hashed regions. Finally, removing virtual transitions within $(\mathbf{t}', \mathbf{x}')$ yields the desired realization (\mathbf{t}, \mathbf{x}) in X .

3.1 An auxiliary observation-variable sampler

A uniformization oriented approach can enable the construction of a Gibbs sampler targeting the conditional distribution $f_X((\mathbf{t}, \mathbf{x}) | \boldsymbol{\theta}, \mathbf{O}_k, x_0)$. For such purpose, authors Rao and Teh (2013) show it is possible to recycle groups of real transition times within each iteration. The method applies well to many families of Markov jump processes, but it is insufficient in order to tackle complex systems such as QNs due to a quadratic cost on the number of states when producing \mathbf{x} . This is a known problem in discrete-time systems with large state spaces (such as dynamic Bayesian networks or infinite-state hidden Markov models), and proposed solutions include approximate inference methods (Boyer and Koller, 1998; Ng et al., 2002) or the use of slice sampling techniques for exact inference (Van Gael et al., 2008).

However, QNs contain strong serial dependencies, and transitions over an infinite set of states are triggered by a very reduced number of rates; hence, this can render techniques aimed at Dirichlet mixture models (Walker, 2007; Kalli et al., 2011) or hidden Markov models unusable. A viable approach would ideally consider limited divergences in network paths X over subsequent steps in a sampler; yet allowing for considerable deviations in the routing of a reduced set of tasks. Here, we describe a sampling scheme that achieves this goal by employing random auxiliary mappings to the space of task transitions Γ . Intuitively:

- In each iteration we will first produce additional auxiliary evidence, resulting from task transitions within the current trajectory of X .
- This evidence will be used next in order to significantly restrict the explorable range of network paths in the the following sampler iteration.

This approach poses a computationally tractable technique suited for the analysis of system transitions in QNs, and will construct a Markov chain of posterior trajectories over the entire range of paths in full agreement with the original observed evidence, where reasonably distant samples in the chain are statistically unrelated.

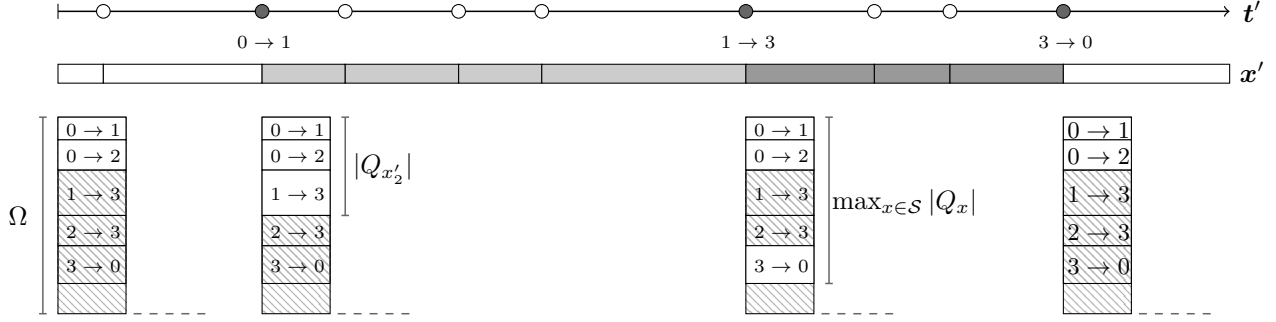


Figure 4: Graphical representation of times \mathbf{t}' , states \mathbf{x}' and transition probabilities for a uniformization-based simulation in a single-class bottleneck network. We observe a single task routed from entry to departure, with virtual transitions represented by empty dots. Vertical rectangles are proportionally split according the likelihood of the various possible services and arrivals.

3.1.1 Preliminaries

Set $\Omega > \max_{x \in \mathcal{S}} |Q_x|$ and let \mathbf{t}' and \mathbf{x}' define some auxiliary frames of transition times and states in X , including both real and virtual values. Arrival, departure and job service observations must come at transition times in \mathbf{t}' ; hence, we may define an augmented set of observations $\mathbf{O}'_k = \{O_{t'_i} \in \mathcal{O} \cup \mathcal{O}_3 : i = 1, \dots, m\}$ at times t'_1, \dots, t'_m , with

$$\mathcal{O}_3 = \{(i, j, k) \in \Gamma : i, j \neq 0\} \cup \{\emptyset\}$$

and such that $\mathbf{O}_k \subseteq \mathbf{O}'_k$. This accounts for missing observations; note that since arrivals and departures are always observed, a missing observation offers evidence for either an inner transition or virtual jump in the network. For simplicity, we assume that no state is reachable from itself in a transition, so that $\mathcal{T}(x, x) = \emptyset$; however, the framework naturally extends to networks where self-transitions are a possibility. Now, denote by \mathbf{u} an auxiliary family of subsets of $\Gamma \cup \emptyset$, such that

$$\mathbb{P}(u_i = u | x'_{i-1}, x'_i) = \begin{cases} p, & \text{if } u = \{\mathcal{T}(x'_{i-1}, x'_i)\}, \\ 1 - p, & \text{if } u = \Gamma \cup \emptyset, \end{cases} \quad (2)$$

with some fixed $p \in [0, 1]$, for all $u_i \in \mathbf{u}$, $i = 1, \dots, m$. Hence, auxiliary variables $u \in \mathbf{u}$ will refer to either the entire space of task transitions or sets with a single element in Γ ; we note that these single element sets will be further contained within a larger observation-set $O \in \mathbf{O}'_k$.

Recall that in queueing networks a task transition may follow from an infinite number of network configurations; that is, there may exist an infinite amount of task orderings across the stations so that a specific job is serviced in one given server and routed to another. However, any network state can only transition to a finite space, by relocating one task in a new queue after a service or an arrival. Thus

$$\{(x, x')\} \subset \mathcal{T}^{-1}\mathcal{T}(x, x') \subset \mathcal{S}^2,$$

strictly, for all $x, x' \in \mathcal{S}$. Moreover, any $u \in \mathbf{u}$ such that $u \neq \Gamma \cup \emptyset$ can only be produced by a limited set of uniformized

paths in X , and compatibility definitions in Definition 2.2 extend naturally to these auxiliary-observation variables. Restrictions are of two types:

- *Transition triplets* impose a transition for an identifiable task. The transition probability is identical over all pairs of compatible states $(x, x') \in \mathcal{S}$.
- *Null sets* impose virtual jumps. The transition probability (lack thereof) depends both on the network configuration and dominating rate Ω .

3.1.2 Sampler

Let (\mathbf{t}, \mathbf{x}) denote a network path in X fully compatible with \mathbf{O}_k , with $\mathbf{t} = \{t_1, \dots, t_n\}$ and $\mathbf{x} = \{x_1, \dots, x_n\}$; then, marginalizing over $(\mathbf{x}', \mathbf{u})$ the frame \mathbf{t}' is independent of any observations and such that (cf. Rao and Teh (2013))

$$f_{\mathbf{t}'}(t'_1, \dots, t'_m | (\mathbf{t}, \mathbf{x}), \Omega, \boldsymbol{\theta}, x_0) = \mathbb{I}(\mathbf{t} \subseteq \{t'_1, \dots, t'_m\}) \times \prod_{i=0}^n (\Omega + Q_{x_i})^{\#\{t'_1, \dots, t'_m\} \cap (t_i, t_{i+1})} e^{-(\Omega + Q_{x_i})(t_{i+1} - t_i)}$$

with $t_0 = 0$ and $t_{n+1} = T$. Thus, it may be sampled in a collapsed step incorporating virtual transitions to times in \mathbf{t} , employing a succession of Poisson processes with rates $\{\Omega + Q_{x_i} : x_i \in \mathbf{x}\}$. Note that \mathbf{t}' along with (\mathbf{t}, \mathbf{x}) induces preliminary sequences of missing observations in \mathbf{O}'_k and uniformized transitions in \mathbf{x}' . Next, we target $\mathbf{u} | \mathbf{x}'$ sampling m auxiliary-observation variables from (2).

Finally, we obtain a new path (\mathbf{t}, \mathbf{x}) in full agreement with both real and auxiliary observations, producing $\mathbf{t}, \mathbf{x}, \mathbf{x}'$ in a blocked step. This simplifies to sampling a sequence $\mathbf{x}' | \mathbf{t}', \mathbf{u}, \Omega, \boldsymbol{\theta}, \mathbf{O}'_k, x_0$ and removing virtual entries; it is achieved by employing dynamic arrays within a procedure for discrete-time state-space models as shown in Algorithm 2. Alternatively, note it is possible to employ a particle filtering approach within a forward procedure, in order to impose further memory constraints.

Algorithm 2 Forward filtering backward sampling with dynamic arrays

1: Set initial state x_0 and let $\alpha_0(x) = \mathbb{I}(x = x_0)$, $x \in \mathcal{S}$.

2: **for** $i \in \{1, \dots, m\}$ **do**

3: **for** $x \in \mathcal{S}$ s.t. $\alpha_{i-1}(x) > 0$ **do**

4: **for** $x' \in \mathcal{S}$ s.t. $|Q_{x,x'}| > 0, (x, x') \in \mathcal{T}^{-1}(O_{t'_i}) \cap \mathcal{T}^{-1}(u_i)$ **do**

5: Update:

$$\alpha_i(x') \leftarrow \alpha_i(x') + (1 - q)^{\mathbb{I}(\mathcal{T}(x,x') \neq \emptyset)} \left(\mathbb{I}(x = x') + \frac{Q_{x,x'}}{\Omega} \right) \alpha_{i-1}(x)$$

6: **end for**

7: **end for**

8: **end for**

9: Sample x_m from $x \in \mathcal{S}$ with probability in proportion to $\alpha_m(x)$.

10: **for** $i \in \{m-1, \dots, 1\}$ **do**

11: **for** $x \in \mathcal{S}$ s.t. $|Q_{x,x_{i+1}}| > 0, \alpha_i(x) > 0, (x, x_{i+1}) \in \mathcal{T}^{-1}(O_{t'_{i+1}}) \cap \mathcal{T}^{-1}(u_{i+1})$ **do**

12: Update:

$$\beta(x) \leftarrow \alpha_i(x)(1 - q)^{\mathbb{I}(\mathcal{T}(x,x_{i+1}) \neq \emptyset)} \left(\mathbb{I}(x = x_{i+1}) + \frac{Q_{x,x_{i+1}}}{\Omega} \right)$$

13: **end for**

14: Sample x_i from $x \in \mathcal{S}$ in proportion to $\beta(x)$.

15: **end for**

16: Remove virtual transitions in (t', x') and return (t, x) .

Forward Filtering

Backward Sampling

3.1.3 Properties and considerations

Along with observations and naturally restrictive constraints on state transitions within QNs, auxiliary variables in \mathbf{u} allow us to limit the space a sampler is allowed to explore within each iteration. These restrictions apply both within forward and backward procedures and leave the underlying filtering equations unaltered, up to proportionality. Increasing the value of p will make computationally expensive iterations less likely, at the cost of a higher dependence between subsequent realizations of X . Also, the term $(1-q)$ enters the forward procedure penalizing network paths with unobserved transitions and is only proportionally relevant when no observation exists.

In Figure 5 we observe a task transition diagram with a single iteration in the proposed sampler, for the bottleneck network in Figure 1 (top). In this example, two tasks (numbered 1 and 2) are observed entering and leaving the network at different times; however, there exists no information regarding job services within the network. In each iteration, the sampler begins with a network path whose task transitions are fully compatible with the existing evidence. In an initial step, the existing path is supplemented with virtual transitions at the corresponding Poisson rates. In the Figure, we observe that nodes for both virtual jumps and the unobserved job services are superimposed over shaded boxes; the boxes represent further evidence for the lack of task arrivals or departures at these times. Next, auxiliary variables are produced across real and virtual jumps, the subsets are loosely represented by ticks ($\Gamma \cup \emptyset$) and crosses ($\{\mathcal{T}(x'_{i-1}, x'_i)\}$) for open and clamped nodes respectively. Then, the uniformized frame is emptied and both real and auxiliary evidence is propagated, imposing task transitions

or virtual jumps within clamped nodes and resulting in a restricted frame for possible network paths. Finally, a new compatible path is sampled via forward filtering backward sampling as summarized in Algorithm 2; this will consider the imposed task transitions and weight successive network states over the clamped epochs. The resulting path is fully compatible with the observed evidence, however, notice that task transitions at arrival or departure times may change between iterations.

Note that by choosing Ω strictly greater than all values in the diagonal of Q , the resulting Markov chain over posterior network transitions is irreducible. Increasing the dominating rate will improve mixing in exchange for higher computational requirements. Finally, we note that a high value of p may hinder the sampler from fully exploring the posterior range of network paths.

3.1.4 Parameter sampling

Finally, given a new family of network realizations $\mathbf{X} = \{X^k\}_{k=1, \dots, K}$ fully compatible with observation sequences $\mathbf{O} = \{O_k\}_{k=1, \dots, K}$, we may obtain posterior samples of arrival and service rate parameters. For traditional FCFS stations this is such that

$$\lambda_c | \mathbf{X} \sim \text{Gamma}(\delta_c, \sum_{k=1}^K T_k)$$

and

$$\mu_i^c | \mathbf{X} \sim \text{Gamma}(\gamma_i^c, \tau_i^c),$$

for $c \in \mathcal{C}, i = 1, \dots, M$; assuming independent network parameters and uninformative priors. Here δ_c, γ_i^c and τ_i^c denote respectively the number of class c arrivals, class c jobs served at station i and the time server i has been occupied

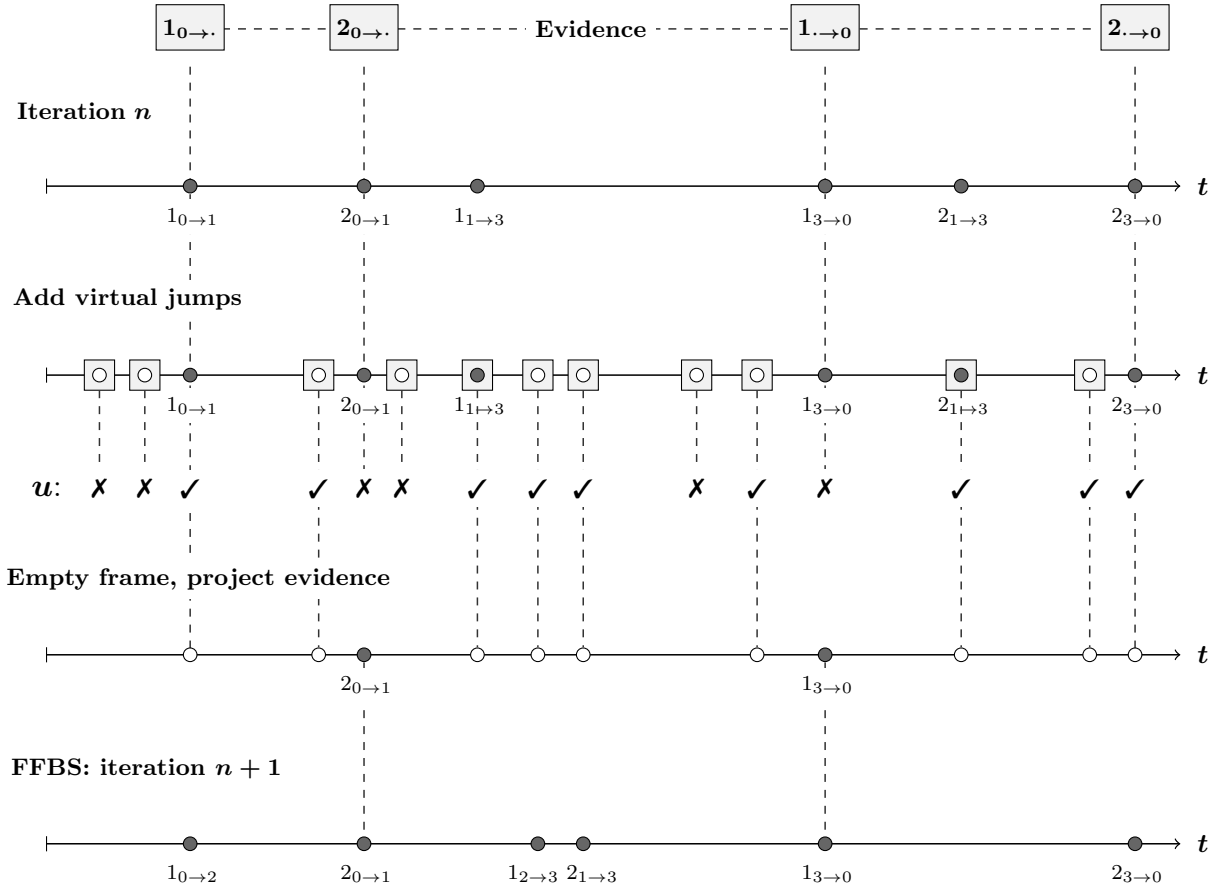


Figure 5: Task transition diagram with a single iteration in the proposed sampler, for a bottleneck network with three servers. Here, tasks 1 and 2 are observed entering and leaving the network. First, start with a path whose task transitions are fully compatible with the evidence. Then, supplement it with virtual jumps at the corresponding rates, and produce auxiliary variables across real and virtual nodes. Next, empty the uniformized frame and propagate both real and auxiliary evidence; imposing task transitions or virtual jumps within clamped nodes. Finally, repopulate the frame via forward filtering backward sampling, hence maintaining agreement with the existing evidence.

by a class c job, in all realizations in \mathbf{X} . Finally, posterior probability vectors for class c routings in every node $i = 1, \dots, M$ are given by

$$P_{i,\cdot}^c | \mathbf{X} \sim \text{Dir}(\mathbf{1} + \boldsymbol{\kappa}_i^c)$$

where $\boldsymbol{\kappa}_i^c$ defines a vector of transition counts from server i in \mathbf{X} . Arrival posteriors in \mathbf{p} are defined the same way. We note that in order to ease identifiability in the inferential problem, it is also possible to fix parameters, incorporate conjugate priors or to impose inequality constraints and bounds across parameters; we will show examples in Section 4 below. Also, the above expressions must be altered when stations respond to prioritization regimes other than FCFS (see Example 3 in Section 4).

4 Examples

In the following, we discuss results obtained across three example networks with both synthetic and real data, in or-

der of increasing difficulty. In all cases, results are obtained through a JAVA implementation of the proposed sampler, and starting compatible network paths have been manually assigned.

The examples demonstrate the ability of the proposed algorithm in order to handle missing data in multi-class inferential problems with varying service disciplines, class switching and imposed prior constraints. Hence, the sampler offers the means to overcome necessary assumptions linked to the common use of product form equilibrium expressions for QNs. To the best of our knowledge, there exists no alternative approach overcoming these restrictions when drawing exact inference in general open Markovian networks.

4.1 Tandem network

In the simplest example, we analyse simulated data on a tandem network with two $M/M/1$ stations, FCFS service




Figure 6: Graphical summary of output for a single-class tandem network fitted to synthetic data. On the left, we observe the posterior distribution of service rate parameters with original values $\mu_1 = 0.2$ and $\mu_2 = 0.5$. Black bars in the center correspond to autocorrelation values. On the right, a contour plot for the joint posterior density of rates (dashed lines represent real values).

disciplines and a single task class. Data is generated so that true service rates are $\mu_1 = 0.2$ and $\mu_2 = 0.5$, arrivals are given by $\lambda = 0.12$ and the network topology is defined by a routing probability matrix P such that $P_{1,2} = 1$ and $P_{2,0} = 1$. Also, jobs enter directly into the first queue and $p_{0,1} = 1$.

For the inferential problem, job service observations (in first station) are always ignored and the only source of information are end-to-end measurements. Thus, available knowledge is limited to the times when tasks enter the queue on the first station and when they depart through the second station. Overall, we examine 5000 realizations totalling 17827 tasks during 115601 time units. For the purpose, the network topology in P is fixed deterministic, since there exists a unique route from start to completion of tasks. Also, in order to ensure identifiability we impose an inequality constraint on service rates and assign fairly uninformative parameter priors, so that

$$\pi(\mu_1, \mu_2) \propto \mathbb{I}(\mu_1 \leq \mu_2) \times \exp(-10^{-3}(\mu_1 + \mu_2)).$$

Note that the problem directly links to the inferential task with two exponentially distributed random variables when only its sum is observed, with the further complexity that unknown waiting times have to be discounted from the empirical observations.

In Figure 6 (right) we observe a contour plot for the joint posterior kernel density estimation over service rates, and we notice a significant negative correlation in values (the dashed vertical and horizontal lines represent the original parameters values in the network path simulations). Results are obtained across two chains with 100000 iterations each, a 10000 burn-in stage, varying starting rates

and different scales for dominating rates and probabilities producing auxiliary-observations, so that $p_1 = 0$ and $\Omega_1 = 2 \max_{x \in \mathcal{S}} |Q_x|$ and $p_2 = 0.25, \Omega_2 = 1.5 \max_{x \in \mathcal{S}} |Q_x|$. Note that the second chain is produced employing restrictive auxiliary-observations as opposed to the first; hence, stronger serial dependencies across subsequent latent paths in the network should be expected. Yet, the remainder plots show marginal posterior kernel density estimations for both service rates, along with an autocorrelation summary across a thinned sample in the second chain, showing a satisfactory mixing.

A discussion on the effects and computational gains resulting from employing restrictive auxiliary observations follows in the next example. In general, networks of interest are complex and $p = 0$ would pose a computationally infeasible problem. Also, even in simple networks such as this example, computing times can be excessive, and considerable reductions can be traded at the cost of higher serial dependences.

4.2 Bottleneck network

We examine simulated data in the bottleneck network in Figure 1 (top), with 3 FCFS stations and 3 different task classes. The true service rates can be observed in Table 1, and task arrivals are given by $\lambda_1 = 0.08, \lambda_2 = 0.06$ and $\lambda_3 = 0.04$. In this case, along with end-to-end measurements, approximately half of all generated job service observations are retrieved so that $q = 0.5$. The network topology is

Table 1: Summary statistics for posterior service rates along with computing times across three chains tuned differently in the bottleneck network in Figure 1 (top).

	Real	Summary		Quantiles				
		Mean	StDev	2.5%	25%	52%	75%	97.5%
μ_1^1	0.3	0.311	0.019	0.273	0.297	0.310	0.323	0.350
μ_1^2	0.25	0.237	0.017	0.204	0.225	0.236	0.247	0.271
μ_1^3	0.2	0.185	0.016	0.154	0.173	0.184	0.196	0.218
μ_2^1	0.7	0.709	0.052	0.613	0.673	0.707	0.743	0.818
μ_2^2	0.5	0.537	0.040	0.462	0.509	0.536	0.564	0.620
μ_2^3	0.3	0.297	0.028	0.245	0.277	0.296	0.315	0.354
μ_3^1	1.5	1.626	0.088	1.458	1.566	1.625	1.684	1.802
μ_3^2	1.2	1.206	0.069	1.074	1.159	1.204	1.252	1.346
μ_3^3	0.8	0.728	0.052	0.630	0.692	0.727	0.762	0.833

	p	Ω	Run Time	ESS:	Mean	Min
1	0.7	$2 \max_{x \in \mathcal{S}} Q_x $	6056.5s		14575	7246
2	0.5	$1.5 \max_{x \in \mathcal{S}} Q_x $	24709.4s		26069	14750
3	0.2	$1.2 \max_{x \in \mathcal{S}} Q_x $	73993.2s		42101	22460

defined by $\{P^c, p^c : c \in \mathcal{C}\}$, where

$$P^c = \begin{matrix} & 0 & 1 & 2 & 3 \\ \begin{matrix} 1 \\ 2 \\ 3 \end{matrix} & \begin{pmatrix} 0 & 0 & 0 & 1 \\ 0 & 0 & 0 & 1 \\ 1 & 0 & 0 & 0 \end{pmatrix} \end{matrix},$$

is identical for all three classes and assumed to be known. In addition, job entries are split evenly, i.e. $p_{0,1}^c = 0.5$ and $p_{0,2}^c = 0.5$ for all $c \in \mathcal{C}$.

In total, we analyse 500 network realizations totalling 1281 tasks during 5083 time units. In order to ease identifiability we assume the existence of a slow, medium and fast server; and assign rather uninformative parameter priors, i.e.

$$\pi(\mu_1^c, \mu_2^c, \mu_3^c) \propto \mathbb{I}(\mu_1^c \leq \mu_2^c \leq \mu_3^c) \times \exp(-10^{-3}(\mu_1^c + \mu_2^c + \mu_3^c))$$

for all $c \in \mathcal{C}$. Note that this network type may not be analysed by means of product-form representations centred around figures of queue-lengths (c.f Wang et al. (2016)). This is because traditional BCMP networks require FCFS stations to share service rates across task classes. On the other hand, an MCMC sampler as presented in (Sutton and Jordan, 2011) can be extended in order to handle general service distributions and target network path transitions; however, the framework is not designed for such aim, it would require an additional Metropolis-Hastings step and it is likely to perform poorly.

In Table 1 we observe summary statistics, computing times and effective sample sizes across three chains with 100000 iterations each, a 10000 burn-in stage, varying starting rates and different scales for dominating rates Ω and

Table 2: Correlation matrix between service rate parameters in a bottleneck network.

	μ_1^2	μ_1^3	μ_2^1	μ_2^2	μ_2^3	μ_3^1	μ_3^2	μ_3^3
μ_1^1	-0.02	-0.01	-0.05	0.01	0.01	-0.05	0.01	0.02
μ_1^2		-0.02	0.00	-0.03	0.01	0.01	-0.04	0.01
μ_1^3			0.01	0.01	-0.09	0.00	-0.01	-0.06
μ_2^1				-0.03	-0.05	-0.13	0.01	0.01
μ_2^2					-0.02	0.00	-0.10	0.01
μ_2^3						0.00	0.00	-0.10
μ_3^1							0.01	0.01
μ_3^2								0.01

probabilities p producing auxiliary-observations. There, we notice a good trade-off between effective samples and the drastic decrease in computing times required when imposing strong serial relations on network paths across subsequent iterations in the sampler. This is the case even when the volume of virtual jumps produced is reduced, and emphasizes the need for such slice sampling techniques in inferential problems with QNs. In addition, Table 2 displays the overall posterior correlation matrix between service rate parameters, and shows very mild relations in rates for each task class. There, we notice the importance of employing posterior samples from the produced chains in order to answer extrapolation-type questions in network systems. Finally, note it is possible to ease imposed restrictions on the network topology and to employ different service disciplines across servers (see next example).

4.3 Feedback network

Finally, we show how the proposed sampling scheme may be used to analyse a real data set. For the purpose, we employ work-logs for medical clinicians. Briefly, the data set includes task requests and completions for individual doctors outside the 9:00-17:00 Monday to Friday *in hours* settings. It belongs to two jointly coordinated university hospitals in the United Kingdom, together servicing a geographical region with over 2.5 million residents. In Figure 7 we show




Figure 7: Sample diagram with a subset of tasking data linked to a clinician during a shift

a diagram with a small of subset of data linked to a clinician during a shift; there we observe three overlapping tasks recorded in the system (from request to completion), and each belonging to a different class. Note that it is not possible to know when the clinician was engaged with each duty; as individual jobs for tasks are not registered when queueing or being routed across teams of administrative staff, nurses and doctors. An extended description of the data set may be found in Perez et al. (2016).

Multiple tasks are grouped across 14 categories and analysed with a *feedback network* as shown in Figure 8. There, we notice the presence of two M/M/1 servers with alternative disciplines and route switching among classes. Task observations for doctors are of roughly two kinds, based on whether they require *engagement* or not. Many tasks are recorded and erased within doctor work-logs in a short time span, due to no need for action; on the other hand, the remainder of tasks exhibit long processing times indicating the need for considerable doctor activity.




Figure 8: Feedback network with two M/M/1 servers and route switching among task classes.

In the proposed example, arrival jobs are buffered within an administrative FCFS priority type queue and depart to a transition center where they either leave the system or get routed for processing with some unknown probability.

Once they are assigned to further processing, they join the doctor's processing centre and switch their routing mechanism; so they will depart the network next time they undergo administrative processing in the first queue. The service station aimed to capture strain on doctor workload is assigned a *processor sharing* (PS) discipline with a single worker, aiming to accommodate doctors attending concurrent duties outside standard working hours. No job service observations are available, so that $q = 0$ and only the arrival and departure times for tasks to the network are observed.

In total, we analyse a reduced subset of 10000 doctor shifts roughly distributed across 4 years of observations.

Table 3: Summary statistics for rate and routing parameters in a feedback network.

		Summary	
		Mean	StDev
Buffer	μ_1	3.397	0.072
Admission	μ_2^1	0.587	0.038
Certification death	$P_{1,0}^1$	0.250	0.026
	μ_2^2	0.723	0.159
Check patient	$P_{1,0}^2$	0.601	0.061
	μ_2^{12}	0.650	0.051
Clerking	$P_{1,0}^{12}$	0.399	0.033
	μ_2^3	0.731	0.018
Clinical review	$P_{1,0}^3$	0.061	0.007
	μ_2^4	0.608	0.015
Address relatives	$P_{1,0}^4$	0.333	0.014
	μ_2^5	0.361	0.079
Drug prescribing	$P_{1,0}^5$	0.470	0.098
	μ_2^6	0.920	0.028
Early warning	$P_{1,0}^6$	0.547	0.012
	μ_2^7	0.488	0.017
Fall of patient	$P_{1,0}^7$	0.346	0.017
	μ_2^8	0.575	0.080
None of above	$P_{1,0}^8$	0.318	0.054
	μ_2^{10}	0.339	0.013
Other services	$P_{1,0}^{10}$	0.393	0.018
	μ_2^9	0.209	0.012
Procedure request	$P_{1,0}^9$	0.435	0.024
	μ_2^{11}	0.635	0.027
Test request	$P_{1,0}^{11}$	0.387	0.019
	μ_2^{13}	0.565	0.013
Urgent response	$P_{1,0}^{13}$	0.263	0.011
	μ_2^{14}	0.209	0.028
	$P_{1,0}^{14}$	0.306	0.059

The network topology is partially known; i.e. $P_{2,1}^c = p_{0,1}^c = 1$ for all task classes, and $P_{1,0}^c = 1$ after tasks have undergone processing and hence switched routing mechanism. However, $P_{1,0}^c = 1 - P_{1,2}^c$ needs to be determined for all existing task classes. Processing rates for tasks are assumed equal in the first service station and different in the PS server; we assign no constraints and we impose loosely uninformative priors such that

$$\pi(\mu_i^c) \propto \exp(-10^{-3}\mu_i^c)$$

for all $i \in \{1, 2\}$ and $c \in \mathcal{C}$. Also, note that within a PS discipline posterior rates given network realizations are given by

$$\mu^c | \mathbf{X} \sim \text{Gamma}(\gamma^c, \sum_{k=1}^K \int_0^{T_k} \phi_c^k(t) dt),$$

for all $c \in \mathcal{C}$, where γ^c denotes the number of class c jobs served at the station in all realizations in \mathbf{X} ; and

$$\phi_c^k(t) = \frac{\sum_{j=0}^J \mathbb{I}(\text{Job } j \text{ is class } c) \cdot \mathbb{I}(a_j < t < d_j)}{\sum_{j=0}^J \mathbb{I}(a_j < t < d_j)},$$

where summations are across all jobs processed in the PS station in realization k , and a_j, d_j denote the arrival and departure times of the job to the server. In Table 3 we observe summary statistics for parameters across two chains with 100000 iterations each, a 50000 burn-in stage and varying starting rates. In one chain, we use $\Omega = 2 \max_{x \in \mathcal{S}} |Q_x|$ and $p = 0.75$; in the second we have $\Omega = 1.5 \max_{x \in \mathcal{S}} |Q_x|$ and $p = 0.5$.

Table 4: Point estimates and standard errors for average processing times excluding waiting times, across different tasks. These relate to times from entry to departure in network.

	Completion			Completion	
	Mean	StErr		Mean	StErr
Admission	1.796	0.087	Early w.	1.828	0.046
Cert. death	0.985	0.141	Fall	1.701	0.176
Check	1.400	0.073	None	2.263	0.076
Clerking	1.856	0.032	Other	3.173	0.173
Cl. review	1.588	0.027	Procedure	1.442	0.045
Address rel.	1.971	0.374	Test	1.818	0.032
Drug pres.	0.920	0.017	Urgent	3.877	0.491

In addition, Table 4 shows point estimates and standard errors for average completion times in all task types, these correspond to the full processing times from entry to departure in the network (excluding queueing times) and are reported in hour units. Hence, we notice it is possible to assess workload both globally and across single components in the system, thus allowing to answer extrapolation kinds of questions on workload; i.e. in relation to means, variances and extreme values for system strain under likely alterations.

5 Discussion

This paper has presented a flexible approach for carrying exact Bayesian inference within known or hypothesized queueing networks. Its focus is on multi-class, open and Markovian cases and the approach is centred around the underlying continuous-time Markov chains induced by these complex stochastic systems. The proposed method relies on a slice sampling technique with mappings to the space of task transitions across servers in the network. It sits well over uniformization-oriented MCMC approaches introduced in Rao and Teh (2013) and can deal with missing data, imposed prior knowledge and strong serial dependencies posing a complex inferential task (cf. Sutton and Jordan (2011)).

The need for such inferential frameworks with missing data is justified by the ability of general-form networks to allow evaluating response times in complex systems. Overall, recovering measures such as processing times is a technically difficult task when designing increasingly complex IT systems (Liu et al., 2006), or in service delivery networks (such as those in hospitals) due to ethical issues with such an intrusive process (Perez et al., 2016). Yet, QNs provide the tools to assess system alterations, diagnose poor performance or evaluate robustness to spikes in workload.

The advantage of the presented inferential method is that it permits retrospectively assessing the likely status of systems at any point in time; rather than only providing summary information on strain over individual bits. However, limitations relate to tractability restrictions with high-magnitude networks. In such cases, controlling the dimensionality of unobservable state spaces requires imposing strong serial dependencies within simulated latent network paths across steps in the sampler. This however may restrict the produced chain from exploring the posterior range of network paths efficiently. Approximate inferential frameworks relying on reduced product-form simplifications of state beliefs may improve the scalability of the method. Moreover, it is possible to explore the use of particle filtering approaches along with auxiliary variables for this purpose, since clamping explorable spaces within filtering procedures would likely ease the usual challenges regarding particle degeneracy; that is, ending with a very few particles having non-zero weights.

Also, the use of the uniformization technique will limit applications of the present framework to the study of purely Markovian processes. While it is possible to employ Markov-modulated regimes that adapt service and arrival rates to network states, this will greatly expand state spaces under consideration. Also, uniformization may deem the sampler computationally inefficient should service rates vary greatly across queues or job classes, as certain transition types will greatly dominate the underlying discrete time Markov chain.

Finally, the paper assumes that the volume of job service observations retrieved across the network is given by $\%(100 \cdot q)$ of the total processing during a fixed time interval.

For simplicity, q is assumed fixed and known to the user. Many network structures (such as bottleneck networks) will allow for uncertainty regarding this parameter to be quantified by means of the presented sampler, as each iteration will provide a total number of network transitions complementing the observation number as a sufficient statistic. However, it is necessary to impose the knowledge of q in order to ensure model identifiability whenever networks contain either global or self-loops.

Supplementary material

Synthetic data along with a Java implementation of the algorithm can be found in <https://bitbucket.org/ikertxo1986/auxvarsamplerjava> or <https://github.com/IkerPerez/auxVarSampler>. This allows to reproduce the results within the examples above.

Acknowledgements

We would like to thank the anonymous reviewers for their valuable remarks and suggestions that have improved the quality of this paper.

References

- Armero, C. and Bayarri, M. J. (1994). Prior assessments for prediction in queues. *The Statistician*, 43(1):139–153.
- 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(2):248–260.
- Boyen, X. and Koller, D. (1998). Tractable inference for complex stochastic processes. In *Proceedings of the Fourteenth Conference on Uncertainty in Artificial Intelligence*, UAI’98, pages 33–42.
- Buzacott, J. A. and Shanthikumar, J. G. (1993). *Stochastic models of manufacturing systems*, volume 4. Prentice Hall, New Jersey.
- Fan, Y. and Shelton, C. R. (2008). Sampling for approximate inference in continuous time bayesian networks. In *Tenth International Symposium on Artificial Intelligence and Mathematics*, ISAIM’08.
- 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(3):1204–1231.
- Jensen, A. (1953). Markoff chains as an aid in the study of Markoff processes. *Scandinavian Actuarial Journal*, 36:87–91.
- Kalli, M., Griffin, J. E., and Walker, S. G. (2011). Slice sampling mixture models. *Statistics and Computing*, 21(1):93–105.
- Kleinrock, L. (1976). *Queueing Systems Vol II: Computer Applications*. Wiley, New York.
- Koole, G. and Mandelbaum, A. (2002). Queueing models of call centers: An introduction. *Annals of Operations Research*, 113(1):41–59.
- Lippman, S. A. (1975). Applying a new device in the optimization of exponential queuing systems. *Operations Research*, 23(4):687–710.
- 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(1):36 – 60.
- Neal, R. M. (2003). Slice sampling. *The Annals of Statistics*, 31(3):705–767.
- Ng, B., Peshkin, L., and Pfeffer, A. (2002). Factored particles for scalable monitoring. In *Proceedings of the Eighteenth Conference on Uncertainty in Artificial Intelligence*, UAI’02, pages 370–377.
- Nodelman, U., Shelton, C. R., and Koller, D. (2002). Continuous time bayesian networks. In *Proceedings of the Eighteenth Conference on Uncertainty in Artificial Intelligence*, UAI’02, pages 378–387.
- 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(3):996 – 1007.
- Perez, I., Brown, M., Pinchin, J., Martindale, S., Sharples, S., Shaw, D., and Blakey, J. (2016). Out of hours workload management: Bayesian inference for decision support in secondary care. *Artificial Intelligence in Medicine*, 73:34 – 44.
- 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.
- Spinner, S., Casale, G., Brosig, F., and Kounev, S. (2015). Evaluating approaches to resource demand estimation. *Performance Evaluation*, 92:51 – 71.
- Sutton, C. and Jordan, M. I. (2011). Bayesian inference for queueing networks and modeling of internet services. *The Annals of Applied Statistics*, 5(1):254–282.
- Van Gael, J., Saatchi, Y., Teh, Y. W., and Ghahramani, Z. (2008). Beam sampling for the infinite hidden markov model. In *Proceedings of the 25th International Conference on Machine Learning*, ICML’08, pages 1088–1095. ACM.

Walker, S. G. (2007). Sampling the dirichlet mixture model with slices. *Communications in Statistics - Simulation and Computation*, 36(1):45–54.

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(1):2:1–2:26.