

EXACT MONTE CARLO SIMULATION FOR FORK-JOIN NETWORKS

HONGSHENG DAI,* *University of Brighton*

Abstract

In a fork-join network each incoming job is split into K tasks and the K tasks are simultaneously assigned to K parallel service stations for processing. For the distributions of response times and queue lengths of fork-join networks, no explicit formulae are available. Existing methods provide only analytic approximations for the response time and the queue length distributions. The accuracy of such approximations may be difficult to justify for some complicated fork-join networks. In this paper we propose a perfect simulation method based on *coupling from the past* to generate exact realisations from the equilibrium of fork-join networks. Using the simulated realisations, Monte Carlo estimates for the distributions of response times and queue lengths of fork-join networks are obtained. Comparisons of Monte Carlo estimates and theoretical approximations are also provided. The efficiency of the sampling algorithm is shown theoretically and via simulation.

Keywords: Coupling from the past; fork-join network; perfect sampling; read-once coupling from the past

2010 Mathematics Subject Classification: Primary 65C05; 65C50
Secondary 60K20; 60K25

1. Introduction

In a fork-join network jobs arrive according to a Poisson process with rate λ and each incoming job is split into K tasks at the fork station. Then the K tasks are simultaneously assigned to K parallel service stations for processing. The i th service station has s_i servers. Each server has exponential service times with rate μ_i . When all the K tasks of a job are completed, they will be joined immediately at the join station and leave the system.

Such fork-join networks are commonly used to model computer and manufacturing systems. For example, each logical computer I/O request is split into several physical I/O requests across the disk device and after all physical I/O requests are completed they are joined. Another example is that of manufacturing the assembly of a product, where the system requires several parts to be processed simultaneously at parallel work stations.

A fork-join queueing model is shown in Figure 1, where each service station has finite capacity N (the number of task-waiting places), i.e. the i th service station can hold at most $N + s_i$ tasks (including tasks in service and tasks in the queue). Other fork-join networks may have infinite capacity ($N = \infty$) or be such that the servers in each service station have different service rates. For various types of fork-join network, our main interests are the probability distributions of the response time (the time to complete a job) and the queue length (the number of jobs in the network) in equilibrium. The difficulty is that the stationary distribution of the

Received 30 April 2010; revision received 9 December 2010.

* Postal address: Department of Mathematics, School of Computing, Engineering and Mathematics, University of Brighton, Lewes Road, Brighton BN2 4GJ, UK. Email address: h.dai@brighton.ac.uk




FIGURE 1: A fork-join network.

fork-join network, even if the network is as simple as that given in Figure 1, is intractable (see [9]).

Various methods for approximating or bounding the mean response times and approximating the queue length distributions for fork-join networks have been developed. Flatto and Hahn [6] and Flatto [5] derived the generating functions of the queue length distribution for a fork-join network having two service stations, each with a single server. Baccelli *et al.* [2], Balsoma *et al.* [3], and Nelson and Tantawi [10] gave bounds for the mean response time of an M/M/1 fork-join network. Ko and Serfozo [9] established a closed-form formula for approximating the distributions of the response time and queue length for certain fork-join networks in equilibrium. Raghavan and Viswanadhan [12] and Xia *et al.* [16] presented bounds for the mean response time of several types of fork-join network. Squillante *et al.* [14] considered a more general fork-join queueing system in which dynamic policies are considered for scheduling multiple tasks to maintain effective server utilization. However, most existing methods focus on analytical approximations for the equilibrium of fork-join networks. Simulation results in [9] demonstrated that the accuracy of such approximations may decrease as $s_i \mu_i$ decreases to λ .

In this paper we consider the use of Monte Carlo simulations to estimate the distributions of the response time and queue length when N is finite. This work is important since existing approximation methods are valid for only $N = \infty$. The accuracy of the Monte Carlo estimates depend on the number of simulated realisations; the more simulated realisations, the more accurate the Monte Carlo estimate. Therefore, such Monte Carlo estimates are important for fork-join networks with $\lambda \approx s_i \mu_i$ where existing analytic approximations may not be good enough. Another advantage of the Monte Carlo method is that we can obtain queue length distribution estimates for fork-join networks with multiple servers in each service station. However, analytic approximations of the queue length distribution for such fork-join networks are not available using existing methods, such as the method in [9]. In addition, we can also use the Monte Carlo distribution estimate to justify the accuracy of existing analytical approximations for fork-join networks. Finally, most existing methods focus on analytic approximations of mean response times; however, the response time distribution characteristics, such as the median and quartiles, are sometimes more important, as, for example, in our simulation scenarios where the response time distribution is highly skewed.

Traditional Monte Carlo simulation methods run the fork-join network as a continuous-time Markov chain starting from a given state. Then realisations are collected after the burn-in stage. Such methods only generate random realisations approximately from the equilibrium of the network. Thus, the estimates highly depend on the quality of the random realisations. In this paper we propose a perfect simulation method based on *coupling from the past* (CFTP),

introduced in [11], to generate exact realisations from the equilibrium of fork-join networks given in Figure 1. The CFTP algorithm developed in this paper is very efficient (a polynomial-time algorithm), i.e. the mean coalescence time of the algorithm is bounded by a polynomial function of KN . Based on the simulated realisations we provide Monte Carlo estimates for the distributions of the response time and queue length. Comparisons of the Monte Carlo estimates based on perfect sampling and approximation results in [9] demonstrate that the proposed perfect sampling method works very well.

This paper is organized as follows. In Section 2 we introduce the model notation and an algorithm to simulate a fork-join network from any given starting state. Then we develop a CFTP algorithm with bounding chains in Section 3 to simulate exactly from the equilibrium of the fork-join network. Simulation studies and estimates of the distributions of the queue length and response times are provided in Section 4. Complexity of the algorithm is discussed in Section 5. Section 6 contains a discussion.

2. Preliminaries and continuous-time Markov chain simulation

2.1. Notation

Consider the fork-join network shown in Figure 1. Jobs arrive at the system according to a Poisson process with rate λ and each incoming job is split into K tasks at the fork station. Then the K tasks are simultaneously assigned to K parallel service stations for processing. The i th station has s_i servers, each with exponential service times with mean $1/\mu_i$. When all the K tasks of a job are completed, they are joined immediately at the join station (the service time at the join station is 0) and leave the system. Service station i can hold $N_i = N + s_i$ tasks at most. Define $\mathbf{N} = (N_1, \dots, N_K)$.

If the tasks of a job are not all completed, the completed tasks will wait in the corresponding buffer of the join station. We also assume that the join station has enough places to hold all completed waiting tasks.

Let $Q_i(t)$ be the number of tasks at service station i (including tasks in service and tasks waiting) at time t . We know that $Q_i(t)$ can take a value in $\{0, 1, \dots, N_i + s_i\}$. The fork-join network can be represented by $\mathbf{Q}(t) = (Q_1(t), \dots, Q_K(t))$, which is a K -dimensional continuous-time Markov chain.

From the results in [9] we know that when $N = \infty$, the transition rate for the K -dimensional continuous-time Markov chain $\mathbf{Q}(t)$ from a state $\mathbf{n} = (n_1, \dots, n_K)$ to another state \mathbf{n}' is given by

$$q(\mathbf{n}, \mathbf{n}') = \begin{cases} \lambda & \text{if } \mathbf{n}' = \mathbf{n} + \mathbf{1}, \\ \mu_i \min\{n_i, s_i\} & \text{if } \mathbf{n}' = \mathbf{n} - \mathbf{e}_i, \\ 0 & \text{otherwise.} \end{cases}$$

Here \mathbf{e}_i is a K -dimensional vector with 1 in the i th component and 0s elsewhere, and $\mathbf{1}$ is a K -dimensional vector of all 1s. The Markov chain with $N = \infty$ is ergodic when $\lambda < \mu_i s_i$ for all values of i , $1 \leq i \leq K$.

Here we consider N to be a finite number, which means that each service station can only hold up to a finite number of tasks. This is reasonable in practice. If service station i holds the maximum number of tasks, $N + s_i$, the newly arrived job will not be split into K tasks since the buffer of service station i is full. Therefore, when N is finite, the transition rate for $\mathbf{Q}(t)$ is

given by

$$q(\mathbf{n}, \mathbf{n}') = \begin{cases} \lambda & \text{if } \mathbf{n}' = \mathbf{n} + \mathbf{1} \text{ and } \mathbf{n} + \mathbf{1} \leq N, \\ \mu_i \min\{n_i, s_i\} & \text{if } \mathbf{n}' = \mathbf{n} - \mathbf{e}_i, \\ 0 & \text{otherwise.} \end{cases} \tag{1}$$

2.2. Continuous-time Markov chain simulation

We now introduce an algorithm to simulate the Markov chain $Q(t)$ with a given starting state. This algorithm will be used in the next section to develop a CFTP algorithm.

Using the continuous-time Markov chain theory, we know that the amount of time that $Q(t)$ spends in the current state \mathbf{n} (holding time of a state) is exponentially distributed with rate

$$r_n := \begin{cases} \sum_{\mathbf{n}' \neq \mathbf{n}} q(\mathbf{n}, \mathbf{n}') = \lambda + \sum_i \mu_i \min\{n_i, s_i\} & \text{if } \mathbf{n} + \mathbf{1} \leq N, \\ \sum_{\mathbf{n}' \neq \mathbf{n}} q(\mathbf{n}, \mathbf{n}') = \sum_i \mu_i \min\{n_i, s_i\} & \text{if } \mathbf{n} + \mathbf{1} \not\leq N. \end{cases}$$

Then at the end of the holding time, $Q(t)$ jumps to another state \mathbf{n}' with transition probability

$$p_{\mathbf{n}, \mathbf{n}'} = \begin{cases} \frac{\lambda}{r_n} & \text{if } \mathbf{n}' = \mathbf{n} + \mathbf{1} \text{ and } \mathbf{n} + \mathbf{1} \leq N, \\ \frac{\mu_i \min\{n_i, s_i\}}{r_n} & \text{if } \mathbf{n}' = \mathbf{n} - \mathbf{e}_i, \\ 0 & \text{otherwise.} \end{cases} \tag{2}$$

Note that the holding times of each state and the transitions of $Q(t)$ are independent.

We may simulate the continuous-time Markov chain using the following method. Suppose that $Q(t) = \mathbf{n}$. Simulate the next transition time $t' = T + t$, where T is drawn from an exponential distribution with rate r_n . This can be done by simulating $\xi_0 \sim \exp(\lambda)$ and $\xi_{ij} \sim \exp(\mu_i)$, $j = 1, \dots, s_i$, $i = 1, \dots, K$, and setting

$$T = \min\{\xi_0, \xi_{ij}, j = 1, \dots, \min\{n_i, s_i\}, i = 1, \dots, K\}$$

if $\mathbf{n} + \mathbf{1} \leq N$ or setting $T = \min\{\xi_{ij}, j = 1, \dots, \min\{n_i, s_i\}, i = 1, \dots, K\}$ if $\mathbf{n} + \mathbf{1} \not\leq N$. This is because such T is from $\exp(r_n)$. Note that if $n_i = 0$ then define $\min_{j \leq \min\{n_i, s_i\}} \xi_{ij} = \infty$.

At time t' the process Q will jump to \mathbf{n}' with probability $p_{\mathbf{n}, \mathbf{n}'}$. The transitions of the chain can be simulated as follows. Provided that $\mathbf{n} + \mathbf{1} \leq N$, if $T = \xi_0$ then let $\mathbf{n}' = \mathbf{n} + \mathbf{1}$; if $T = \min_{j \leq \min\{n_i, s_i\}} \xi_{ij}$ for some i then let $\mathbf{n}' = \mathbf{n} - \mathbf{e}_i$. This is because $P(T = \min_{j \leq \min\{n_i, s_i\}} \xi_{ij} \text{ for some } i) = \mu_i \min\{n_i, s_i\} / r_n$ and $P(T = \xi_0) = \lambda / r_n$, which correspond to the transition probabilities in (2). Similarly, given $\mathbf{n} + \mathbf{1} \not\leq N$, if $T = \min_{j \leq \min\{n_i, s_i\}} \xi_{ij}$ for some i then let $\mathbf{n}' = \mathbf{n} - \mathbf{e}_i$. On the other hand, with such a simulation method, the transitions are independent of the holding times. So the Markov chain $Q(t)$ is correctly simulated.

The above simulation method can be proved using the properties of exponential variables. Details can be found in [13, Chapter 5]. Although this simulation idea is easy to implement, in order to develop a CFTP method later, we use the following more complicated algorithm instead to simulate $\{Q(t), t \leq \mathcal{T}\}$. The algorithm outputs a sequence of times $\tau = \{\tau_1, \tau_2, \dots\}$ which is a superset of the transition times $t = \{t_1, t_2, \dots\}$ of $Q(t)$. A thinning method is used to identify the transition times t from τ .

Algorithm 1. (Markov chain simulation algorithm.) Outputs: τ_l , $\mathbf{Q}(\tau_l)$, and t_k .

01. Set $t_0 = \tau_0 = 0, k, l = 0$, and $\mathbf{Q}(t_0) =$ starting state.
02. **While** ($t_k < \mathcal{T}$)
03. Simulate $\xi_0^{(l)} \sim \exp(\lambda)$ and $\xi_{ij}^{(l)} \sim \exp(\mu_i)$ for $j = 1, \dots, s_i$ and $i = 1, \dots, K$.

Note that we use $\mathcal{U}^{(l)}$ to denote all simulated $\xi^{(l)}$ s.

04. $\text{MC}(\tau_l, \mathbf{Q}(\tau_l), \mathcal{U}^{(l)}; \tau_{l+1}, \mathbf{Q}(\tau_{l+1}), I)$ (given in Algorithm 2 below).
05. **If** $I = 1$ **then** $t_{k+1} = \tau_{l+1}$ (the next transition time) and $k = k + 1$.
06. $l = l + 1$.

Algorithm 2. (Markov chain updating algorithm.) This subroutine is given by $\text{MC}(\tau_l, \mathbf{Q}(\tau_l), \mathcal{U}^{(l)}; \tau_{l+1}, \mathbf{Q}(\tau_{l+1}), I)$.

Inputs: $\tau_l, \mathbf{Q}(\tau_l)$, and $\mathcal{U}^{(l)}$. Outputs: $\tau_{l+1}, \mathbf{Q}(\tau_{l+1})$, and I .

For simplicity, denote $\mathbf{Q}(\tau_l)$ as $\mathbf{n}^{(l)} = (n_1^{(l)}, \dots, n_K^{(l)})$. Let $T_l^{\min} = \min\{\xi_0^{(l)}, \xi_{ij}^{(l)}, j = 1, \dots, s_i, i = 1, \dots, K\}$ and $T_l = \min\{\xi_0^{(l)}, \xi_{ij}^{(l)}, j = 1, \dots, \min\{n_i^{(l)}, s_i\}, i = 1, \dots, K\}$.

01. $\tau_{l+1} = \tau_l + T_l^{\min}$
02. **If** $T_l = T_l^{\min}$ **then**
03. **if** $T_l = \xi_0^{(l)}$ **then**
04. $\mathbf{Q}(\tau_{l+1}) = \mathbf{Q}(\tau_l) + \mathbf{1}$ and $I = 1$
05. **if** $\mathbf{Q}(\tau_{l+1}) \not\leq N$ **then** $\mathbf{Q}(\tau_{l+1}) = \mathbf{Q}(\tau_l)$ and $I = 0$.
06. **else if** $T_l = \xi_{ij}^{(l)}$ for some i **then**
07. $\mathbf{Q}(\tau_{l+1}) = \mathbf{Q}(\tau_l) - \mathbf{e}_i$ and $I = 1$
08. **end if**
09. **else**
10. $I = 0$ and $\mathbf{Q}(\tau_{l+1}) = \mathbf{Q}(\tau_l)$
11. **end if**

Lemma 1. Algorithm 1 outputs $\boldsymbol{\tau} = \{\tau_0, \tau_1, \dots, \tau_l, \dots\}$ and $\{\mathbf{Q}(\tau_l), l = 0, 1, \dots\}$. The algorithm also outputs a subsequence of $\boldsymbol{\tau}$, as $\{t_0, t_1, \dots, t_k, \dots\}$.

The simulated continuous-time Markov process is given by $\{t_k, \mathbf{Q}(t_k)\}, k = 0, 1, \dots$. The time t_k is the k th transition time of $\mathbf{Q}(t)$ and the process $\mathbf{Q}(t)$ is constant in $[t_k, t_{k+1})$. The simulated process has transition rates given in (1).

Lemma 1 is true since Algorithm 1 simulates potential holding times T_l^{\min} of $\mathbf{Q}(t)$ from an exponential distribution with rate parameter $\lambda + \sum_i \mu_i s_i$, which is larger than the rates of exponential holding times of all states. Then a thinning method is used to check whether time τ_l is a transition time point of $\mathbf{Q}(t)$. See Appendix A for a detailed proof.

Note that the updating rules in Algorithm 2 can be denoted as a deterministic function ϕ as

$$\begin{aligned}
 \mathbf{Q}(\tau_l + a) &= \phi(\mathbf{Q}(\tau_l), \mathcal{U}^{(l)}, a) \\
 &= \begin{cases} \mathbf{Q}(\tau_l) & \text{if } 0 < a < T_l^{\min}, \\ \mathbf{Q}(\tau_l) + \mathbf{1} & \text{if } a = T_l^{\min}, \mathbf{Q}(\tau_l) + \mathbf{1} \leq N, \text{ and } T_l^{\min} = T_l = \xi_0^{(l)} \\ & \text{(line 04 of Algorithm 2),} \\ \mathbf{Q}(\tau_l) & \text{if } a = T_l^{\min}, \mathbf{Q}(\tau_l) + \mathbf{1} \not\leq N, \text{ and } T_l^{\min} = T_l = \xi_0^{(l)} \\ & \text{(line 05 of Algorithm 2),} \\ \mathbf{Q}(\tau_l) - \mathbf{e}_i & \text{if } a = T_l^{\min} \text{ and } T_l^{\min} = T_l = \xi_{ij}^{(l)} \text{ for some } i \\ & \text{(line 07 of Algorithm 2),} \\ \mathbf{Q}(\tau_l) & \text{if } a = T_l^{\min} \text{ and } T_l^{\min} < T_l \text{ (line 10 of Algorithm 2),} \end{cases} \quad (3)
 \end{aligned}$$

where $\mathcal{U}^{(l)}$ is defined in Algorithm 1.

3. Simulating from equilibrium using CFTP with bounding chains

CFTP was introduced in the landmark paper by Propp and Wilson [11], who showed how to provide perfect samples from the limiting distribution of a Markov chain. The idea is to run Markov chains starting from all states simultaneously from the past, using the same random numbers in the updating procedure, until all chains coalesce into a single chain, then keep running the coalesced chain and collect a sample at time 0.

By running all Markov chains simultaneously, CFTP may result in very heavy computational costs. The fork-join network can be described by the K -dimensional continuous-time Markov chain $\mathbf{Q}(t)$ with state space \mathcal{Q} . The number of states in \mathcal{Q} is larger than K^N . It is impossible to run so many chains simultaneously. The computational cost of the algorithm can be reduced if, for the state space \mathcal{Q} , there is a partial order preserved by the process $\mathbf{Q}(t)$. This is called monotone CFTP for which we only need to run the upper chain starting from the maximum state and the lower chain starting from the minimum state (see [11] for more details). If the upper and lower chains, which bound all the other chains, coalesce then all chains coalesce. Thus, the efficiency of the CFTP algorithm can be improved significantly.

Although monotone CFTP is easy to perform, finding a partial order preserved by the Markov chains is a nontrivial task in many cases. An alternative improvement is CFTP with bounding chains or dominating processes; see [4], [7], and [8]. The advantages of bounding-chain CFTP are that the partial order preserved by the Markov chains is not required and that only several bounding chains, which bound all Markov chains, are required to run simultaneously.

In Section 3.1 we first define a partial order for \mathcal{Q} . Then we show that the partial order is preserved by $\mathbf{Q}(t)$ in some cases, but not always. Then based on this result we provide a CFTP algorithm using bounding chains in Sections 3.2 and 3.3.

3.1. A partial order

For $\mathbf{n}, \mathbf{n}' \in \mathcal{Q}$, define $\mathbf{n} < \mathbf{n}'$ if $n_j \leq n'_j, j = 1, \dots, K$, and $n_j < n'_j$ for at least one value of j . Define $\mathbf{n} = \mathbf{n}'$ if $n_j = n'_j$ for $j = 1, \dots, K$.

The partial order is preserved by $\mathbf{Q}(t)$ meaning that if $\mathbf{Q}(\tau_l) \leq \mathbf{Q}'(\tau_l)$ then $\phi(\mathbf{Q}(\tau_l), \mathcal{U}^{(l)}, a) \leq \phi(\mathbf{Q}'(\tau_l), \mathcal{U}^{(l)}, a)$ for all a such that $0 < a \leq T_l^{\min}$. The following lemma gives the cases where the partial order is preserved and the cases where the partial order is not preserved.

Lemma 2. Suppose that the current time is τ_l and that $\mathbf{Q}(\tau_l) = \mathbf{n} < \mathbf{Q}'(\tau_l) = \mathbf{n}'$. Define

$$T_l = \min\{\xi_0^{(l)}, \xi_{ij}^{(l)}, j = 1, \dots, \min\{n_i, s_i\}, i = 1, \dots, K\}$$

and

$$T'_l = \min\{\xi_0^{(l)}, \xi_{ij}^{(l)}, j = 1, \dots, \min\{n'_i, s_i\}, i = 1, \dots, K\}.$$

Then we have the following results.

(a) For all a such that $0 < a < T_l^{\min}$, we have

$$\mathbf{Q}(\tau_l + a) = \phi(\mathbf{Q}(\tau_l), \mathcal{U}^{(l)}, a) \leq \mathbf{Q}'(\tau_l + a) = \phi(\mathbf{Q}'(\tau_l), \mathcal{U}^{(l)}, a).$$

Define the events $\mathbb{A}^{(l)} = \{T_l^{\min} \neq \xi_0^{(l)}\}$, $\mathbb{B}_1^{(l)} = \{(\mathbf{n}, \mathbf{n}') : \mathbf{n}' + \mathbf{1} \leq N\}$, $\mathbb{B}_2^{(l)} = \{(\mathbf{n}, \mathbf{n}') : \mathbf{n} + \mathbf{1} \not\leq N\}$, and $\mathbb{B}_3^{(l)} = (B_1^{(l)} \cup B_2^{(l)})^c$. Obviously, $(\mathbf{Q}(\tau_l), \mathbf{Q}'(\tau_l))$ must belong to $\mathbb{B}_1^{(l)} \cup \mathbb{B}_2^{(l)} \cup \mathbb{B}_3^{(l)}$.

For $a = T_l^{\min}$, the relation between $\mathbf{Q}(\tau_l + a)$ and $\mathbf{Q}'(\tau_l + a)$ must be such that one of the following cases holds.

(b) If $T_l^{\min} \neq \xi_0^{(l)}$, i.e. $\mathbb{A}^{(l)}$ is true, we have

$$\phi(\mathbf{Q}(\tau_l), \mathcal{U}^{(l)}, a) \leq \phi(\mathbf{Q}'(\tau_l), \mathcal{U}^{(l)}, a).$$

(c) If $T_l^{\min} = \xi_0^{(l)}$, i.e. $\mathbb{A}^{(l)}$ is not true, and $(\mathbf{Q}(\tau_l), \mathbf{Q}'(\tau_l)) \in \mathbb{B}_1^{(l)}$, then

$$\phi(\mathbf{Q}(\tau_l), \mathcal{U}^{(l)}, a) \leq \phi(\mathbf{Q}'(\tau_l), \mathcal{U}^{(l)}, a).$$

(d) If $\mathbb{A}^{(l)}$ is not true and $(\mathbf{Q}(\tau_l), \mathbf{Q}'(\tau_l)) \in \mathbb{B}_2^{(l)}$, then

$$\phi(\mathbf{Q}(\tau_l), \mathcal{U}^{(l)}, a) \leq \phi(\mathbf{Q}'(\tau_l), \mathcal{U}^{(l)}, a).$$

(e) If $\mathbb{A}^{(l)}$ is not true and $(\mathbf{Q}(\tau_l), \mathbf{Q}'(\tau_l)) \in \mathbb{B}_3^{(l)}$, then

$$\phi(\mathbf{Q}(\tau_l), \mathcal{U}^{(l)}, a) \leq \phi(\mathbf{Q}'(\tau_l), \mathcal{U}^{(l)}, a)$$

is not guaranteed.

Proof. See Appendix B.

Note that result Lemma 2(a) implies that the partial order is always preserved by the process in (τ_l, τ_{l+1}) . If $a = T_l^{\min}$ then $\tau_l + a = \tau_{l+1}$. Thus, Lemma 2(b) means that the partial order is preserved if $\mathbb{A}^{(l)}$ is true. Similarly, at time τ_{l+1} the partial order is also preserved under the conditions of Lemma 2(c) and (d). But the partial order may not be preserved at τ_{l+1} under the conditions of Lemma 2(e). Therefore, we cannot simply use a monotone CFTP algorithm, but it is possible to use a bounding-chain algorithm to deal with this problem.

3.2. Bounding chains

We run Markov chains starting from different states from time 0. Suppose that $\underline{\mathbf{Q}}(0) \leq \overline{\mathbf{Q}}(0)$ bounds all other Markov chains. We construct two bounding chains $\underline{\mathbf{Q}}(t)$ and $\overline{\mathbf{Q}}(t)$ as follows.

Suppose that the current time is τ_l . After simulating random numbers $\mathcal{U}^{(l)}$, we update $\underline{\mathbf{Q}}(t)$ and $\overline{\mathbf{Q}}(t)$ in (τ_l, τ_{l+1}) according to Algorithm 2. We know that $\underline{\mathbf{Q}}(t)$ and $\overline{\mathbf{Q}}(t)$ will bound all Markov chains in (τ_l, τ_{l+1}) since, according to Lemma 2(a), the partial order is preserved.

At time τ_{l+1} , if one of the conditions in Lemma 2(b), (c), or (d) is satisfied, we also update $\underline{\mathbf{Q}}(t)$ and $\overline{\mathbf{Q}}(t)$ at τ_{l+1} according to Algorithm 2. We know that $\underline{\mathbf{Q}}(\tau_{l+1})$ and $\overline{\mathbf{Q}}(\tau_{l+1})$ will

bound all Markov chains $\mathbf{Q}(\tau_{l+1})$ since the partial order is preserved according to Lemma 2(b), (c), and (d).

If the condition in Lemma 2(e) is true, we let

$$\begin{aligned} \underline{\mathbf{Q}}(\tau_{l+1}) &= \underline{\mathbf{Q}}(\tau_l), \\ \overline{\mathbf{Q}}(\tau_{l+1}) &= (\min\{\overline{Q}_1(\tau_l) + 1, N_1\}, \dots, \min\{\overline{Q}_K(\tau_l) + 1, N_K\}). \end{aligned} \tag{4}$$

Lemma 3. *The above bounding chains will bound all Markov chains, i.e. $\underline{\mathbf{Q}}(t) \leq \mathbf{Q}(t) \leq \overline{\mathbf{Q}}(t)$ for all t and any Markov chains $\mathbf{Q}(t)$.*

Proof. We only need to prove that if the condition in Lemma 2(e) is true then $\underline{\mathbf{Q}}(\tau_l) \leq \mathbf{Q}(\tau_l) \leq \overline{\mathbf{Q}}(\tau_l)$ implies that $\underline{\mathbf{Q}}(\tau_{l+1}) \leq \mathbf{Q}(\tau_{l+1}) \leq \overline{\mathbf{Q}}(\tau_{l+1})$.

If the condition in Lemma 2(e) is true, we know that $\mathbf{Q}(\tau_{l+1}) = \mathbf{Q}(\tau_l) + \mathbf{1}$ or $\mathbf{Q}(\tau_{l+1}) = \mathbf{Q}(\tau_l)$. Obviously, given $\underline{\mathbf{Q}}(\tau_l) \leq \mathbf{Q}(\tau_l) \leq \overline{\mathbf{Q}}(\tau_l)$, we have

$$\underline{\mathbf{Q}}(\tau_l) \leq \mathbf{Q}(\tau_{l+1}) \quad \text{and} \quad \mathbf{Q}(\tau_{l+1}) \leq (\min\{\overline{Q}_1(\tau_l) + 1, N_1\}, \dots, \min\{\overline{Q}_K(\tau_l) + 1, N_K\}).$$

So (4) guarantees that $\underline{\mathbf{Q}}(\tau_{l+1}) \leq \mathbf{Q}(\tau_{l+1}) \leq \overline{\mathbf{Q}}(\tau_{l+1})$. This completes the proof.

When $\overline{\mathbf{Q}}(t)$ and $\underline{\mathbf{Q}}(t)$ coalesce, all Markov chains coalesce. Based on this, a CFTP algorithm is given in the next subsection.

3.3. CFTP with bounding chains

Define an upper-bound chain starting from N and a lower-bound chain starting from $\mathbf{0} = (0, \dots, 0)$. When the upper-bound chain and lower-bound chain coalesce, all Markov chains coalesce. When using CFTP, we need to re-use old random numbers. A simple variation on CFTP is the so called read-once CFTP proposed in [15], which outputs random realisations from equilibrium using just a read-once source of random numbers. Here we use the read-once CFTP algorithm, as Wilson [15] suggested that read-once CFTP has the advantage if many independent realisations are desired. According to Wilson [15], the following read-once CFTP algorithm returns realisations from the equilibrium of $\mathbf{Q}(t)$.

Algorithm 3. *(Read-once CFTP (NumberofRealisations).)*

01. State := (0, 0, . . . , 0) (it can be an arbitrary state)
02. Repeat
03. ApplyCompositeMap(State,CoalescenceFlag) (given in Algorithm 4 below)
04. Until CoalescenceFlag
05. For $i = 1$ to NumberofRealisations
06. Repeat
07. OldState = State
08. ApplyCompositeMap(State,CoalescenceFlag)
09. Until CoalescenceFlag
10. Output OldState

Algorithm 4. (*ApplyCompositeMap(State,CoalescenceFlag)*.) Inputs: State and CoalescenceFlag. Outputs: State and CoalescenceFlag.

01. $\bar{n}_1 = N, \underline{n}_1 = \mathbf{0}, \bar{n}_2 = N, \underline{n}_2 = \mathbf{0}, l = 0, \tau_l = 0,$ and $\tau'_l = 0.$
02. **while** $\bar{n}_2 \neq \underline{n}_2$
03. Simulate $\mathcal{U}^{(l)} = \{\xi_0^{(l)} \sim \exp(\lambda), \xi_{ij}^{(l)} \sim \exp(\mu_i), j = 1, \dots, s_i, i = 1, \dots, K\}.$
04. $MC(\tau_l, \text{State}, \mathcal{U}^{(l)}; \tau_{l+1}, \text{State}, I)$
05. **If** $T_l^{\min} = \xi_0^{(l)}$ and $(\bar{n}_1, \underline{n}_1) \in \mathbb{B}_3^{(l)}$ ($\mathbb{B}_3^{(l)}$ is defined in Lemma 2) **then**
06. $\bar{n}_1 = (\min\{\bar{n}_{1,1} + 1, N_1\}, \dots, \min\{\bar{n}_{1,K} + 1, N_K\}), \underline{n}_1 = \underline{n}_1,$ and $\tau_{l+1} = \tau_l + T_l^{\min}$
07. **else**
08. $MC(\tau_l, \bar{n}_1, \mathcal{U}^{(l)}; \tau_{l+1}, \bar{n}_1, I)$ and $MC(\tau_l, \underline{n}_1, \mathcal{U}^{(l)}; \tau_{l+1}, \underline{n}_1, I)$
09. Simulate $\mathcal{U}^{(l)} = \{\xi_0^{(l)} \sim \exp(\lambda), \xi_{ij}^{(l)} \sim \exp(\mu_i), j = 1, \dots, s_i, i = 1, \dots, K\}.$
10. **If** $T_l^{\min} = \xi_0^{(l)}$ and $(\bar{n}_2, \underline{n}_2) \in \mathbb{B}_3^{(l)}$ **then**
11. $\bar{n}_2 = (\min\{\bar{n}_{2,1} + 1, N_1\}, \dots, \min\{\bar{n}_{2,K} + 1, N_K\}), \underline{n}_2 = \underline{n}_2,$ and $\tau'_{l+1} = \tau'_l + T_l^{\min}$
12. **else**
13. $MC(\tau'_l, \bar{n}_2, \mathcal{U}^{(l)}; \tau'_{l+1}, \bar{n}_2, I)$ and $MC(\tau'_l, \underline{n}_2, \mathcal{U}^{(l)}; \tau'_{l+1}, \underline{n}_2, I)$
14. $l = l + 1$
15. **If** $\bar{n}_1 = \underline{n}_1$ and $\tau_l \leq \tau'_l$ **then** CoalescenceFlag = 1
16. **Else then** CoalescenceFlag = 0

Theorem 1. *Algorithm 3 returns realisations from the equilibrium of $Q(t)$.*

Theorem 1 follows from Lemma 3 and the fact that Algorithm 3 is the read-once CFTP algorithm in [15].

4. Simulation results

The number of jobs in the network at time t is $Q(t) = \max_{1 \leq i \leq K} Q_i(t)$. Let W_{mi} be the time from entering the system to exiting the service station of the i th task of the m th job. Then the total time the m th job spends in the network is $W_m = \max_{1 \leq i \leq K} W_{mi}$, which is the response time of the m th job. When $Q(t)$ is at equilibrium, we denote task quantities at each service station by $Q = (Q_1, \dots, Q_K)$ and we define W to be the response time. We use the proposed perfect sampling methods to estimate probability distributions for $Q = \max_{1 \leq i \leq K} Q_i$ and the response time W .

In this section, all Monte Carlo simulation results are based on 100 000 realisations.

4.1. Simulation results for response times

Scenario 1. In Table 1 we present the Monte Carlo estimates of $E W$ for the fork-join network with $K = 2, \lambda = 1, s_i = 1,$ and different values of μ_i and N . From Table 1 we can see that the approximations, given in [9] with $N = \infty,$ are very close to the Monte Carlo estimates using perfect simulation with a finite value of N . Obviously, the fork-join network with $N = \infty,$ on average, has a longer queue at equilibrium than the fork-join network with a finite value of N .

TABLE 1: Approximations, from [9], and estimates with $N = 300$ and $N = 100$. The estimates are based on the 100 000 realisations simulated using perfect sampling methods. The values in brackets are 100 times the standard errors of the estimates.

μ_1	μ_2	Approximation of E W with $N = \infty$	Estimates of E W with $N = 300$	Estimates of E W with $N = 100$
10.00	15.00	0.1384	0.137 945 (0.034)	0.137 299 (0.035)
10.00	20.00	0.1276	0.127 368 (0.033)	0.126 932 (0.033)
5.00	7.50	0.3057	0.303 399 (0.077)	0.301 999 (0.076)
5.00	10.00	0.2825	0.279 991 (0.076)	0.279 578 (0.074)
3.33	5.00	0.5138	0.509 327 (0.130)	0.503 866 (0.130)
3.33	6.67	0.4762	0.470 525 (0.130)	0.467 724 (0.130)
2.50	3.75	0.7822	0.766 465 (0.200)	0.760 212 (0.200)
2.50	5.00	0.7280	0.716 088 (0.200)	0.703 631 (0.200)
2.00	3.00	1.1458	1.124 944 (0.300)	1.103 505 (0.290)
2.00	4.00	1.0729	1.049 793 (0.300)	1.037 470 (0.300)




FIGURE 2: The response time distribution function with $K = 2, s_i = 1, \lambda = 1, \mu_1 = 10,$ and $\mu_2 = 20$. The dotted lines are the Monte Carlo estimates ± 20 times the standard error. The solid line is the Monte Carlo estimate with $N = 300$ and the dashed line is the approximation in [9] with $N = \infty$.

Since a longer queue results in a longer response time, we expect that the fork-join network with $N = \infty$ has a slightly larger mean response time. The results in Table 1 confirm this: the mean response time decreases as N decreases.

We can also compare the Monte Carlo estimates and the approximation in [9] for $F(t) = P(W \leq t)$. With $\mu_1 = 10$ and $\mu_2 = 20$, the results are shown in Figure 2, where the solid line (the Monte Carlo estimate) and the dashed line (the approximation given in [9]) are almost the same. If we zoom into the graph (see Figure 3), we can see that the approximation is slightly lower. With $\mu_1 = 2$ and $\mu_2 = 3$, the results are shown in Figures 4 and 5. We can see that the difference between the Monte Carlo estimate and the theoretical approximations becomes larger. Note that the simulated sample size 100 000 is very large and the 95% confidence interval will be almost the same as the estimate. Thus, in these graphs we output the Monte Carlo estimate ± 20 times its standard error (dotted lines).

Scenario 2. In Table 2 we present the Monte Carlo estimates of E W for the fork-join network with $K = 2, \lambda = 1, s_1 = 2, s_2 = 3,$ and different values of μ_i and N . From the results we




FIGURE 3: A zoomed-in view of the response time distribution function of Figure 2, with $K = 2$, $s_i = 1$, $\lambda = 1$, $\mu_1 = 10$, and $\mu_2 = 20$. The dotted lines are the Monte Carlo estimates ± 20 times the standard error. The solid line is the Monte Carlo estimate with $N = 300$ and the dashed line is the approximation in [9] with $N = \infty$.

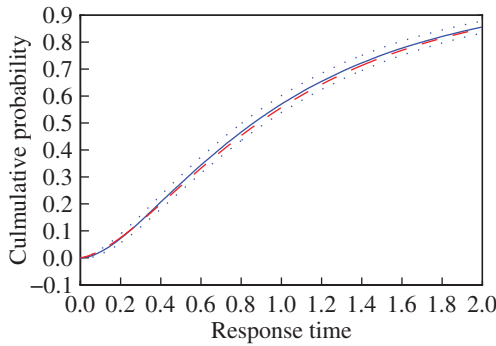


FIGURE 4: The response time distribution function with $K = 2$, $s_i = 1$, $\lambda = 1$, $\mu_1 = 2$, and $\mu_2 = 3$. The dotted lines are the Monte Carlo estimates ± 20 times the standard error. The solid line is the Monte Carlo estimate with $N = 300$ and the dashed line is the approximation in [9] with $N = \infty$.




FIGURE 5: A zoomed-in view of the response time distribution function of Figure 4, with $K = 2$, $s_i = 1$, $\lambda = 1$, $\mu_1 = 2$, and $\mu_2 = 3$. The dotted lines are the Monte Carlo estimates ± 20 times standard error. The solid line is the Monte Carlo estimate with $N = 300$ and the dashed line is the approximation in [9] with $N = \infty$.

TABLE 2: Approximations, from [9], and estimates with $N = 300$ and $N = 100$. The estimates are based on the 100 000 realisations simulated using perfect sampling methods. The values in brackets are 100 times the standard errors of the estimates.

μ_1	μ_2	Approximation of E W with $N = \infty$	Estimates of E W with $N = 300$	Estimates of E W with $N = 100$
5.00	5.00	0.2989	0.302 923 (0.071)	0.300 916 (0.071)
5.00	10.00	0.2342	0.234 578 (0.060)	0.234 795 (0.061)
5.00	15.00	0.2181	0.218 539 (0.061)	0.218 199 (0.061)
1.00	1.00	1.7208	1.756 273 (0.400)	1.746 010 (0.400)
1.00	2.00	1.4477	1.443 872 (0.370)	1.437 890 (0.370)
1.00	3.00	1.3822	1.375 673 (0.380)	1.369 130 (0.380)
0.56	0.56	9.7814	8.624 601 (2.400)	8.122 390 (2.250)
0.56	1.11	9.5285	8.310 150 (2.450)	7.807 320 (2.300)
0.56	1.67	9.4974	8.266 610 (2.470)	7.811 830 (2.290)

can see that the difference between the Monte Carlo estimate with a finite value of N and the theoretical approximation with $N = \infty$ increases as $s_i \mu_i / \lambda$ decreases. When $s_i \mu_i / \lambda$ is large (for example, $\lambda = 1$ and $\mu_i = 5$), the queue length in equilibrium is small, and $N = 300$ and $N = 100$ are large enough compared to the queue length. New jobs will not be rejected since N is very large, which is the same as the fork-join network with $N = \infty$. Therefore, in such situations, the mean response time with $N = \infty$ is almost equal to the mean response time with a large finite value of N . On the other hand, when $s_i \mu_i / \lambda$ is close to (but larger than) 1 (for example, $\lambda = 1$ and $\mu_i = 0.56$), the queue length in equilibrium may be very large, and $N = 300$ and $N = 100$ are not large enough. In such situations, new jobs may be rejected due to limited waiting space. The queue length in equilibrium with $N = \infty$ will be randomly larger than that with a finite value of N since with $N = \infty$ new jobs will never be rejected. Therefore, the mean response time with $N = \infty$ is larger than the mean response time with a finite value of N .

We plot the estimates for $F(t) = P(W \leq t)$ with $\mu_1 = 0.56, \mu_2 = 0.56$ and $\mu_1 = 5, \mu_2 = 15$ in Figures 6 and 7, respectively. We can see that the mean response times decreases as $s_i \mu_i / \lambda$




FIGURE 6: The response time distribution function with $K = 2, s_1 = 2, s_2 = 3, \lambda = 1, \mu_1 = 0.56,$ and $\mu_2 = 0.56$. The dotted lines are the Monte Carlo estimates ± 20 times the standard error.




FIGURE 7: The response time distribution function with $K = 2, s_1 = 2, s_2 = 3, \lambda = 1, \mu_1 = 5,$ and $\mu_2 = 15$. The dotted lines are the Monte Carlo estimates ± 20 times the standard error.




FIGURE 8: Histogram of response times, with $N = 300, K = 2, s_1 = 2, s_2 = 3, \lambda = 1, \mu_1 = 0.56,$ and $\mu_2 = 0.56$.

increases. Under this scenario, the distribution approximation for W is not readily available by using the methods in [9].

From Figure 8 we can see that the distribution of response times is highly skewed. Therefore, the other characteristics of the response time distribution, such as the median and quartiles, will be important. Existing methods cannot give approximations for these characteristics for multi-server fork-join networks, but we can easily estimate the median or quartiles of the response time using the simulated realisations. For example, under simulation Scenario 1 with $\mu_1 = 2, \mu_2 = 4,$ and $N = 300,$ the 0.5% and 99.5% quartiles for W are 0.000 963 and 12.065 317, respectively ($E W = 1.049 793$); under simulation Scenario 2 with $\mu_1 = \mu_2 = 0.56,$ and $N = 300,$ the 0.5% and 99.5% quartiles for W are 0.004 539 and 89.327 042, respectively ($E W = 8.624 601$). For simplicity, the results with other values of μ_i are not shown here.

4.2. Simulation for the queue length

Now we consider the simulation for the probability distribution of Q under the following three scenarios.

Scenario 3. In Table 3 we present the Monte Carlo estimates for the probability distribution of Q for the fork-join network with $K = 2, \lambda = 1, s_i = 1,$ and $\mu_1 = \mu_2 = 2$.

TABLE 3: Estimation of $P(Q = K)$. Scenario I gives the theoretical approximation from [9], and scenarios II and III give Monte Carlo estimates based on the 100 000 realisations simulated using perfect sampling methods. In scenario II $N = 300$ and in scenario III $N = 100$.

k	Scenario		
	I	II	III
0	0.3547	0.3646	0.3721
1	0.2767	0.2763	0.2796
2	0.1680	0.1679	0.1643
3	0.0951	0.0914	0.0896
4	0.0505	0.0495	0.0458
5	0.0264	0.0245	0.0245
6	0.0135	0.0128	0.0127
7	0.0070	0.0066	0.0062
8	0.0037	0.0030	0.0026
9	0.0020	0.0020	0.0012
10	0.0010	0.0007	0.0008

TABLE 4: Estimation of $P(Q = K)$. Scenario I gives the theoretical approximation from [9], and scenarios II and III give Monte Carlo estimates based on the 100 000 realisations simulated using perfect sampling methods. In scenario II $N = 300$ and in scenario III $N = 100$.

k	Scenario		
	I	II	III
0	0.4297	0.4388	0.4511
1	0.2804	0.2829	0.2806
2	0.1488	0.1451	0.1406
3	0.0739	0.0693	0.0693
4	0.0359	0.0339	0.0308
5	0.0174	0.0149	0.0146
6	0.0085	0.0078	0.0068
7	0.0041	0.0037	0.0034
8	0.0020	0.0017	0.0017
9	0.0010	0.0010	0.0007
10	0.0005	0.0005	0.0002

Scenario 4. In Table 4 we present the Monte Carlo estimates for the probability distribution of Q for the fork-join network with $M = 2, \lambda = 1, s_i = 1, \mu_1 = 2,$ and $\mu_2 = 3$.

From the results of both scenarios we can see that, with a finite value of N , the number of waiting jobs is randomly smaller than the number of waiting jobs with $N = \infty$, but the difference is small.

Scenario 5. In Figure 9 we plot the Monte Carlo estimates for the probability distribution of Q for the fork-join network with $N = 300, M = 3, \lambda = 1, s_1 = s_2 = 2, s_3 = 3, \mu_1 = \mu_2 = 1,$ and $\mu_3 = 1.5$. The approximation method in [9] does not work for this scenario.




FIGURE 9: Plot of the distribution of Q , with $K = 3, s_1 = s_2 = 2, s_3 = 3, \lambda = 1, \mu_1 = \mu_2 = 1.0$, and $\mu_3 = 1.5$.

5. Computational complexity of the CFTP algorithm

The complexity of the CFTP algorithm depends on the number of simulated $\mathcal{U}^{(l)}$ until coalescence. Thus, we only need to consider $\overline{Q}(t)$ and $\underline{Q}(t)$ at discrete time points $\tau = \{\tau_1, \tau_2, \dots\}$. For simplicity, let $\overline{Q}^{(l)} = \overline{Q}(\tau_l)$ and $\underline{Q}^{(l)} = \underline{Q}(\tau_l)$ for $l = 1, 2, \dots$. The computational complexity depends on the number of steps, v_c , needed for $\overline{Q}^{(l)}$ and $\underline{Q}^{(l)}$ to coalesce. In this section we will find a bound for $E v_c$, given $\lambda < \mu_i s_i$ for all $i < K$.

Let $d(\overline{Q}^{(l)}, \underline{Q}^{(l)})$ be the distance between the upper chain and lower chain, which is defined as

$$d(\overline{Q}^{(l)}, \underline{Q}^{(l)}) := \sum_{i=1}^K |d_i^{(l)}|,$$

where $d_i^{(l)}$ is the difference between $\overline{Q}^{(l)}$ and $\underline{Q}^{(l)}$ at the i th component. Obviously,

$$d(\overline{Q}^{(0)}, \underline{Q}^{(0)}) = \sum_i N_i,$$

since $\overline{Q}^{(0)} = (N_1, \dots, N_K)$ and $\underline{Q}^{(0)} = (0, \dots, 0)$. Define $I(\cdot)$ as the indicator function. Recalling the notation for $\mathbb{A}^{(l)}$ and $\mathbb{B}_i^{(l)}$ in Lemma 2, we have

$$\begin{aligned} & E[d(\overline{Q}^{(l+1)}, \underline{Q}^{(l+1)})] \\ &= E[E[d(\overline{Q}^{(l+1)}, \underline{Q}^{(l+1)}) \mid \overline{Q}^{(l)}, \underline{Q}^{(l)}]] \\ &= \sum_{j=1}^3 E[I((\underline{Q}^{(l)}, \overline{Q}^{(l)}) \in \mathbb{B}_j^{(l)}) E[I(\mathcal{U}^{(l)} \notin \mathbb{A}^{(l)}) d(\overline{Q}^{(l+1)}, \underline{Q}^{(l+1)}) \mid \overline{Q}^{(l)}, \underline{Q}^{(l)}]] \\ &\quad + E[E[I(\mathcal{U}^{(l)} \in \mathbb{A}^{(l)}) d(\overline{Q}^{(l+1)}, \underline{Q}^{(l+1)}) \mid \overline{Q}^{(l)}, \underline{Q}^{(l)}]] \\ &\leq P(\mathcal{U}^{(l)} \notin \mathbb{A}^{(l)}) E[I((\underline{Q}^{(l)}, \overline{Q}^{(l)}) \in \mathbb{B}_1^{(l)} \cup \mathbb{B}_2^{(l)}) d(\overline{Q}^{(l)}, \underline{Q}^{(l)})] \\ &\quad + P(\mathcal{U}^{(l)} \notin \mathbb{A}^{(l)}) E[I((\underline{Q}^{(l)}, \overline{Q}^{(l)}) \in \mathbb{B}_3^{(l)}) [d(\overline{Q}^{(l)}, \underline{Q}^{(l)}) + K]] \\ &\quad + P(\mathcal{U}^{(l)} \in \mathbb{A}^{(l)}, T_l^{\min} = \overline{T}_l < \underline{T}_l) E[d(\overline{Q}^{(l)}, \underline{Q}^{(l)}) - 1] \\ &\quad + P(\mathcal{U}^{(l)} \in \mathbb{A}^{(l)}, T_l^{\min} < \overline{T}_l \leq \underline{T}_l) E[d(\overline{Q}^{(l)}, \underline{Q}^{(l)})] \end{aligned}$$

$$\begin{aligned}
 &+ P(\mathcal{U}^{(l)} \in \mathbb{A}^{(l)}, T_l^{\min} = \bar{T}_l = \underline{T}_l) E[d(\bar{\mathcal{Q}}^{(l)}, \underline{\mathcal{Q}}^{(l)})] \\
 &= E[d(\bar{\mathcal{Q}}^{(l)}, \underline{\mathcal{Q}}^{(l)})] + K P(\mathcal{U}^{(l)} \notin \mathbb{A}^{(l)}) P((\underline{\mathcal{Q}}^{(l)}, \bar{\mathcal{Q}}^{(l)}) \in \mathbb{B}_3^{(l)}) \\
 &\quad - P(\mathcal{U}^{(l)} \in \mathbb{A}^{(l)}, T_l^{\min} = \bar{T}_l < \underline{T}_l), \tag{5}
 \end{aligned}$$

where $\bar{T}_l = \min\{\xi_0^{(l)}, \xi_{ij}^{(l)}, j = 1, \dots, \min\{\bar{\mathcal{Q}}_i^{(l)}, s_i\}, i = 1, \dots, K\}$ and $\underline{T}_l = \min\{\xi_0^{(l)}, \xi_{ij}^{(l)}, j = 1, \dots, \min\{\underline{\mathcal{Q}}_i^{(l)}, s_i\}, i = 1, \dots, K\}$. In (5) the first equality sign is because $\mathcal{U}^{(l)}$ belongs to $\mathbb{A}^{(l)}$ or $(\mathbb{A}^{(l)})^c$ and $(\underline{\mathcal{Q}}^{(l)}, \bar{\mathcal{Q}}^{(l)})$ must be in $\mathbb{B}_1^{(l)} \cup \mathbb{B}_2^{(l)} \cup \mathbb{B}_3^{(l)}$. The inequality sign in (5) follows from (3) and the definitions of $\mathbb{A}^{(l)}$ and $\mathbb{B}_i^{(l)}$.

Using properties of exponential variables, it is easy to show that

$$P(\mathcal{U}^{(l)} \in \mathbb{A}^{(l)}, T_l^{\min} = \bar{T}_l < \underline{T}_l) = P(T_l^{\min} = \bar{T}_l < \underline{T}_l) \geq \frac{\min_{i \leq K} \mu_i}{\Lambda} := \Delta, \tag{6}$$

where $\Delta > 0$ is a constant and $\Lambda = \lambda + \sum_i s_i \mu_i$.

We also have the following lemma.

Lemma 4. *Given $\lambda < \mu_i s_i$ for all $i < K$, for a large value of l , we find that $P((\underline{\mathcal{Q}}^{(l)}, \bar{\mathcal{Q}}^{(l)}) \in \mathbb{B}_3^{(l)})$ converges to 0 at an exponential rate, denoted as ε_l .*

Proof. See Appendix C.

Define $M = KN$. According to Lemma 4, for a large value of M and $l > M$, we have $\varepsilon_l < [\log(l)]^{-1}$. Then, from (5), (6), and Lemma 4, we have

$$E[d(\bar{\mathcal{Q}}^{(l+1)}, \underline{\mathcal{Q}}^{(l+1)})] \leq E[d(\bar{\mathcal{Q}}^{(l)}, \underline{\mathcal{Q}}^{(l)})](1 + O(\varepsilon_l M^{-1}) - O(M^{-1})). \tag{7}$$

Then we have the following theorem.

Theorem 2. *The expected number of coalescence steps is such that*

$$E v_c \leq O(M^2 \log(M)).$$

Proof. According to (7), for a large value of M and $l > M$, we have

$$\begin{aligned}
 P(\bar{\mathcal{Q}}^{(l)} \neq \underline{\mathcal{Q}}^{(l)}) &= P(d(\bar{\mathcal{Q}}^{(l)}, \underline{\mathcal{Q}}^{(l)}) \geq 1) \leq E[d(\bar{\mathcal{Q}}^{(l)}, \underline{\mathcal{Q}}^{(l)})] \\
 &\leq M[1 - O([M \log(M)]^{-1})]^{l-M}.
 \end{aligned}$$

The expected number of coalescence steps is such that

$$\begin{aligned}
 E v_c &= \sum_{l=1}^{\infty} P(v_c \geq l) \\
 &= \sum_{l=1}^{\infty} P(\bar{\mathcal{Q}}^{(l)} \neq \underline{\mathcal{Q}}^{(l)}) \\
 &\leq M + M O(M \log(M)) \\
 &= O(M^2 \log(M)).
 \end{aligned}$$

TABLE 5: Running time comparisons (in seconds) with $\mu_i = 1$ and $s_i = 3$.

K	N		
	500	1000	1500
5	336	654	983
10	1147	2267	3377
15	2564	4979	7254

Therefore, the expected number of steps needed to achieve coalescence is bounded by $O(M^2 \log(M))$, which means that the complexity of the algorithm increases as a polynomial function of $M = KN$. The algorithm is efficient.

We consider the running time comparisons for different values of N and K . We choose $\lambda = 1$, $\mu_i = 1$, and $s_i = 3$ for all values of i . In Table 5 we present the running times (in seconds) for simulating 10 000 realisations on a desktop with a 2.66 GHz Intel® Core™ Duo processor. We can see that the running time increases as N and K increase. The algorithm is efficient for reasonably large values of N and K .

6. Discussion

We presented a perfect sampling method based on CFTP to draw perfect simulations from the equilibrium of a fork-join network with finite capacity and zero service time at the join station. This work is important since most existing works focus on the analytic approximation for fork-join networks with infinite capacity. The accuracy of existing approximation methods is difficult to justify in some cases. For example, the analytic approximation formula of $F(t) = P(W \leq t)$ in [9] depends on some constant coefficients. The accurate values of the coefficients are sometimes found by simulation experiments (details can be found in [9]). The proposed simulation method can provide a Monte Carlo estimate, whose accuracy is controlled by the number of simulated realisations. The proposed perfect sampling method brings new insight to this area.

The proposed CFTP method makes use of the memoryless property of exponential service times and simulates continuous-time Markov chains using the *thinning* idea. When service times do not follow an exponential distribution, the *thinning* idea is not valid because the memoryless property does not hold. Thus, future research work could seek perfect sampling methods for fork-join networks with general independent service times (see [1] and [17]).

Other future research work could seek perfect sampling methods for fork-join networks with infinite capacity, i.e. $N = \infty$, although the assumption of finite capacity is reasonable in practice. This problem may be solved by extending the dominated CFTP method in [8]. Future research work in this area could also seek perfect sampling methods for fork-join networks with a nonzero service time at the join station or for networks with dynamic policies for scheduling multiple tasks.

Appendix A. Proof of Lemma 1

Algorithm 1 obviously simulates the transition probability (steps 03, 05, and 06 of Algorithm 2) from (2). We only need to show that it simulates holding times correctly.

Note that Algorithm 1 outputs $\tau = \{\tau_0, \tau_1, \dots\}$, $\{Q(\tau_0), Q(\tau_1), \dots\}$ (step 01 of Algorithm 2), and all transition times $\{t_0, t_1, \dots\}$ (a subset of τ). Since

$$T_l^{\min} = \min\{\xi_0^{(l)}, \xi_{ij}^{(l)}, j = 1, \dots, s_i, i = 1, \dots, K\}$$

is exponentially distributed with rate $\Lambda := \lambda + \sum_{i=1}^K s_i \mu_i$, then, obviously, τ_0, τ_1, \dots are jump times of a Poisson process with rate Λ .

Now consider the thinning algorithm for simulating a Poisson process with a rate $r, r < \Lambda$. Suppose that I_k is a Bernoulli random variable independent of T^{\min} and is such that $P(I_l = 1) = r/\Lambda$. Then the subsequence $\{\tau_l : I_l = 1\}$ will be the jump times of a Poisson process with rate r . From this thinning idea we can simulate an exponential variable with mean $1/r$ as follows. Suppose that $I_l = 1$, and that we simulate I_{l+1}, I_{l+2}, \dots until we obtain the first $I_{l'} = 1$. Then the interarrival time $\tau_{l'} - \tau_l$ must come from an exponential distribution with mean $1/r$. We use this idea to simulate the holding times of the Markov chain.

Suppose that $t_k = \tau_l$ and $\mathbf{Q}(\tau_l) = \mathbf{n}^{(l)}$. First consider $\mathbf{n}^{(l)} + \mathbf{1} \leq \mathbf{N}$. Since $T_l = \min\{\xi_0^{(l)}, \xi_{ij}^{(l)}, j = 1, \dots, \min\{n_i^{(l)}, s_i\}, i = 1, \dots, K\}$, we know that

$$P(T_l^{\min} = T_l) = \frac{r_{\mathbf{n}^{(l)}}}{\Lambda} = \left(\lambda + \sum_i \mu_i \min\{n_i, s_i\} \right) / \Lambda.$$

Thus, we can define $I_l = 1$ if $T_l^{\min} = T_l$ and $I_l = 0$ otherwise. Algorithm 1 generates I_{l+1}, I_{l+2}, \dots until we obtain the first $I_{l'} = 1$, i.e. $T_{l'}^{\min} = T_{l'}$ and outputs the time $t_{k+1} = \tau_{l'}$ (step 05 of Algorithm 1). Therefore, $t_{k+1} - t_k$ has an exponential distribution with mean $1/r_{\mathbf{n}^{(l)}}$ and it is the holding times of the process at $\mathbf{Q}(\tau_l) = \mathbf{n}^{(l)}$.

If $\mathbf{n}^{(l)} + \mathbf{1} \not\leq \mathbf{N}$, $T_l^{\min} = T_l = \xi_0^{(l)}$ implies that there is no jump at time $\tau_{l+1} = \tau_l + T_l^{\min}$ (step 05 of Algorithm 2 will reset $I = 0$). Therefore, the probability of running step 07 of Algorithm 2 is $P(T_l^{\min} = T_l = \xi_{ij}^{(l)})$ for some $i) = r_{\mathbf{n}^{(l)}}/\Lambda = \sum_i \mu_i \min\{n_i, s_i\}/\Lambda$. With similar arguments as above, we know that the holding times are simulated correctly.

Appendix B. Proof of Lemma 2

Algorithm 1 outputs a Poisson process with rate $\Lambda = \lambda + \sum_i s_i \mu_i$ having jump times at $\boldsymbol{\tau} = \{\tau_k, k = 1, \dots\}$. According to the updating rules of Algorithm 1, a simulated Markov chain can only have possible jumps at time τ_k .

We know that $\tau_{l+1} = \tau_l + T_l^{\min}$, where T_l^{\min} is defined in Algorithm 2. Note that the partial order is preserved in the interval (τ_l, τ_{l+1}) since there is no jump for both $\mathbf{Q}(t)$ and $\mathbf{Q}'(t)$. Therefore, result (a) is proved.

Now consider $a = T_l^{\min}$. Since the generated variables $\xi_0^{(l)}, \xi_{ij}^{(l)}$ for $j = 1, \dots, s_i$, and $i = 1, \dots, K$, will be mutually unequal with probability 1, given $T_l^{\min} \neq \xi_0^{(l)}$, we know that, for some i, j , one of the following cases must be true: (i) $T_l^{\min} = T_l' = T_l = \xi_{ij}^{(l)}$, (ii) $T_l^{\min} = T_l' = \xi_{ij}^{(l)} < T_l$, or (iii) $T_l^{\min} < T_l' \leq T_l$. According to (3) if (i) is true, $\mathbf{Q}(\tau_{l+1}) = \mathbf{n} - \mathbf{e}_i < \mathbf{Q}'(\tau_{l+1}) = \mathbf{n}' - \mathbf{e}_i$ (the partial order is preserved). If (ii) is true then $n_i' < n_i$. Therefore, according to (3), $\mathbf{Q}(\tau_{l+1}) = \mathbf{n} \leq \mathbf{Q}'(\tau_{l+1}) = \mathbf{n}' - \mathbf{e}_i$ (the partial order is preserved). If (iii) is true then neither $\mathbf{Q}(t)$ or $\mathbf{Q}'(t)$ has a jump at time $\tau_{l+1} = \tau_l + T_l^{\min}$. Therefore, the partial order is preserved since $\mathbf{Q}(\tau_{l+1}) = \mathbf{n} < \mathbf{Q}'(\tau_{l+1}) = \mathbf{n}'$. Thus, result (b) is proved.

For $a = T_l^{\min}$, if $T_l^{\min} = \xi_0^{(l)}$ then $T_l^{\min} = T_l' = T_l = \xi_0^{(l)}$. On the other hand, if $(\mathbf{Q}(\tau_l), \mathbf{Q}'(\tau_l)) \in \mathbb{B}_1^{(l)}$, i.e. $\mathbf{n}' + \mathbf{1} \leq \mathbf{N}$, we have $\mathbf{Q}(\tau_{l+1}) = \mathbf{n} + \mathbf{1} < \mathbf{Q}'(\tau_{l+1}) = \mathbf{n}' + \mathbf{1}$ according to (3) (the partial order is preserved). Result (c) is proved.

For result (d), the conditions are $T_l^{\min} = T_l' = T_l = \xi_0^{(l)}$ and $(\mathbf{Q}(\tau_l), \mathbf{Q}'(\tau_l)) \in \mathbb{B}_2^{(l)}$, i.e. $\mathbf{Q}(\tau_l) + \mathbf{1} \not\leq \mathbf{N}$. According to (3), we have $\mathbf{Q}(\tau_{l+1}) = \mathbf{n} < \mathbf{Q}'(\tau_{l+1}) = \mathbf{n}'$ (the partial order is preserved). Result (d) is proved.

For result (e), the conditions are $T_l^{\min} = T_l' = T_l = \xi_0^{(l)}$ and $(\mathbf{Q}(\tau_l), \mathbf{Q}'(\tau_l)) \in \mathbb{B}_3^{(l)}$. Then, according to (3), $\phi(\mathbf{Q}(\tau_l), \mathcal{U}^{(l)}, a) = \mathbf{n} + \mathbf{1} \leq \phi(\mathbf{Q}'(\tau_l), \mathcal{U}^{(l)}, a) = \mathbf{n}'$ may not be true (the partial order may not be preserved). Result (e) is proved.

Appendix C. Proof of Lemma 4

For the definition of $\mathbb{B}_3^{(l)}$ in Lemma 2 we have

$$P(\underline{\mathcal{Q}}^{(l)}, \overline{\mathcal{Q}}^{(l)}) \in \mathbb{B}_3^{(l)} \leq P(\overline{\mathcal{Q}}_i^{(l)} = N_i \text{ for at least one } i) \leq K \max_i P(\overline{\mathcal{Q}}_i^{(l)} = N_i).$$

Let $h = \lfloor \sqrt{l} \rfloor$, where $\lfloor \sqrt{l} \rfloor$ is the floor function. Let event \mathbb{D} be the event that from time 0 to τ_l there are more than h jumps for $\overline{\mathcal{Q}}_i(t)$.

First we show that $P(\mathbb{D}^c)$ converges to 0 at an exponential rate. At each time τ_l , $\overline{\mathcal{Q}}_i(t)$ has a jump with probability less than or equal to $(\mu_i s_i + \lambda)/\Lambda$. Therefore, we have

$$\begin{aligned} P(\mathbb{D}^c) &= P(\text{there are } k \text{ jumps at times } \tau_1, \dots, \tau_l, \text{ with } k \leq h) \\ &\leq \sum_{k=1}^h \binom{l}{k} \left(\frac{\lambda + \mu_i s_i}{\Lambda}\right)^k \left(\frac{\lambda + \sum_{j \neq i} \mu_j s_j}{\Lambda}\right)^{l-k} \\ &\leq \frac{l!}{h!(l-h)!} h \delta_1^l, \end{aligned}$$

where $0 < \delta_1 = \max_i \{(\lambda + \mu_i s_i)/\Lambda, (\lambda + \sum_{j \neq i} \mu_j s_j)/\Lambda\} < 1$. Then, using Stirling’s approximation for $l!/h!(l-h)!$, we have

$$P(\mathbb{D}^c) \leq O\left(h^h \left(\frac{l}{l-h}\right)^{l-h}\right) h^{1/2} \delta_1^l = O((he\delta_1^h)^h h^{1/2}),$$

which converges to 0 at an exponential rate.

The discrete transitions for $\overline{\mathcal{Q}}_i(t)$ are the same as the path of a simple random walk $\mathbf{R} = \{R_k, k = 0, \dots, h\}$ starting from N_i with barriers 0 and N_i . The transition probabilities of \mathbf{R} may be different at different states, but the positive jump probabilities p of \mathbf{R} are such that $p \leq \lambda/(\mu_i s_i + \lambda) < \frac{1}{2}$ when $\lambda < \mu_i s_i$. It is easy to show that, given $R_0 = N_i$, for a large value of k ,

$$\begin{aligned} P(R_k = N_i) &\leq \sum_{j=0}^{\lfloor (k+1)/2 \rfloor} \binom{k}{j} \left(\frac{\lambda}{\lambda + \mu_i s_i}\right)^{k-j} \left(\frac{\mu_i s_i}{\lambda + \mu_i s_i}\right)^j \\ &\leq O(2^k k^{-1/2}) \sum_{j=0}^{k/2} \left(\frac{\lambda}{\lambda + \mu_i s_i}\right)^{k-j} \left(\frac{\mu_i s_i}{\lambda + \mu_i s_i}\right)^j \\ &= O(2^k k^{1/2}) \left(\frac{\lambda}{\lambda + \mu_i s_i}\right)^{k/2} \left(\frac{\mu_i s_i}{\lambda + \mu_i s_i}\right)^{k/2} \\ &= O\left(k^{1/2} \left[\frac{4\lambda\mu_i s_i}{(\lambda + \mu_i s_i)^2}\right]^{k/2}\right) \\ &\leq O(k^{1/2} \delta_2^{k/2}), \end{aligned}$$

where $0 < \delta_2 = \max_i \{4\lambda\mu_i s_i/(\lambda + \mu_i s_i)^2\} < 1$. Thus, $P(R_k = N_i)$ converges to 0 at an exponential rate as $\lambda < \mu_i s_i$. Therefore,

$$\begin{aligned} P(\overline{\mathcal{Q}}_i^{(l)} = N_i) &= P(\overline{\mathcal{Q}}_i^{(l)} = N_i, \mathbb{D}) + P(\overline{\mathcal{Q}}_i^{(l)} = N_i, \mathbb{D}^c) \\ &\leq l P(R_h = N_i) + P(\mathbb{D}^c) \\ &= O(l^{5/4} \delta_2^{h/2}) + O((he\delta_1^h)^h h^{1/2}), \end{aligned}$$

which converges to 0 at an exponential rate, denoted as ε_l .

References

- [1] AYHAN, H. AND SEO, D.-W. (2001). Laplace transform and moments of waiting times in Poisson driven (max,+) linear systems. *Queueing Systems* **37**, 405–438.
- [2] BACCELLI, F., MASSEY, W. A. AND TOWSLEY, D. (1989). Acyclic fork-join queuing networks. *J. Assoc. Comput. Math.* **36**, 615–642.
- [3] BALSOMA, S., DONATIETLLO, L. AND VAN DIJK, N. M. (1998). Bound performance models of heterogeneous parallel processing systems. *IEEE Trans. Parallel Distributed Systems* **9**, 1041–1056.
- [4] DAI, H. (2008). Perfect sampling methods for random forests. *Adv. Appl. Prob.* **40**, 897–917.
- [5] FLATTO, L. (1985). Two parallel queues created by arrivals with two demands. II. *SIAM J. Appl. Math.* **45**, 861–878.
- [6] FLATTO, L. AND HAHN, S. (1984). Two parallel queues created by arrivals with two demands. I. *SIAM J. Appl. Math.* **44**, 1041–1053.
- [7] HUBER, M. (2004). Perfect sampling using bounding chains. *Ann. Appl. Prob.* **14**, 734–753.
- [8] KENDALL, W. S. AND MØLLER, J. (2000). Perfect simulation using dominating processes on ordered spaces, with application to locally stable point processes. *Adv. Appl. Prob.* **32**, 844–865.
- [9] KO, S.-S. AND SERFOZO, R. F. (2004). Response times in M/M/s fork-join networks. *Adv. Appl. Prob.* **36**, 854–871.
- [10] NELSON, R. AND TANTAWI, A. N. (1988). Approximate analysis of fork-join synchronization in parallel queues. *IEEE Trans. Software Eng.* **37**, 739–743.
- [11] PROPP, J. G. AND WILSON, D. B. (1996). Exact sampling with coupled Markov chains and applications to statistical mechanics. *Random Structures Algorithms* **9**, 223–252.
- [12] RAGHAVAN, N. R. S. AND VISWANADHAM, N. (2001). Generalized queueing network analysis of integrated supply chains. *Internat. J. Production Res.* **39**, 205–224.
- [13] ROSS, S. M. (2007). *Introduction to Probability Models*, 8th edn. Academic Press, Burlington, MA.
- [14] SQUILLANTE, M. S., ZHANG, Y., SIVASUBRAMANIAM, A. AND GAUTAM, N. (2008) Generalized parallel-server fork-join queues with dynamic task scheduling. *Ann. Operat. Res.* **160**, 227–255.
- [15] WILSON, D. B. (2000). How to couple from the past using a read-once source of randomness. *Random Structures Algorithms* **16**, 85–113.
- [16] XIA, C. H., LIU, Z., TOWSLEY, D. AND LELARGE, M. (2007). Scalability of fork/join queueing networks with blocking. In *ACM SIGMETRICS Performance Evaluation Review*, Association for Computing Machinery, New York, pp. 133–144.
- [17] ZHANG, Z. (1990). Analytical results for waiting time and system size distributions in two parallel queueing systems. *SIAM J. Appl. Math.* **50**, 1176–1193.