REPORT DOCUMENTATION PAGE

Form Approved OMB NO. 0704-0188

1. REPORT DATE (DD-MM-YYYY)	2. REPORT TYPE	3. DATES COVERED (From - To)
22-12-2014	Final Report	6-Jan-2014 - 5-Oct-2014

4. TITLE AND SUBTITLE
Final Report: STIR: Gaussian Random Fields Methods for Stochastic Fork-Join Networks with Synchronization Constraints

5a. CONTRACT NUMBER
W911NF-14-1-0019

5b. GRANT NUMBER

5c. PROGRAM ELEMENT NUMBER
611102

6. AUTHORS
Guodong Pang

5d. PROJECT NUMBER

5e. TASK NUMBER

5f. WORK UNIT NUMBER

7. PERFORMING ORGANIZATION NAMES AND ADDRESSES
Pennsylvania State University
110 Technology Center Building

University Park, PA 16802 -7000

8. PERFORMING ORGANIZATION REPORT NUMBER

9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS (ES)
U.S. Army Research Office
P.O. Box 12211
Research Triangle Park, NC 27709-2211

10. SPONSOR/MONITOR'S ACRONYM(S)
ARO

11. SPONSOR/MONITOR'S REPORT NUMBER(S)
64487-MA-II.3

12. DISTRIBUTION AVAILIBILITY STATEMENT
Approved for Public Release; Distribution Unlimited

13. SUPPLEMENTARY NOTES
The views, opinions and/or findings contained in this report are those of the author(s) and should not contrued as an official Department of the Army position, policy or decision, unless so designated by other documentation.

14. ABSTRACT
The objective of this research is to develop Gaussian random field methods for the design, modeling, analysis and control of stochastic fork-join networks (FJNs). The synchronization constraints require that tasks can only be synchronized only if all tasks of the same job are completed. The main mathematical challenge lies in the resequencing of arrival orders after service completion at each station, which requires an infinite dimensional state space to track the status of all parallel tasks for each job. It was an extremely difficult open problem. We have

15. SUBJECT TERMS
Gaussian Random Fields; Stochastic Fork-Join Networks with Synchronization Constraints; Functional Limit Theorems

16. SECURITY CLASSIFICATION OF: a. REPORT	b. ABSTRACT	c. THIS PAGE	17. LIMITATION OF ABSTRACT	15. NUMBER OF PAGES	19a. NAME OF RESPONSIBLE PERSON
UU	UU	UU	UU		Guodong Pang
					19b. TELEPHONE NUMBER
					863-100-1

Standard Form 298 (Rev 8/98)
Prescribed by ANSI Std. Z39.18

Report Title

Final Report: STIR: Gaussian Random Fields Methods for Stochastic Fork-Join Networks with Synchronization Constraints

ABSTRACT

The objective of this research is to develop Gaussian random field methods for the design, modeling, analysis and control of stochastic fork-join networks (FJNs). The synchronization constraints require that tasks can only be synchronized only if all tasks of the same job are completed. The main mathematical challenge lies in the resequencing of arrival orders after service completion at each station, which requires an infinite dimensional state space to track the status of all parallel tasks for each job. It was an extremely difficult open problem. We have developed a novel method using multiparameter sequential empirical processes driven by service vectors of parallel tasks of each job to describe the system dynamics of FJNs. We have proved functional law of large numbers and functional central limit theorems for the service and queueing dynamics for synchronization jointly in an asymptotic regime where the arrivals of jobs and the numbers of servers get large appropriately. This research has produced two research papers, under review in Mathematics of Operations Research (minor revision) and Annals of Applied Probability. One paper was on the finalist of the INFORMS 2014 JFIG Paper Competition. We have given three conference and three seminar presentations for this work.

Enter List of papers submitted or published that acknowledge ARO support from the start of the project to the date of this printing. List the papers, including journal references, in the following categories:

(a) Papers published in peer-reviewed journals (N/A for none)

Received Paper

TOTAL:

Number of Papers published in peer-reviewed journals:

(b) Papers published in non-peer-reviewed journals (N/A for none)

Received Paper

TOTAL:

Number of Papers published in non peer-reviewed journals:

(c) Presentations

1. Gaussian approximations for a fork-join network with non-exchangeable synchronization. 2014 INFROMS Manufacturing and Service Operations Management (MSOM) Conference. 06/20-06/21, 2014. (Ph.D. Student, Hongyuan Lu, presented)

2. Heavy-traffic limits for a fork-join network in the Halfin-Whitt regime. 2014 INFORMS Annual Meeting. 11/12-/11/19, 2014. (Ph.D. Student, Hongyuan Lu, presented)

3. A large-scale multi-server fork-join network with non-exchangeable synchronization. 2014 INFORMS Annual Meeting JFIG Paper Competition. 11/12-/11/19, 2014.

4. Fork-join networks with non-exchangeable synchronization in heavy traffic. University of Texas at Austin, Random Structures Seminar, 05/09/2014

5. Gaussian limits for a fork-join network with non-exchangeable synchronization in heavy traffic. Penn State University, Probability and Financial Mathematics Seminar. 02/07/2014 (Ph.D. Student, Hongyuan Lu, presented)

6. Heavy-traffic limits for a fork-join network in the Halfin-Whitt regime. Penn State University, Probability and Financial Mathematics Seminar. 10/17/2014 (Ph.D. Student, Hongyuan Lu, presented)

Number of Presentations: 6.00

Non Peer-Reviewed Conference Proceeding publications (other than abstracts):

<u>Received</u> <u>Paper</u>

TOTAL:

Number of Non Peer-Reviewed Conference Proceeding publications (other than abstracts):

Peer-Reviewed Conference Proceeding publications (other than abstracts):

<u>Received</u> <u>Paper</u>

TOTAL:

Number of Peer-Reviewed Conference Proceeding publications (other than abstracts):

(d) Manuscripts

Received		Paper
12/22/2014	2.00	Hongyuan Lu, Guodong Pang. Heavy Traffic Limits for A Fork-Join Network In the Halfin-Whitt Regime, Annals of Applied Probability (11 2014)
12/22/2014	1.00	Guodong Pang, Hongyuan Lu. Gaussian Limits for a fork-join network with non-exchangeable synchronization in heavy traffic, Mathematics of Operations Research (12 2013)

TOTAL: **2**

Number of Manuscripts:

Books

Received	Book

TOTAL:

Received	Book Chapter

TOTAL:

Patents Submitted

Patents Awarded

Awards

2014 INFORMS Junior Faculty Interest Group (JFIG) Paper Competition, Finalist.

Graduate Students

NAME	PERCENT_SUPPORTED	Discipline
Hongyuan Lu	0.50	
Yuhang Zhou	0.10	
FTE Equivalent:	**0.60**	
Total Number:	**2**	

Names of Post Doctorates

NAME	PERCENT_SUPPORTED
FTE Equivalent:	
Total Number:	

Names of Faculty Supported

NAME	PERCENT_SUPPORTED	National Academy Member
Guodong Pang	0.06	
FTE Equivalent:	**0.06**	
Total Number:	**1**	

Names of Under Graduate students supported

NAME	PERCENT_SUPPORTED
FTE Equivalent:	
Total Number:	

Student Metrics

This section only applies to graduating undergraduates supported by this agreement in this reporting period

The number of undergraduates funded by this agreement who graduated during this period: 0.00

The number of undergraduates funded by this agreement who graduated during this period with a degree in science, mathematics, engineering, or technology fields:...... 0.00

The number of undergraduates funded by your agreement who graduated during this period and will continue to pursue a graduate or Ph.D. degree in science, mathematics, engineering, or technology fields:...... 0.00

Number of graduating undergraduates who achieved a 3.5 GPA to 4.0 (4.0 max scale):...... 0.00

Number of graduating undergraduates funded by a DoD funded Center of Excellence grant for Education, Research and Engineering:...... 0.00

The number of undergraduates funded by your agreement who graduated during this period and intend to work for the Department of Defense 0.00

The number of undergraduates funded by your agreement who graduated during this period and will receive scholarships or fellowships for further studies in science, mathematics, engineering or technology fields:...... 0.00

Names of Personnel receiving masters degrees

NAME

Total Number:

Names of personnel receiving PHDs

NAME

Hongyuan LU (to be awarded in 2015

Total Number: **1**

Names of other research staff

NAME	PERCENT_SUPPORTED

FTE Equivalent:

Total Number:

Sub Contractors (DD882)

Inventions (DD882)

Scientific Progress

See Attachment

Technology Transfer

N/A

Gaussian Random Fields Methods For Fork-Join Network with Synchronization Constraints

Final Report

Scientific Progress and Accomplishments

Guodong Pang

The Harold and Inge Marcus Department of Industrial and Manufacturing Engineering

Pennsylvania State University, University Park, PA 16802

gup3@psu.edu

Forward

The proposed research is to develop Gaussian random fields methods to study fork-join networks (FJNs) with synchronization constraints. FJNs arise from many military operations, e.g., Army force deployment and counter-terrorism, where commands come from one or multiple types of operations and each operation requires multiple parallel and/or sequential tasks to be processed in service stations with multiple servers, and to be rejoined for further processing with synchronization constraints, e.g., non-exchangeability. In this research, we focus on the non-exchangeable synchronization constraint, which requires that tasks can only be synchronized only if all tasks of the same job are completed. The main mathematical challenge lies in the resequencing of arrival orders after service completion at each station, which requires an infinite dimensional state space to track the status of all parallel tasks for each job. That was an extremely difficult open problem.

We have developed a novel method using multiparameter sequential empirical processes driven by service vectors of parallel tasks of each job to describe the system dynamics of FJNs. This research has produced two research papers, focusing on a single class FJN in two asymptotic regimes, where the arrival rate of jobs and the number of servers in each station get large appropriately. We consider the number of tasks in each waiting buffer for synchronization, jointly with the number of tasks in each parallel service station and the number of synchronized jobs. In the first paper, we consider the quality-driven regime, and show that all the limiting processes are functionals of two independent processes - the limiting arrival process and a generalized Kiefer process driven by the service vector of each job. We characterize the transient and stationary distributions of the limiting processes. In the second paper, we consider the quality-and-efficiency-driven regime (Halfin-Whitt regime), and show that all the limit processes in the functional central limit theorem are also characterized via functionals of the initial limit quantities, the arrival limit process and a generalized multiparameter Kiefer process driven by the service vectors. This new framework is being further generalized to analyze fork-join networks with multiple classes of jobs, and study control, reliability and provisioning problems.

Contents

List of Figures

List of Tables

1 Statement of the Research Problem

Fork-join networks consist of a set of service stations that serve job requests simultaneously and sequentially according to pre-designated deterministic precedence constraints. Such networks have many applications in manufacturing and telecommunications [4, 16, 25, 26, 27, 43, 53, 36, 37, 49], patient flow analysis in healthcare [22, 1, 2, 57, 58], parallel computing [47, 52, 51, 32], military deployment operations [24, 56], and law enforcement systems [29]. Two types of synchronization constraints are of particular interest. One is called *exchangeable synchronization* (ES) in which tasks are not tagged with a particular job and can be synchronized for a service completion once the necessary tasks are completed. This type of synchronization constraint is often used in manufacturing systems; for example, in many assembly systems, different parts of a product are processed at separate workstations or plant locations and a product will be assembled once all of its necessary parts are completed. In this case, the parts are not tagged with a particular product, since they are standardized for the same type of product. The second type is called *non-exchangeable synchronization* (NES). Tasks are tagged with a particular job and can only be synchronized when all the parallel tasks of the same job are completed.

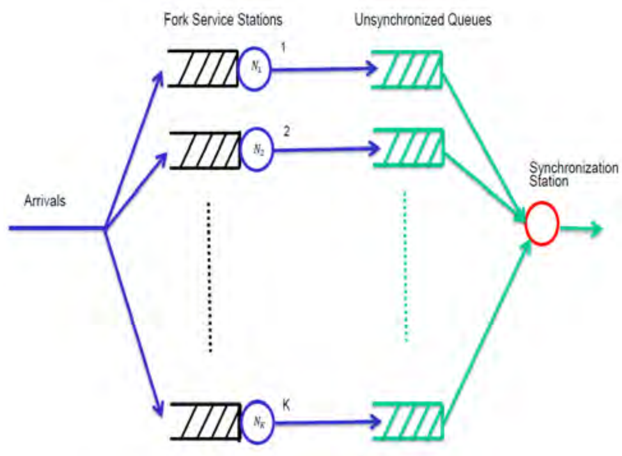

Figure 1: A fundamental fork-join network

Fork-join networks with NES are used in many applications, including healthcare systems, parallel computing, MapReducing scheduling (e.g., large-scale parallel Web search), disassembly and reassembly systems in manufacturing and so on. In patient flows of hospitals [1, 2, 22, 57, 58], the treatment and discharge processes are typical examples of fork-join networks with NES: a patient must have all test results ready before a doctor examination and these tests are conducted in different units/laboratories and can never be mixed; a patient, after the discharge decision is made, must wait for necessary procedures, pharmacy, transportation, etc., before being physically discharged. In MapReduce scheduling [11, 32, 51, 54], jobs are processed in two phases: in the map phase, a large-scale data input (e.g., Web processing data) is distributed into individual computation nodes, and each node processes one block of input data, and after the execution of all blocks of the same data input, they will be joined as an output in the reduce phase.

Despite the vast appealing applications of such networks, very little has been known about their behaviors in the many-server heavy-traffic regimes. We start considering a fundamental fork-join network model with a single class of jobs and NES, where each arriving job is forked into several parallel tasks upon arrival and each of the tasks is processed in parallel at a dedicated service station with multiple servers under the non-idling FCFS discipline. Upon service completion, each task will join a buffer associated with its service station, and wait for synchronization, such that each job is synchronized only if all of its tasks have been completed. Figure 1 depicts such a model. In this model, in addition to the service dynamics, we are interested in the waiting buffer dynamics for synchronization. One important performance measure is the response time of a job, namely, the time from arrival to synchronization. The response time may also include the time required for the synchronization process, but we do not consider that in this work. Thus, the response time includes two delays, waiting time for service and waiting time for synchronization. Since each service station can be regarded as a separate many-server queue, the waiting time for service has been well understood. However, the waiting time for synchronization, which is our focus in this paper, has not been studied. Specifically, we investigate the waiting buffer dynamics for synchronization jointly with the service dynamics.

The main mathematical challenge lies in the resequencing of the arrival orders after service completion at each service station, due to the randomness of the service times and the multi-server setting. When there is a single server in each of the parallel service station and the service discipline is FCFS, the service completion order is preserved to be the same as the arrival order of tasks in each service station, so that the two types of synchronization constraints are equivalent. However, the arrival order of tasks in each service station can be *resequenced* at the service completion epochs when the number of servers in a service station is larger than one or the service discipline is not FCFS. Resequencing has been one of the most difficult obstacles in the study of fork-join networks. Some limited work has been dedicated to the study of such challenging problems. For example, substantial efforts were dedicated to the study of the max-plus recursions [21, 3, 12]. More recently, Atar et al. [2] have studied a fork-join network with single-server service stations where tasks may reenter for service at some service stations in a Bernoulli mechanism so that the arrival orders of tasks at each service station are resequenced after service completion. They show that under a priority discipline, the system dynamics with NES is asymptotically equivalent to that with ES in the conventional (single-server) heavy-traffic regime. For a Markovian fork-join network with multiple servers, Zviran [58] shows that the system dynamics with NES is also asymptotically equivalent to that with ES in the conventional heavy-traffic regime. However, the two types of synchronization constraints lead to very different system dynamics when the service stations have many parallel servers in the Halfin-Whitt regime, as conjectured in [2, 58]. To the best of our knowledge, our work is the first to tackle the resequencing problem in non-Markovian fork-join networks with NES and multiple-server service stations in the many-server heavy-traffic regimes. We will consider both cases when each service station is operating in the quality-driven (QD) regime, or in the quality-and-efficiency-driven (QED, Halfin-Whitt) regimes.

When all the service stations operate in the QD regime, this is equivalent to a model which has infinite numbers of servers at all service stations asymptotically. To describe the system dynamics, we can start with a graphical representation as shown in Figure 2(a) for a system of two parallel tasks. At each job's arrival epoch, we mark the arrival time on the horizontal line (x-axis) and the service times of all parallel tasks on the vertical line (y-axis). At each time t, by drawing a negative forty-five degree line, we can count the numbers of tasks in each service station and each waiting buffer for synchronization. When the arrival process is Poisson, we can apply Poisson

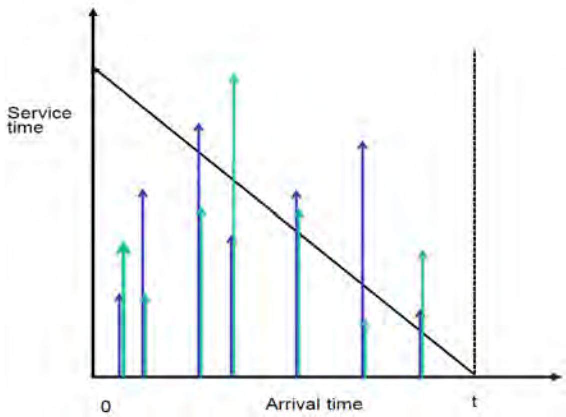

(a) The QD Regime

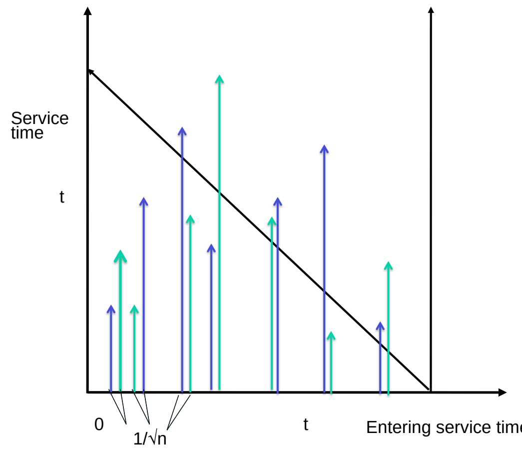

(b) The QED Regime

Figure 2: Graphical representations of the system dynamics in the QD and QED regimes

random measure theory, similarly as in the "physics" of $M/GI/\infty$ queues [14]. It can be shown that at each time t, the numbers of tasks in each service station and each waiting buffer for synchronization all have Poisson distributions and their parameter values and covariances can also be obtained; see Proposition 2.1. However, when the arrival process is more general, this Poisson random measure approach does not work, and we cannot obtain the exact distributions for these performance measures. Thus, we consider heavy-traffic approximations of the system dynamics when the arrival rate is relatively large. For that, the graphical representation in Figure 2(a) also plays an important role; see the system's dynamic equations in §2.

Here we develop a new approach to describe the system dynamics. Both the service dynamics and the waiting buffer dynamics for synchronization are represented as functionals of the multiparameter sequential empirical process driven by the service vector of all parallel tasks. Their diffusion-scaled processes converge weakly to limit processes that can be all represented as functionals of two independent processes - the limiting arrival process and the multiparameter generalized Kiefer process driven by the service vector. When the limiting arrival process is Brownian motion, we show that the aforementioned limiting processes are a multidimensional continuous Gaussian process, and thus characterize the joint transient and stationary distributions of these processes. We also study the impact of the correlation among the service vector upon these distributions.

There are several advantages with this new approach. It gives a clean and elegant representation of the limiting processes, involving only two independent stochastic processes arising from the arrival and service processes. Moreover, the characterization of the limiting processes as Gaussian and their transient and stationary distributional properties can be easily obtained. Furthermore, this new approach paves the way to study the fork-join network with all the service stations operating in the QED regime. We believe that this new approach launches a new framework to study more general fork-join networks, for example, multiclass models, and when the service vectors for parallel tasks form a stationary and weakly dependent sequence.

When all the service stations are in the QED regime, we exploit the delicate relationship between finite-server models and its corresponding infinite-server models. This was exploited to prove an FCLT for the $GI/GI/n$ queue by Reed [45]. We make an important observation that the multidimensional processes of the waiting buffer dynamics for synchronization and the service dynamics in the fork-join network can be represented through the corresponding processes in the infinite-server case. Thus, our results from the QD regime can be extended to establish the FCLT for the fork-join network in the QED regime. To illustrate, we can also use a similar graphical representation as in Figure 2(a) to describe the system dynamics. In particular, as shown in Figure 2(b), we mark the entering service times of all parallel tasks for each job on the horizontal line (x-axis), and the service times of them on the vertical line (y-axis). However, unlike the infinite-server case, tasks of the same job may not enter service simultaneously. Fortunately, it is well known that the delay for service in the QED regime is $O(1/\sqrt{n})$; see, e.g., [45, 50]. This asymptotically negligible difference among entering service times helps us to establish the FCLT for the fork-join network in the QED regime.

An important implication of our results is that the size of the waiting buffer for synchronization is of the same order as that of the total number of tasks at each service station, and thus, the waiting time for synchronization is of the same order as the service time, $O(1)$. Namely, the response time in the QED regime includes the delay for service $O(1/\sqrt{n})$, the service time $O(1)$ and the delay for synchronization $O(1)$. It remains to establish the FCLT for the (virtual) waiting time process for synchronization. More importantly, it remains to find an optimal scheduling policy that will minimize the delay for synchronization in the single-class case. We believe that our methods and results will provide useful insights towards that direction.

In the development of approximations to the fork-join system, we make a fundamental contribution to the study of multiparameter sequential empirical processes driven by random vectors. Sequential empirical processes driven by a sequence of random vectors (allowing for correlation among random variables in the vector) and their limits as generalized Kiefer processes have been studied in the statistics literature; see e.g., [42, 6, 8, 9, 13], but the convergence is proved in the space $D([0,T]^k,\mathbb{R})$ of real-valued càdlàg functions defined on $[0,T]^k$, $k \geq 2$, endowed with the generalized Skorohod J_1 topology in [35] and [48]. In our setting, it is necessary to prove the convergence in the space $D([0,T], D([0,T]^k,\mathbb{R}))$ of function-valued càdlàg functions defined on $[0,T]$, endowed with the standard Skorohod J_1 topology for $D([0,T]^k,\mathbb{R})$-valued càdlàg functions.

Literature review. Most of the literature on fork-join networks is on models with single-server service stations. We only give a brief summary here on relevant work in heavy traffic. These studies are in the conventional (single-server) heavy-traffic regime. In Varma's dissertation [53], the diffusion-scaled workload processes and unsynchronized queueing processes in some fork-join network models with ES are shown to converge weakly to certain multi-dimensional reflected Brownian motions. The stationary distributions of the system response time and the processes counting the number of tasks in unsynchronized queues are specified by some partial differential equations (PDEs). Nguyen [36] shows the diffusion-scaled processes counting the queue lengths at each service station of a single-class fork-join network model with ES converge to a reflected Brownian motion in a polyhedral cone of the nonnegative orthant. Nguyen [37] discusses the difficult challenges with multiclass fork-join models with ES. As we have noted above, for a fork-join network with feedback and NES, Atar et al. [2] show that a dynamic priority discipline achieves throughput optimality asymptotically in the conventional heavy-traffic regime, as a consequence of the asymptotic equivalence between NES and ES constraints.

Very little work has been done for fork-join networks with multi-server service stations. Ko and Serfozo [25] consider a fork-join network model with a single class of Poisson arrivals and K

parallel service stations with multiple servers at each station and exponential service times, and obtain an approximation for the distribution of the system response time in equilibrium under the NES constraint. Dai [10] provides an exact simulation algorithm to approximate the system response time in equilibrium for the same Markovian model in [25] by using a "coupling from the past" method. Zviran [58] studies optimal control of multi-server feedforward fork-join networks with exponential service times in the conventional heavy-traffic regime and shows that FCFS is asymptotically optimal and the resequencing disruption becomes asymptotically negligible. Zaied [57] calculates mean offered-load functions of fork-join networks with NES and multiple processing stages when the arrival process is time-inhomogeneous Poisson and service times for parallel tasks are independent, and studies staffing of time-varying emergency departments and synchronization delays under Markovian assumptions. Both dissertations of Zviran [58] and Zaied [57] are motivated from applications in patient flow analysis. Gurvich and Ward [17] study optimal matching policies for a pure join model (Markovian) with multiple classes of jobs under certain matching constraints.

This work contributes to the recent development for non-Markovian many-server queueing models. We only mention those that are most relevant to our work due to the large volume of papers on many-server models. Krichagina and Puhalskii [28] first observe that the system dynamics of an infinite-server queueing model can be represented by an integral functional of a sequential empirical process driven by service times. They show that the diffusion-scaled processes counting the number of jobs in the system can be approximated by a functional of a standard Kiefer process driven by service times. Pang and Whitt [39, 41] generalize that approach to establish two-parameter process limits for $G/G/\infty$ queues when the service times are i.i.d. and weakly dependent, respectively. Reed [45] and Puhalskii and Reed [44] have observed a relationship between finite-server and infinite-server queues and generalized the approach in [28] to obtain the diffusion limits for $G/GI/N$ queues in the Halfin-Whitt regime. Mandelbaum and Momcilovic [33] generalize the approach by Reed [45] to study $G/GI/N+GI$ queues with abandonment. All these papers use sequential empirical processes driven by a sequence of univariate random variables. Our approach to study fork-join networks with NES uses multiparameter sequential empirical processes driven by a sequence of i.i.d. random vectors and properties of multiparameter processes and martingales.

Notation Throughout the paper, the following notation will be used. $\mathbb{R}$ and $\mathbb{R}_+$ ($\mathbb{R}^d$ and $\mathbb{R}^d_+$, respectively) denote sets of real and real non-negative numbers (d-dimensional vectors, respectively, $d \geq 2$). $\mathbb{Z}_+$ is the set of non-negative integers. $\mathbb{N}$ denotes the set of natural numbers. For $a, b \in \mathbb{R}$, we denote $a \wedge b := \min(a, b)$ and $a \vee b := \max(a, b)$. For $x \in \mathbb{R}$, let $x^+ := \max\{x, 0\}$ and $x^- := -\min\{x, 0\}$. For any $x \in \mathbb{R}_+$, $\lfloor x \rfloor$ is used to denote the largest integer less than or equal to x. We use bold letter to denote a vector, e.g., $\boldsymbol{x} := (x_1, ..., x_N) \in \mathbb{R}^N$. $\mathbf{0}$ denotes the vector whose components are all 0. For $\boldsymbol{x}, \boldsymbol{y} \in \mathbb{R}^N$, we denote $\boldsymbol{x} \leq \boldsymbol{y}$, $\boldsymbol{x} \geq \boldsymbol{y}$ and $\boldsymbol{x} > \boldsymbol{y}$ in the componentwise sense, and let $\boldsymbol{x} \wedge \boldsymbol{y} = (x_1 \wedge y_1, ..., x_N \wedge y_N)$. We use $\mathbf{1}(A)$ to denote the indicator function of a set A. The abbreviation *a.s.* means almost surely. For any univariate distribution function $F(\cdot)$, we denote $F^c(\cdot) = 1 - F(\cdot)$. For $\boldsymbol{\alpha} \in \mathbb{R}^2_+$ and $\alpha \in \mathbb{R}_+$, we call $\Delta_{\boldsymbol{\alpha}}(\delta)$ (*resp.* $\Delta_\alpha(\delta)$) is a δ-grid of $[0, \alpha_1] \times [0, \alpha_2]$ (*resp.* $[0, \alpha]$), if $\Delta_{\boldsymbol{\alpha}}(\delta)$ (*resp.* $\Delta_\alpha(\delta)$) is a finite partition of $[0, \alpha_1] \times [0, \alpha_2]$ (*resp.* $[0, \alpha]$), where each element of the partition is the rectangle $[s_1, t_1) \times [s_2, t_2)$ (*resp.* $[s, t)$), satisfying $0 \leq s_k < t_k < \alpha_k$ for $k = 1, 2$ (*resp.* $0 \leq s < t$), and $\min_{k=1,2}(t_k - s_k) \geq \delta$ (*resp.* $t - s \geq \delta$). For two real-valued functions f and g, we write $f(x) = O(g(x))$ if $\limsup_{x\to\infty} |f(x)/g(x)| < \infty$.

All random variables and processes are defined on a common probability space $(\Omega, \mathcal{F}, P)$. For any two complete separable metric spaces $\mathcal{S}_1$ and $\mathcal{S}_2$, we denote $\mathcal{S}_1 \times \mathcal{S}_2$ as their product space, endowed with the maximum metric, i.e., the maximum of two metrics on $\mathcal{S}_1$ and $\mathcal{S}_2$. $\mathcal{S}^k$ is used to represent k-fold product space of any complete and separable metric space $\mathcal{S}$ for $k \in \mathbb{N}$. For a complete separable metric space $\mathcal{S}$, $\mathbb{D}([0,\infty),\mathcal{S})$ denotes the space of all $\mathcal{S}$-valued càdlàg functions on $[0,\infty)$, and is endowed with the Skorohod J_1 topology (see, e.g., [5, 15, 55]). Denote $\mathbb{D} \equiv \mathbb{D}([0,\infty),\mathbb{R})$. The space $\mathbb{D}([0,\infty),\mathbb{D})$, denoted as $\mathbb{D}_{\mathbb{D}}$, is endowed with the Skorohod J_1 topology, that is, both inside and outside $\mathbb{D}$ spaces are endowed with the Skorohod J_1 topology. For a complete separable metric space $\mathcal{S}$, the space $\mathbb{D}([0,\infty)^2,\mathcal{S})$ is the space of all $\mathcal{S}$-valued "continuous from above with limits from below" functions on $[0,\infty)^2$, and is endowed with the same metric as defined by [18]. $\mathbb{D}_2 \equiv \mathbb{D}([0,1]^2,\mathbb{R})$ is denoted as the space of all "continuous from above with limits from below" functions on the unit square $[0,1]^2$ in the sense of Neuhaus [35], and is endowed with the same metric $d_{\mathbb{D}_2}$ as in [35]. Weak convergence of probability measures μ_n to μ will be denoted as $\mu_n \Rightarrow \mu$. For a sequence of processes $\{\mathcal{X}^n : n \geq 1\}$ and a process $\mathcal{X}$, we use notation $\mathcal{X}^n \stackrel{d_f}{\Rightarrow} \mathcal{X}$ to denote the convergence in finite-dimensional distributions of $\mathcal{X}^n$ to $\mathcal{X}$.

2 The Infinite-Server Fork-Join Network Model

2.1 Model and Assumptions

In this section, we present a detailed description of our infinite-server fork-join network model and the assumptions. As shown in Figure 1, there is a single class of jobs, and each job is forked into K parallel tasks, $K \geq 2$. Each task is processed in a service station with multiple servers under the FCFS discipline. There is an infinite number of servers at each station. After service completion, each task will join a waiting buffer for synchronization associated with each service station, and when all tasks of the same job are completed, they will be synchronized and leave the system. Here we assume that the synchronization process takes zero amount of time.

Let $A := \{A(t) : t \geq 0\}$ be the arrival process of jobs with τ_i representing the arrival time of the i^{th} job, $i \in \mathbb{N}$. Let $\{\boldsymbol{\eta}^i : i \geq 1\}$ denote the i.i.d. service time vectors of the parallel tasks. The joint distribution of the service time vector for the i^{th} job $\boldsymbol{\eta}^i$ is $F(\boldsymbol{x}) := F(x_1, ..., x_K)$ for $x_k \geq 0$, $k = 1, ..., K$. Their marginal distributions are $F_k(x)$, for $x \geq 0$, $k = 1, ..., K$. The joint distribution of any two service times η^i_j and η^i_k is $F_{j,k}(x_j, x_k) := P(\eta^i_j \leq x_j, \eta^i_k \leq x_k)$ for $x_j, x_k \geq 0$, $j, k = 1, ..., K$. Note $F_{j,k}(\cdot,\cdot) = F_k(\cdot)$ when $j = k$ for $j, k = 1, ..., K$. We denote $F^c_{j,k}(x_j, x_k) := P(\eta^i_j > x_j, \eta^i_k > x_k) = 1 - F_j(x_j) - F_k(x_k) + F_{j,k}(x_j, x_k)$ for $x_j, x_k \geq 0$, $j, k = 1, ..., K$. Note $F^c_{j,k}(\cdot,\cdot) = F^c_k(\cdot)$ when $j = k$ for $j, k = 1, ..., K$. Let $\eta^i_{\text{m}} := \max\{\eta^i_1, ..., \eta^i_K\}$ be the maximum of the components in the service vector $\boldsymbol{\eta}^i$, and $F_{\text{m}}(x) := P(\eta^i_{\text{m}} \leq x) = F(x, ..., x)$ for $x \geq 0$. (Throughout the paper, we use subscript "m" to index quantities and processes associated with the maximum.) The service process is assumed to be independent of the arrivals. We exclude the case of perfectly positively correlated parallel services since that will lead to empty waiting buffers for synchronization.

Let $X_k := \{X_k(t) : t \geq 0\}$ be the process counting the number of tasks in service at the service station k, and $Y_k = \{Y_k(t) : t \geq 0\}$ be the process counting the number of tasks in the waiting buffer for synchronization (unsynchronized queue) after service completion at service station k, $k = 1, ..., K$. Let $S := \{S(t) : t \geq 0\}$ be the process counting the number of synchronized jobs and $D_k := \{D_k(t) : t \geq 0\}$ be the process counting the number of tasks that have completed service at station k, $k = 1, ..., K$. Denote $\boldsymbol{X} := (X_1, ..., X_K)$, $\boldsymbol{Y} := (Y_1, ..., Y_K)$ and $\boldsymbol{D} := (D_1, ..., D_K)$. We assume that the system starts empty.

Assuming that the arrival process $A(t)$ is Poisson with rate λ, by Poisson random measure theory, we can easily obtain the following properties on the processes $\boldsymbol{X}(t)$, $\boldsymbol{Y}(t)$ and $S(t)$ at each time t.

Proposition 2.1. *If the arrival process $A(t)$ is Poisson with rate λ, then at each time $t \geq 0$, for $k = 1, ..., K$, $X_k(t)$ has a Poisson distribution with rate $\lambda \int_0^t F_k^c(s)ds$, $Y_k(t)$ has a Poisson distribution with rate $\lambda \int_0^t (F_{\mathrm{m}}^c(s) - F_k^c(s))ds$, and $S(t)$ has a Poisson distribution with rate $\lambda \int_0^t F_{\mathrm{m}}(s)ds$. For each time $t \geq 0$ and $j, k = 1, ..., K$,*

$$Cov(X_j(t), X_k(t)) = \lambda \int_0^t F_{j,k}^c(s, s)ds, \tag{2.1}$$

$$Cov(Y_j(t), Y_k(t)) = \lambda \int_0^t (F_{j,k}(s, s) - F_{\mathrm{m}}(s))ds, \tag{2.2}$$

$$Cov(X_j(t), Y_k(t)) = \lambda \int_0^t (F_k(s) - F_{j,k}(s, s))ds. \tag{2.3}$$

For each time $t \geq 0$ and $k = 1, ..., K$, $S(t)$ is independent of $X_k(t)$ and $Y_k(t)$. When $K = 2$, $Y_1(t)$ and $Y_2(t)$ are independent for each $t \geq 0$.

When the arrival process $A(t)$ is general, we will obtain heavy-traffic limits for the fluid and diffusion scaled processes of $(\boldsymbol{X}, \boldsymbol{Y}, S)$ jointly. We will let the arrival rate grow large for the system to be in heavy traffic. For that, we consider a sequence of such systems indexed by n and use superscript n for the processes $A, \boldsymbol{X}, \boldsymbol{Y}, \boldsymbol{D}, S$, and the arrival times $\{\tau_i : i \geq 1\}$, but we let the service times $\{\boldsymbol{\eta}^i : i \geq 1\}$ and their distribution functions be independent of n. We make the following assumption on the arrival process A^n.

Assumption 1: FCLT for arrivals. There exist: (i) a continuous nondecreasing deterministic real-valued function $\bar{a}$ on $[0, \infty)$ with $\bar{a}(0) = 0$ and (ii) a stochastic process $\hat{A}$ with continuous sample paths, such that

$$\hat{A}^n := n^{-\frac{1}{2}}(A^n - n\bar{a}) \Rightarrow \hat{A} \quad \text{in} \quad \mathbb{D} \quad \text{as} \quad n \to \infty. \tag{2.4}$$

■

It follows from (3.6) that we have the associated FWLLN

$$\bar{A}^n := \frac{A^n}{n} \Rightarrow \bar{a} \quad \text{in} \quad \mathbb{D} \quad \text{as} \quad n \to \infty. \tag{2.5}$$

When the arrival process is renewal, the limit in (2.5) is $\bar{a}(t) = \lambda t$, for $t \geq 0$ and some positive constant λ, and the limit in (3.6) is $\hat{A} = \sqrt{\lambda c_a^2} B_a$, where c_a^2 is the squared coefficient of variation (SCV) of an interarrival time, and B_a is a standard Brownian motion (BM).

We also make a regularity assumption on the joint service-time distribution function $F(\boldsymbol{x})$.

Assumption 2: Service time distributions. The joint distribution function $F(\boldsymbol{x})$ of the service time vectors $\{\boldsymbol{\eta}^i : i \in \mathbb{N}\}$ is continuous. ■

From the graphical representation of the system dynamics in Figure 2(a), we can write, for each $t \geq 0$ and $k = 1, ..., K$,

$$X_k^n(t) = \sum_{i=1}^{A^n(t)} \mathbf{1}(\tau_i^n + \eta_k^i > t), \tag{2.6}$$

$$
\begin{aligned}
Y_k^n(t) &= \sum_{i=1}^{A^n(t)} \mathbf{1}(\tau_i^n + \eta_k^i \le t \quad \text{and} \quad \tau_i^n + \eta_{k'}^i > t \quad \text{for some } k' \neq k) \\
&= \sum_{i=1}^{A^n(t)} \big(\mathbf{1}(\tau_i^n + \eta_k^i \le t) - \mathbf{1}(\tau_i^n + \eta_{\mathrm{m}}^i \le t)\big) \\
&= \sum_{i=1}^{A^n(t)} \big(\mathbf{1}(\tau_i^n + \eta_{\mathrm{m}}^i > t) - \mathbf{1}(\tau_i^n + \eta_k^i > t)\big)\,,
\end{aligned}
\tag{2.7}
$$

$$
S^n(t) = \sum_{i=1}^{A^n(t)} \mathbf{1}(\tau_i^n + \eta_{\mathrm{m}}^i \le t) = \sum_{i=1}^{A^n(t)} \mathbf{1}(\tau_i^n + \eta_k^i \le t,\ \forall k), \tag{2.8}
$$

$$
D_k^n(t) = \sum_{i=1}^{A^n(t)} \mathbf{1}(\tau_i^n + \eta_k^i \le t). \tag{2.9}
$$

The following balanced equations hold for each $t \ge 0$ and $k = 1, ..., K$,

$$
D_k^n(t) = A^n(t) - X_k^n(t), \tag{2.10}
$$

$$
Y_k^n(t) = D_k^n(t) - S^n(t). \tag{2.11}
$$

As we have remarked in the introduction, by previous work on $G/GI/\infty$ queues [28], each individual process X_k^n and D_k^n (*resp.* S^n) can be represented by an integral of a sequential empirical process driven by a sequence of i.i.d. random variables $\{\eta_k^i : i \ge 1\}$ (*resp.* $\{\eta_{\mathrm{m}}^i : i \ge 1\}$) for each $k = 1, ..., K$. Thus, Gaussian limits for the diffusion-scaled processes X_k^n, D_k^n and S^n in heavy traffic for each k can be established, and as a consequence, a Gaussian limit for the diffusion-scaled process Y_k^n can be obtained from those of D_k^n and S^n, $k = 1, ..., K$. However, that approach does not give a characterization of the joint Gaussian distribution of the limiting processes of the diffusion-scaled processes $(\boldsymbol{X}^n, \boldsymbol{Y}^n, S^n)$.

We will represent all the processes $\boldsymbol{X}^n, \boldsymbol{Y}^n, S^n$ as integrals of a multiparameter sequential empirical process $\bar{K}^n := \{\bar{K}^n(t, \boldsymbol{x}) : t \ge 0, \boldsymbol{x} \in \mathbb{R}_+^K\}$ driven by the sequence of service vectors $\{\boldsymbol{\eta}^i : i \ge 1\}$:

$$
\bar{K}^n(t, \boldsymbol{x}) := \frac{1}{n} \sum_{i=1}^{\lfloor nt \rfloor} \mathbf{1}(\boldsymbol{\eta}^i \le \boldsymbol{x}), \quad t \ge 0, \quad \boldsymbol{x} \in \mathbb{R}_+^K. \tag{2.12}
$$

That is, we write, for $t \ge 0$ and $k = 1, ..., K$,

$$
X_k^n(t) = n \int_0^t \int_{\mathbb{R}_+^K} \mathbf{1}(s + x_k > t) d\bar{K}^n\left(\bar{A}^n(s), \boldsymbol{x}\right), \tag{2.13}
$$

$$
Y_k^n(t) = n \int_0^t \int_{\mathbb{R}_+^K} \left(\mathbf{1}(s + x_k \le t) - \mathbf{1}(s + x_j \le t,\ \forall j)\right) d\bar{K}^n\left(\bar{A}^n(s), \boldsymbol{x}\right), \tag{2.14}
$$

and

$$
S^n(t) = n \int_0^t \int_{\mathbb{R}_+^K} \mathbf{1}(s + x_j \le t,\ \forall j) d\bar{K}^n\left(\bar{A}^n(s), \boldsymbol{x}\right). \tag{2.15}
$$

The integrals in (2.13), (2.14) and (2.15) are well-defined as a Stieltjes integral for functions of bounded variation as integrators.

2.2 An FCLT for Multiparameter Sequential Empirical Processes

We present an FCLT for multiparameter sequential empirical processes $\hat{U}^n := \{\hat{U}^n(t,\boldsymbol{x}) : t \geq 0,\ \boldsymbol{x} \in [0,1]^K\}$ driven by a sequence of i.i.d. random vectors with uniform marginals:

$$\hat{U}^n(t,\boldsymbol{x}) := \frac{1}{\sqrt{n}} \sum_{i=1}^{\lfloor nt \rfloor} \left(\mathbf{1}(\boldsymbol{\xi}^i \leq \boldsymbol{x}) - H(\boldsymbol{x})\right), \quad t \geq 0, \quad \boldsymbol{x} \in [0,1]^K, \tag{2.16}$$

where for each $i \in \mathbb{N}$, $\boldsymbol{\xi}^i := (\xi_1^i, ..., \xi_K^i)$ is a vector of nonnegative random variables with continuous joint distribution function $H(\cdot)$ and uniform marginals over $[0,1]$.

The convergence for the processes $\hat{U}^n(t,\boldsymbol{x})$ is established in the space $\mathbb{D}([0,\infty),\mathbb{D}([0,1]^K,\mathbb{R}))$. We remark that this theorem is in the same spirit as Lemma 3.1 in [28], where an FCLT is proved for the two-parameter process $\hat{U}^n(t,x)$ in the univariate case in the space $\mathbb{D}([0,\infty),\mathbb{D}([0,1],\mathbb{R}))$. We generalize that result to the multivariate setting.

Theorem 2.1. *The multiparameter sequential empirical processes $\hat{U}^n(t,\boldsymbol{x})$ defined in (2.16) converge weakly to a continuous Gaussian limit,*

$$\hat{U}^n(t,\boldsymbol{x}) \Rightarrow U(t,\boldsymbol{x}) \quad \textit{in} \quad \mathbb{D}([0,\infty),\mathbb{D}([0,1]^K,\mathbb{R})) \quad \textit{as} \quad n \to \infty, \tag{2.17}$$

where $U(t,\boldsymbol{x})$ is a continuous Gaussian random field with mean function $E[U(t,\boldsymbol{x})] = 0$ and covariance function

$$Cov(U(t,\boldsymbol{x}),U(s,\boldsymbol{y})) = (t \wedge s)(H(\boldsymbol{x} \wedge \boldsymbol{y}) - H(\boldsymbol{x})H(\boldsymbol{y})), \quad t,s \geq 0, \quad \boldsymbol{x},\boldsymbol{y} \in [0,1]^K.$$

To show the FCLT for the processes $(\boldsymbol{X}^n, \boldsymbol{Y}^n, S^n)$, we define the diffusion-scaled multiparameter sequential empirical processes $\hat{K}^n := \{\hat{K}^n(t,\boldsymbol{x}) : t \geq 0, \boldsymbol{x} \in \mathbb{R}_+^K\}$ by

$$\hat{K}^n(t,\boldsymbol{x}) := \frac{1}{\sqrt{n}} \sum_{i=1}^{\lfloor nt \rfloor} \left(\mathbf{1}(\boldsymbol{\eta}^i \leq \boldsymbol{x}) - F(\boldsymbol{x})\right), \quad t \geq 0, \quad \boldsymbol{x} \in \mathbb{R}_+^K. \tag{2.18}$$

Theorem 2.1 can be applied to show an FCLT for the processes $\hat{K}^n(t,\boldsymbol{x})$. Define $\boldsymbol{F} : \mathbb{R}^K \to [0,1]^K$ with $\boldsymbol{F}(\boldsymbol{x}) = (F_1(x_1), ..., F_K(x_K))$. By Sklar's theorem [46], for any multivariate distribution function F, there exists a unique multivariate distribution function H (called "copula") with uniform marginals on $[0,1]$ such that $F(\boldsymbol{x}) = H(\boldsymbol{F}(\boldsymbol{x}))$ when the marginal distribution functions F_k, $k = 1, ..., K$, are continuous. Then, $\hat{K}^n(\cdot,\cdot)$ can be represented as a composition of $\hat{U}^n(\cdot,\cdot)$ with $\boldsymbol{F}(\cdot)$ in the second component, i.e.,

$$\hat{K}^n(t,\boldsymbol{x}) = \hat{U}^n(t,\boldsymbol{F}(\boldsymbol{x})), \quad t \geq 0, \quad \boldsymbol{x} \in \mathbb{R}_+^K.$$

Thus, it follows from Theorem 2.1 that the processes $\hat{K}^n(t,\boldsymbol{x})$ converge in distribution:

$$\hat{K}^n(t,\boldsymbol{x}) = \hat{U}^n(t,\boldsymbol{F}(\boldsymbol{x})) \Rightarrow \hat{K}(t,\boldsymbol{x}) := U(t,\boldsymbol{F}(\boldsymbol{x})) \quad \text{in} \quad \mathbb{D}([0,\infty),\mathbb{D}_K) \quad \text{as} \quad n \to \infty, \tag{2.19}$$

which implies that

$$\bar{K}^n(t,\boldsymbol{x}) \Rightarrow \bar{k}(t,\boldsymbol{x}) := tF(\boldsymbol{x}) \quad \text{in} \quad \mathbb{D}([0,\infty),\mathbb{D}_K) \quad \text{as} \quad n \to \infty. \tag{2.20}$$

2.3 FWLLN and FCLT

We define fluid-scaled processes $\bar{\boldsymbol{X}}^n$, $\bar{\boldsymbol{Y}}^n$ and $\bar{S}^n$ by

$$\bar{\boldsymbol{X}}^n := \frac{1}{n}\boldsymbol{X}^n, \quad \bar{\boldsymbol{Y}}^n := \frac{1}{n}\boldsymbol{Y}^n, \quad \bar{S}^n := \frac{1}{n}S^n. \tag{2.21}$$

The FWLLN for $(\bar{\boldsymbol{X}}^n, \bar{\boldsymbol{Y}}^n, \bar{S}^n)$ is stated in the following theorem.

Theorem 2.2 (FWLLN). *Under Assumptions* 1 *and* 2*, the fluid-scaled processes converge to deterministic fluid functions,*

$$(\bar{A}^n, \bar{\boldsymbol{X}}^n, \bar{\boldsymbol{Y}}^n, \bar{S}^n) \Rightarrow (\bar{a}, \bar{\boldsymbol{X}}, \bar{\boldsymbol{Y}}, \bar{S}) \tag{2.22}$$

in $\mathbb{D}^{2K+2}$ *as* $n \to \infty$*, where the limits are all deterministic functions:* $\bar{a}$ *is the limit in* (2.5)*, for each* $t \geq 0$*,*

$$\bar{\boldsymbol{X}}(t) := (\bar{X}_1(t), ..., \bar{X}_K(t)), \quad \bar{X}_k(t) := \int_0^t F_k^c(t-s)d\bar{a}(s), \quad \textit{for} \quad k = 1, ..., K, \tag{2.23}$$

$$\bar{\boldsymbol{Y}}(t) := (\bar{Y}_1(t), ..., \bar{Y}_K(t)), \quad \bar{Y}_k(t) := \int_0^t (F_{\mathrm{m}}^c(t-s) - F_k^c(t-s))d\bar{a}(s), \quad \textit{for} \quad k = 1, ..., K, \tag{2.24}$$

$$\bar{S}(t) := \int_0^t F_{\mathrm{m}}(t-s)d\bar{a}(s). \tag{2.25}$$

When $\bar{a}(t) = \lambda t$ *for a constant arrival rate* $\lambda > 0$ *and* $E[\eta_k^1] < \infty$ *for* $k = 1, ..., K$*,*

$$\bar{X}_k(\infty) := \lim_{t\to\infty} \bar{X}_k(t) = \lambda E[\eta_k^1], \quad k = 1, ..., K, \tag{2.26}$$

$$\bar{Y}_k(\infty) := \lim_{t\to\infty} \bar{Y}_k(t) = \lambda(E[\eta_{\mathrm{m}}^1] - E[\eta_k^1]), \quad k = 1, ..., K, \tag{2.27}$$

$$\lim_{t\to\infty} \frac{\bar{S}(t)}{t} = \lambda. \tag{2.28}$$

We define the diffusion scaling of $\boldsymbol{X}^n$, $\boldsymbol{Y}^n$ and S^n by

$$\hat{\boldsymbol{X}}^n := \sqrt{n}(\bar{\boldsymbol{X}}^n - \bar{\boldsymbol{X}}), \quad \hat{\boldsymbol{Y}}^n := \sqrt{n}(\bar{\boldsymbol{Y}}^n - \bar{\boldsymbol{Y}}), \quad \hat{S}^n := \sqrt{n}(\bar{S}^n - \bar{S}). \tag{2.29}$$

We will show the following FCLT for these diffusion-scaled processes.

Theorem 2.3 (FCLT). *Under Assumptions* 1 *and* 2*, the diffusion-scaled processes converge in distribution,*

$$(\hat{A}^n, \hat{K}^n, \hat{\boldsymbol{X}}^n, \hat{\boldsymbol{Y}}^n, \hat{S}^n) \Rightarrow (\hat{A}, \hat{K}, \hat{\boldsymbol{X}}, \hat{\boldsymbol{Y}}, \hat{S}) \tag{2.30}$$

in $\mathbb{D} \times \mathbb{D}([0,\infty), \mathbb{D}_K) \times \mathbb{D}^{2K+1}$ *as* $n \to \infty$*, where* $\hat{A}$ *is the limit in* (3.6)*,* $\hat{K}$ *is the limit in* (2.19)*, which is independent of* $\hat{A}$*, and for* $t \geq 0$ *and* $k = 1, ..., K$*,*

$$\hat{\boldsymbol{X}}(t) := \hat{\boldsymbol{M}}_1(t) + \hat{\boldsymbol{M}}_2(t), \quad \hat{\boldsymbol{M}}_i(t) := (\hat{M}_{1,i}(t), ..., \hat{M}_{K,i}(t)), \ i = 1, 2, \tag{2.31}$$

$$\hat{M}_{k,1}(t) := \int_0^t F_k^c(t-s)d\hat{A}(s) = \hat{A}(t) - \int_0^t \hat{A}(s)dF_k^c(t-s), \tag{2.32}$$

$$
\begin{aligned}
\hat{M}_{k,2}(t) &:= \int_0^t \int_{\mathbb{R}_+^K} \mathbf{1}(s+x_k>t)d\hat{K}(\bar{a}(s),\boldsymbol{x}) = -\int_0^t \int_{\mathbb{R}_+^K} \mathbf{1}(s+x_k\le t)d\hat{K}(\bar{a}(s),\boldsymbol{x}), && (2.33)\\
\hat{S}(t) &:= \hat{V}_1(t)+\hat{V}_2(t), && (2.34)\\
\hat{V}_1(t) &:= \int_0^t F_{\mathrm{m}}(t-s)d\hat{A}(s) = -\int_0^t \hat{A}(s)dF_{\mathrm{m}}(t-s), && (2.35)\\
\hat{V}_2(t) &:= \int_0^t \int_{\mathbb{R}_+^K} \mathbf{1}(s+x_j\le t,\ \forall j)d\hat{K}(\bar{a}(s),\boldsymbol{x}), && (2.36)\\
\hat{\boldsymbol{Y}}(t) &:= \hat{\boldsymbol{Z}}_1(t)+\hat{\boldsymbol{Z}}_2(t), \quad \hat{\boldsymbol{Z}}_i(t) := (\hat{Z}_{1,i}(t),...,\hat{Z}_{K,i}(t)),\ i=1,2, && (2.37)\\
\hat{Z}_{k,1}(t) &:= \int_0^t (F_k(t-s)-F_{\mathrm{m}}(t-s))d\hat{A}(s) = \int_0^t \hat{A}(s)d(F_{\mathrm{m}}(t-s)-F_k(t-s)), && (2.38)\\
\hat{Z}_{k,2}(t) &:= \int_0^t \int_{\mathbb{R}_+^K} (\mathbf{1}(s+x_k\le t)-\mathbf{1}(s+x_j\le t,\ \forall j))d\hat{K}(\bar{a}(s),\boldsymbol{x}) = -\hat{M}_{k,2}(t)-\hat{V}_2(t). && (2.39)
\end{aligned}
$$

The processes $\hat{\boldsymbol{M}}_2$, $\hat{\boldsymbol{Z}}_2$ and $\hat{V}_2$ are defined in the mean-square sense. This is in the same way as the limit process with respect to a standard Kiefer process for the $G/GI/\infty$ queue is defined in [28, 39]. The limit processes are characterized in the next subsection.

2.4 Characterization of the Limit Processes

In this section, we show the Gaussian property of the limiting processes $(\hat{\boldsymbol{X}}, \hat{\boldsymbol{Y}})$ and $\hat{S}$ when the arrival limit process is a Brownian motion.

Theorem 2.4 (Gaussian Property)**.** *Under Assumptions* 1 *and* 2*, when the arrival limit process $\hat{A}$ is a Brownian motion, i.e., $\hat{A}(t) = c_a B_a(\bar{a}(t))$ for a standard Brownian motion B_a, a positive constant $c_a > 0$ and $t \ge 0$, the limiting processes $(\hat{\boldsymbol{X}}, \hat{\boldsymbol{Y}})$ and $\hat{S}$ in Theorem* 2.3 *are well-defined continuous Gaussian processes. For each $t \ge 0$,*

$$
(\hat{\boldsymbol{X}}(t), \hat{\boldsymbol{Y}}(t)) \stackrel{d}{=} N(\mathbf{0}, \Sigma(t)), \quad \textit{and} \quad \hat{S}(t) \stackrel{d}{=} N(0, \sigma^S(t)),
$$

where for $j, k = 1, ..., K$,

$$
\sigma^X_{jk}(t) := Cov(\hat{X}_j(t), \hat{X}_k(t)) = \int_0^t \Big[F^c_{j,k}(t-s,t-s) + (c_a^2-1)F^c_j(t-s)F^c_k(t-s)\Big]d\bar{a}(s), \tag{2.40}
$$

$$
\begin{aligned}
\sigma^Y_{jk}(t) := Cov(\hat{Y}_j(t), \hat{Y}_k(t)) = \int_0^t \Big[&(F_{j,k}(t-s,t-s) - F_{\mathrm{m}}(t-s)) \\
&+ (c_a^2-1)(F_j(t-s)-F_{\mathrm{m}}(t-s))(F_k(t-s)-F_{\mathrm{m}}(t-s))\Big]d\bar{a}(s),
\end{aligned} \tag{2.41}
$$

$$
\begin{aligned}
\sigma^{XY}_{jk}(t) := Cov(\hat{X}_j(t), \hat{Y}_k(t)) = \int_0^t \Big[\, &(F_k(t-s) - F_{j,k}(t-s,t-s)) \\
&+ (c_a^2-1)\left(F^c_j(t-s)(F_k(t-s)-F_{\mathrm{m}}(t-s))\right)\Big]d\bar{a}(s),
\end{aligned} \tag{2.42}
$$

and

$$
\sigma^S(t) := Var(\hat{S}(t)) = \int_0^t F_{\mathrm{m}}(t-s)d\bar{a}(s) + (c_a^2-1)\int_0^t (F_{\mathrm{m}}(t-s))^2 d\bar{a}(s). \tag{2.43}
$$

When the arrival rate function $\bar{a}(t) = \lambda t$ for a positive constant $\lambda > 0$,

$$(\hat{\boldsymbol{X}}(t), \hat{\boldsymbol{Y}}(t)) \Rightarrow (\hat{\boldsymbol{X}}(\infty), \hat{\boldsymbol{Y}}(\infty)) \stackrel{d}{=} N(\mathbf{0}, \Sigma(\infty)) \quad as \quad t \to \infty,$$

$$\lim_{t\to\infty} t^{-1} Var(\hat{S}(t)) = \lambda c_a^2,$$

where for $j, k = 1, ..., K$,

$$\sigma_{jk}^X(\infty) := \lambda \int_0^\infty F_{j,k}^c(s,s)ds + \lambda(c_a^2 - 1)\int_0^\infty F_j^c(s)F_k^c(s)ds, \tag{2.44}$$

$$\sigma_{jk}^Y(\infty) := \lambda \int_0^\infty \Big[(F_{j,k}(s,s) - F_{\mathrm{m}}(s)) + (c_a^2 - 1)(F_j(s) - F_{\mathrm{m}}(s))(F_k(s) - F_{\mathrm{m}}(s))\Big]ds, \tag{2.45}$$

$$\sigma_{jk}^{XY}(\infty) := \lambda \int_0^\infty \Big[(F_k(s) - F_{j,k}(s,s)) + (c_a^2 - 1)\big(F_j^c(s)(F_k(s) - F_{\mathrm{m}}(s))\big)\Big]ds. \tag{2.46}$$

We make the following remarks on the Gaussian property of the limiting processes.

(i) When we set $c_a^2 = 1$, the variance and covariance formulas coincide with those in the Poisson arrival case in Proposition 2.1.

(ii) When $K = 2$ and $c_a^2 = 1$, $Cov(\hat{Y}_j(t), \hat{Y}_k(t)) = 0$, for $t \geq 0$ and $k, j = 1, ..., K$ with $k \neq j$, even if the service times of parallel tasks are correlated, since both terms inside the integral in (2.41) vanish.

(iii) We emphasize the interesting structure of the variances of $\hat{X}_k$ and $\hat{Y}_k$ and their covariances, $k = 1, ..., K$. Recall that for $G/GI/\infty$ queues [28], the steady-state variance formula of the number of jobs in the system is given as the sum of two terms, the mean and the coefficient $(c_a^2 - 1)$ multiplying an integral associated with the service time distribution; for example, when $E[\eta_k^1] < \infty$, the variance of the steady-state number of tasks in the k^{th} service station is

$$Var(\hat{X}_k(\infty)) = \lambda E[\eta_k^1] + \lambda(c_a^2 - 1)\int_0^\infty (F_k^c(s))^2 ds, \quad k = 1, ..., K.$$

It turns out that the steady-state variance formula for the number of tasks in the waiting buffer for synchronization has the same structure; for instance, when $E[\eta_k^1] < \infty$ for $k = 1, ..., K$, the variance of the steady-state waiting buffer size at the k^{th} service station is

$$Var(\hat{Y}_k(\infty)) = \lambda(E[\eta_{\mathrm{m}}^1] - E[\eta_k^1]) + \lambda(c_a^2 - 1)\int_0^\infty (F_{\mathrm{m}}^c(s) - F_k^c(s))^2 ds, \quad k = 1, ..., K.$$

The same structure also exists for the covariances between $\hat{X}_j$ and $\hat{Y}_k$, as shown in (2.42), for $k, j = 1, ..., K$.

(iv) The synchronized process does not have a Brownian motion limit, but its limiting process is Gaussian, and has the same variability as the arrival process when the arrival rate is constant. This can be also explained by regarding the synchronized process as the departure process of a $G/GI/\infty$ queue with the same arrival process and service times as the maximum of the service vectors (see [28, 39]).

■

To explore the impact of the correlation among the service times of each job's parallel tasks on the system dynamics, we consider the case when the service vector $\boldsymbol{\eta}^i$ has the joint continuous distribution function

$$F(\boldsymbol{x}) = (1-\rho)\prod_{k=1}^{K} G(x_k) + \rho G\left(\min_{k=1,...,K}\{x_k\}\right) \tag{2.47}$$

with a marginal continuous distribution function $G(\cdot)$, for $0 \le \rho < 1$, $x_k \ge 0$ and $k = 1, ..., K$. Namely, the service times at the parallel stations have the same distribution, and are symmetrically correlated with a correlation parameter $\rho \in [0, 1)$. We state the mean and covariance functions of the performance measures studied above as functions of the parameter ρ in the following corollary.

Corollary 2.1. *Under the same assumptions in Theorem* 2.4, *when the service vector* $\boldsymbol{\eta}^i$ *has the joint distribution function* F *in* (2.47), *for each* $t \ge 0$ *and* $k = 1, ..., K$, $\bar{X}_k(t)$ *and* $Var(\hat{X}_k(t))$ *are the same as in* (2.23) *and* (2.40), *respectively,*

$$\begin{aligned}
\bar{Y}_k(t) &= (1-\rho)\int_0^t \left[G(t-s)(1-(G(t-s))^{K-1})\right] d\bar{a}(s),\\
Var(\hat{Y}_k(t)) &= \int_0^t \Big[(1-\rho)G(t-s)(1-(G(t-s))^{K-1})\\
&\qquad + (1-\rho)^2(c_a^2-1)(G(t-s))^2\left(1-(G(t-s))^{K-1}\right)^2\Big] d\bar{a}(s),\\
Cov(\hat{X}_k(t), \hat{Y}_k(t)) &= (c_a^2-1)(1-\rho)\int_0^t \left[G^c(t-s)G(t-s)(1-(G(t-s))^{K-1})\right] d\bar{a}(s),
\end{aligned}$$

for $j, k = 1, ..., K$ *and* $j \ne k$,

$$\begin{aligned}
Cov(\hat{X}_j(t), \hat{X}_k(t)) &= \int_0^t \Big[(1-\rho)(G^c(t-s))^2 + \rho G^c(t-s) + (c_a^2-1)(G^c(t-s))^2\Big] d\bar{a}(s),\\
Cov(\hat{Y}_j(t), \hat{Y}_k(t)) &= \int_0^t \Big[(1-\rho)(G(t-s))^2(1-(G(t-s))^{K-2})\\
&\qquad + (1-\rho)^2(c_a^2-1)(G(t-s))^2\left(1-(G(t-s))^{K-1}\right)^2\Big] d\bar{a}(s),\\
Cov(\hat{X}_j(t), \hat{Y}_k(t)) &= (1-\rho)\int_0^t \Big[G(t-s)G^c(t-s)\\
&\qquad + (c_a^2-1)G^c(t-s)G(t-s)\left(1-(G(t-s))^{K-1}\right)\Big] d\bar{a}(s),
\end{aligned}$$

and

$$\begin{aligned}
\bar{S}(t) &= \int_0^t \Big[(1-\rho)(G(t-s))^K + \rho G(t-s)\Big] d\bar{a}(s),\\
Var(\hat{S}(t)) &= \int_0^t \left[(1-\rho)(G(t-s))^K + \rho G(t-s)\right] d\bar{a}(s)\\
&\qquad + (c_a^2-1)\int_0^t \left[(1-\rho)(G(t-s))^K + \rho G(t-s)\right]^2 d\bar{a}(s).
\end{aligned}$$

We make several remarks on the impact of the correlation among the service vector. The mean and the variance of $X_k(t)$ are not affected by the correlation, but the covariances of $X_j(t)$ and $X_k(t)$ increase linearly in ρ for $t \geq 0$ and $j, k = 1, ...K$ with $j \neq k$. The mean of $Y_k(t)$ decreases linearly in ρ and the mean of $S(t)$ increases linearly in ρ for $t \geq 0$ and $k = 1, ..., K$. The covariances of $Y_j(t)$ and $Y_k(t)$ decrease in ρ, in the order of $(1-\rho)^2$, but the covariances of $X_j(t)$ and $Y_k(t)$ decrease linearly in ρ for $t \geq 0$ and $j, k = 1, ...K$. The variance of $S(t)$ increases in ρ, in the order of ρ^2, for $t \geq 0$. The intuitive interpretation for these observations is that positive correlation makes the parallel tasks more likely to finish close to each other so that the waiting time for synchronization becomes less and more jobs are synchronized. It is also important to emphasize that the covariances of $Y_j(t)$ and $Y_k(t)$ and the covariances of $X_j(t)$ and $Y_k(t)$ decrease in different orders in the correlation parameter ρ for $t \geq 0$ and $j, k = 1, ...K$. The same observations hold for the associated steady-state performance measures.

2.5 Comparison with a fork-join network with ES

We make a comparison with an associated fork-join network with ES. We use superscript "ES" in the corresponding processes for this model. Let the arrival and service processes be the same as the model described above. The only difference is the synchronization constraint. Here tasks are not tagged with a particular job, so that whenever there are tasks completed at all parallel service stations, the oldest completed task at each waiting buffer for synchronization will be synchronized. It is evident that when the arrival process $A(t)$ is Poisson, the processes $Y_k^{ES}(t)$ and $S^{ES}(t)$ do not have a Poisson distribution at each time $t \geq 0$, $k = 1, ..., K$. In this case, for each $k = 1, ..., K$, $X_k^{n,ES}$ and $D_k^{n,ES}$ will have the same representations as in (2.6) and (2.9), but the processes $S^{n,ES}$ and $Y_k^{n,ES}$ become

$$S^{n,ES}(t) = \min_{1 \leq j \leq K} \{D_j^{n,ES}(t)\}, \quad t \geq 0, \tag{2.48}$$

and

$$Y_k^{n,ES}(t) = D_k^{n,ES}(t) - S^{n,ES}(t) = D_k^{n,ES}(t) - \min_{1 \leq j \leq K} \{D_j^{n,ES}(t)\}, \quad t \geq 0. \tag{2.49}$$

Thus, at any time, one of the waiting buffers for synchronization should be empty. It is evident that the processes $S^{n,ES}$ and $Y_k^{n,ES}$ cannot be represented as a single integral of the multiparameter sequential empirical process $\bar{K}^n$ as in equations (2.15) and (2.14), respectively.

We now discuss more on the comparison for the steady-state mean values of the fluid limits of these processes when the arrival rate is constant. In the ES model, the synchronization process S^{ES} can be represented as the minimum of the departure processes from all parallel stations, and these departure processes are dependent due to the correlation of service vector of each job. Thus, we are unable to obtain a distributional approximation of the processes S^{ES} and Y_k^{ES}, $k = 1, ..., K$. However, for each $t \geq 0$, by applying the previous results on $G/GI/\infty$ queues [28], we can obtain the mean values of the fluid limit $\bar{Y}_k^{ES}(t)$, $k = 1, ..., K$, and $\bar{S}^{ES}(t)$:

$$\begin{aligned} \bar{Y}_k^{ES}(t) &:= \lambda \left[\int_0^t F_k(s) ds - \min_{1 \leq j \leq K} \left\{ \int_0^t F_j(s) ds \right\} \right] \\ &\longrightarrow \bar{Y}_k^{ES}(\infty) := \lambda \left(\max_{1 \leq j \leq K} \{E[\eta_j^1]\} - E[\eta_k^1] \right) \quad \text{as} \quad t \to \infty, \end{aligned} \tag{2.50}$$

$$\bar{S}^{ES}(t) := \lambda \min_{1\le j\le K}\left\{\int_0^t F_j(s)ds\right\} = \lambda t - \lambda \max_{1\le j\le K}\left\{\int_0^t F_j^c(s)ds\right\}, \quad \text{and} \quad \lim_{t\to\infty}\frac{\bar{S}^{ES}(t)}{t} = \lambda. \tag{2.51}$$

Recall that the steady-state mean value of the waiting buffer for synchronization in our model $\bar{Y}_k(\infty) = \lambda\left(E[\eta_{\mathrm{m}}^1] - E[\eta_k^1]\right)$ in (2.27), $k = 1,...,K$, denoted as $\bar{Y}_k^{NES}(\infty)$ for the comparison purpose. It is evident that the average waiting buffer sizes for synchronization under NES constraint are larger than those under ES constraint, even though the total synchronization throughput rates are the same, $\lim_{t\to\infty}\bar{S}^{ES}(t)/t = \lim_{t\to\infty}\bar{S}^{NES}(t)/t = \lambda$. We also observe that when the parallel service times are perfectly positively correlated, the difference $\bar{Y}_k^{NES}(\infty) - \bar{Y}_k^{ES}(\infty)$ becomes zero for $k = 1,...,K$. We summarize this comparison result in the following proposition.

Proposition 2.2. *Under Assumptions* 1 *and* 2*, when $\bar{a}(t) = \lambda t$ for a positive arrival rate $\lambda > 0$ and $E[\eta_k^1] < \infty$ for $k = 1,...,K$,*

$$\bar{Y}_k^{NES}(\infty) - \bar{Y}_k^{ES}(\infty) = \lambda\big(E[\eta_{\mathrm{m}}^1] - \max_{1\le j\le K}\{E[\eta_j^1]\}\big) \ge 0, \quad \textit{for} \quad k = 1,...,K. \tag{2.52}$$

By the extreme value theory, if the service vector has i.i.d. components such that the service time distribution lies in the domain of attraction for Gumbel extremal distribution, then we have $a_K(\eta_{\mathrm{m}}^1 - b_K) \Rightarrow Z$ as $K \to \infty$, where Z has a Gumbel distribution, and a_K and b_K are constants depending on K; see Chapter 1 in [31]. The Gumbel distribution has cdf $P(Z \le z) = e^{-e^{-z}}$, $z \ge 0$, with mean $E[Z] = \gamma \approx 0.5772$, the Euler-Mascheroni constant, and variance $Var(Z) = \pi/\sqrt{6} \approx 1.2825$. For one example, if the service vector has i.i.d. components of an exponential distribution with rate 1, then $a_K = 1$ and $b_K = \ln K$ (see Example 1.7.2 of [31]), for $k = 1,...,K$,

$$\bar{Y}_k^{NES}(\infty) - \bar{Y}_k^{ES}(\infty) = \lambda\left(\sum_{k=1}^{K}\frac{1}{k} - 1\right) \approx \lambda(\ln(K) - 1) \quad \text{as} \quad K \to \infty. \tag{2.53}$$

For another example, if the service vector has i.i.d. components of a lognormal distribution $LN(0,1)$, we have, for $k = 1,...,K$,

$$\bar{Y}_k^{NES}(\infty) - \bar{Y}_k^{ES}(\infty) \approx \lambda(\gamma/a_K + b_K - e^{1/2}) \quad \text{as} \quad K \to \infty, \tag{2.54}$$

where a_K and b_K are (see Example 1.7.4 of [31]):

$$a_K = (2\ln K)^{1/2}\exp\left\{-(2\ln K)^{1/2} + 0.5(2\ln K)^{-1/2}(\ln\ln K + \ln(4\pi))\right\},$$

and

$$b_K = \exp\left\{(2\ln K)^{1/2} - 0.5(2\ln K)^{-1/2}(\ln\ln K + \ln(4\pi))\right\}.$$

2.6 Numerical Example

In this section, we provide a numerical example with two parallel tasks ($K = 2$), comparing our approximations with simulations. We let the arrival process be renewal with arrival rate $\lambda = 100$ and the SCV $c_a^2 = 5$. The service times of the two parallel tasks are assumed to be a bivariate Marshall-Olkin hyperexponential distribution, which is a mixture of two independent bivariate

Marshall-Olkin exponential distributions [34]. A bivariate Marshall-Olkin exponential distribution function $F_{MO}(x,y)$ for the random vector (X,Y) can be written as $F^c_{MO}(x,y) := P(\mathrm{X} > x, \mathrm{Y} > y) = \exp(-\mu_1 x - \mu_2 y - \mu_{12}(x \vee y))$, $x, y \geq 0$, where three parameters μ_1, μ_2, μ_{12} are such that the two marginals are exponential with rates $\mu_1 + \mu_{12}$ and $\mu_2 + \mu_{12}$ and their correlation $\rho = \mu_{12}/(\mu_1 + \mu_2 + \mu_{12}) \in [0,1]$. We denote $MO(\lambda_1, \lambda_2, \rho)$ for a bivariate Marshall-Olkin exponential distribution, where λ_1 and λ_2 are the rates for the marginals, and ρ is the correlation parameter, for which the parameters $\mu_1 = (\lambda_1 - \rho\lambda_2)/(1+\rho)$, $\mu_2 = (\lambda_2 - \rho\lambda_1)/(1+\rho)$ and $\mu_{12} = (\rho(\lambda_1 + \lambda_2))/(1+\rho)$. In the numerical example, we take a mixture of $MO(4/5, 1, \rho_1)$ with probability 0.4 and $MO(6/5, 6/5, \rho_2)$ with probability 0.6, such that the means of the two hyperexponential marginals are $m_{s,1} = 1$ and $m_{s,2} = 0.9$. By setting $\rho_1 = \rho_2 = 0$, we have two independent parallel service times, and by setting $\rho_1 = 0.7$ and $\rho_2 = 172/679$, we obtain that the correlation (see the correlation formula in §5.2 [40]) between the two parallel service times is equal to 0.5.

In Table 1, we show the approximation values for the mean, variance and covariance of X_k and Y_k, for $k = 1, 2$, and compare them with the corresponding simulated values. To estimate the simulated values, we simulated the system up to time 40 with 4000 independent replications starting with an empty system, which we call one experiment. In each replication, we collected data over the time interval $[20, 40]$ and formed the time average (the system tends to be in steady state in less than 5 time units). We conducted 5 independent experiments and took sample averages as estimations for simulated values. To construct the 95% confidence interval (CI), we used Student t-distribution with four degrees of freedom. The halfwidth of the 95% CI is $2.776 s_5/\sqrt{5}$, where s_5 is the sample deviation.

We make several remarks for the numerical example. First, our approximations match very well with the simulated values. Second, the size of waiting buffers for synchronization is quite large, of the same order as the number of tasks in the service stations. Third, we find that when the two parallel tasks are positively correlated, the mean and the variance of X_k's are the same as those in the independent case, while the covariance between X_1 and X_2 gets larger, the mean and the variance and covariances of Y_k's and the covariances between X_k and Y_j become smaller than those in the independent case, $j, k = 1, 2$. These are also consistent with the observations in Corollary 2.1. Note that this numerical example is more general than that considered in Corollary 2.1.

3 The multi-server fork-join network model

3.1 Model and Assumptions

In this section, we present a detailed description of our multi-server fork-join network model. We consider a fork-join network with a single class of jobs, and each job is forked into K ($K > 1$) parallel tasks. Each task is processed in a service station with finite servers under the non-idling FCFS discipline. Namely, a newly arriving task immediately gets served if there is an idle server in that station, and joins the back of the queue otherwise, and the task waiting for the longest in the queue enters service as soon as a server in that station becomes available. After service completion, each task will join a waiting buffer for synchronization associated with each service station, and when all tasks of the same job are completed, they will be synchronized and leave the system. Here we assume that the synchronization process takes zero amount of time.

Let $A := \{A(t) : t \geq 0\}$ be the arrival process of jobs after time 0. Let τ_i be the arrival time of

Table 1: Comparing approximations with simulations in a stationary model

(X_1, X_2)		$(E[X_1], E[X_2])$	$(Var(X_1), Var(X_2))$	$Cov(X_1, X_2)$
$\rho = 0$	Sim. (95% CI.)	(99.99 ± 0.17 , 89.98 ± 0.12)	(296.26 ± 0.66, 269.46 ± 0.70)	234.14 ± 0.66
	Approx.	(100.00, 90.00)	(296.00, 269.27)	233.99
$\rho = 0.5$	Sim. (95% CI.)	(99.98 ± 0.04, 89.99 ± 0.04)	(296.08 ± 0.57, 269.23 ± 0.80)	256.34 ± 0.43
	Approx.	(100.00, 90.00)	(296.00, 269.27)	256.30

(Y_1, Y_2)		$(E[Y_1], E[Y_2])$	$(Var(Y_1), Var(Y_2))$	$Cov(Y_1, Y_2)$
$\rho = 0$	Sim. (95% CI.)	(43.18 ± 0.05 , 53.20 ± 0.10)	(70.12 ± 0.20, 89.85 ± 0.40)	31.53 ± 0.30
	Approx.	(43.20, 53.20)	(70.31, 90.08)	31.55
$\rho = 0.5$	Sim. (95% CI.)	(20.89 ± 0.01, 30.88 ± 0.02)	(27.14 ± 0.15, 42.23 ± 0.35)	8.36 ± 0.07
	Approx.	(20.89, 30.89)	(27.05, 42.23)	8.31

(X, Y)		$Cov(X_1, Y_1)$	$Cov(X_1, Y_2)$	$Cov(X_2, Y_1)$	$Cov(X_2, Y_2)$
$\rho = 0$	Sim. (95% CI.)	60.80 (± 0.59)	122.87 (± 0.61)	99.21 (± 0.42)	64.56 (± 0.54)
	Approx.	61.09	123.10	99.85	64.57
$\rho = 0.5$	Sim. (95% CI.)	28.72 (± 0.33)	68.37 (± 0.73)	47.51 (± 0.42)	34.49 (± 0.44)
	Approx.	28.67	68.37	47.41	34.44

the i^{th} job, $i \in \mathbb{N}$, that is, $A(t) = \max\{i \geq 1 : \tau_i \leq t\}$ for $t > 0$ and $A(0) = 0$. Let N_k be the number of servers at service station k, $k = 1, ..., K$. Each job brings in a K-dimensional service vector, representing the service time at each service station, which can be correlated. Let $\boldsymbol{\eta}^i := (\eta^i_1, ..., \eta^i_K)$ be the service vector of the job that arrives at time τ_i, $i \in \mathbb{N}$, where η^i_k is the service time at the k^{th} service station. We assume that the sequence $\{\boldsymbol{\eta}^i : i \geq 1\}$ is i.i.d., and let the joint distribution function of $\boldsymbol{\eta}^i$ be $F(\boldsymbol{x}) = F(x_1, ..., x_K)$ for $x_k \geq 0$, $k = 1, ..., K$. Let $F^c(\boldsymbol{x}) := P(\eta^i_1 > x_1, ..., \eta^i_K > x_K)$, for $x_1, ..., x_K \geq 0$. Their marginal distributions are $F_k(\cdot)$ with mean $1/\mu_k \in (0, \infty)$, for $k = 1, ..., K$. Let $\eta^i_m := \max\{\eta^i_1, ..., \eta^i_K\}$ and $F_m(x) := P(\eta^i_m \leq x) = P(\eta^i_j \leq x, \forall j) = F(x, ..., x)$ for $x \geq 0$. (Throughout this paper, we use "m" to index quantities and processes associated with the maximum.) We make a regularity assumption on the service time distributions for the parallel tasks.

Assumption 1. *The joint distribution function $F(\boldsymbol{x})$ of the service time vector $\boldsymbol{\eta}^i$, $i \in \mathbb{N}$, is continuous.*

State Descriptors. Let $X_k := \{X_k(t) : t \geq 0\}$ be the process counting the number of tasks at the service station k, and $Y_k := \{Y_k(t) : t \geq 0\}$ be the process counting the number of tasks in the waiting buffer for synchronization (unsynchronized queue) after service completion at service station k, $k = 1, ..., K$. Denote $\boldsymbol{X} := (X_1, ..., X_K)$ and $\boldsymbol{Y} := (Y_1, ..., Y_K)$. Let $S := \{S(t) : t \geq 0\}$ be the process counting the number of synchronized jobs by each time $t \geq 0$. In addition, let $Q_k := \{Q_k(t) : t \geq 0\}$ and $B_k := \{B_k(t) : k \geq 0\}$ be the processes representing the queue length and the number of tasks in service at station k, respectively, $k = 1, ..., K$. Let $D_k := \{D_k(t) : t \geq 0\}$ be the cumulative service completion (departure) process at service station k, $k = 1, ..., K$. Denote $\boldsymbol{Q} := (Q_1, ..., Q_K)$, $\boldsymbol{B} := (B_1, ..., B_K)$, and $\boldsymbol{D} := (D_1, ..., D_K)$.

A Sequence of Systems. We consider a sequence of the above fork-join networks, indexed by superscript n and let $n \to \infty$. We assume that each service station is operating in the many-server heavy-traffic asymptotic regimes, where the arrival rate of jobs and the number of servers get large appropriately while the service time distributions are fixed. In establishing the FLLN, we allow the arrival rate to be time-dependent. In establishing the FCLT, we will assume that each service station is operating in the Halfin-Whitt (QED) regime, so that it is critically loaded with a constant arrival rate (see Assumption 4 for the precise definition). For any process $\mathcal{X}$, we use $\mathcal{X}^n$ to represent the associated process in the sequence of the fork-join networks.

Some Fundamental Flow Balance Equations. For each service station k, $k = 1, ..., K$, and for each $t \geq 0$, we have the following flow conservation equations:

$$X_k^n(t) = B_k^n(t) + Q_k^n(t), \tag{3.1}$$

$$X_k^n(t) = X_k^n(0) + A^n(t) - D_k^n(t), \tag{3.2}$$

$$Y_k^n(t) = Y_k^n(0) + D_k^n(t) - S^n(t). \tag{3.3}$$

The non-idling condition implies that for each $k = 1, ..., K$ and $t \geq 0$,

$$B_k^n(t) = X_k^n(t) \wedge N_k^n, \quad Q_k^n(t) = (X_k^n(t) - N_k^n)^+. \tag{3.4}$$

In addition, we have the following flow balance equation, for each $k, k' = 1, ..., K$, $k \neq k'$, and $t \geq 0$,

$$X_k^n(t) + Y_k^n(t) = X_{k'}^n(t) + Y_{k'}^n(t), \tag{3.5}$$

that is, the total numbers of tasks in each service station and its associated waiting buffer for synchronization are equal at all time, and are equal to the total number of jobs in the system.

3.2 Fluid Limit

In this section, we present the fluid limit for the fork-join network. We assume that the system starts from empty and allow the arrival rate to be time-dependent.

Assumption 2. *There exists a continuous nondecreasing deterministic real-valued function $\bar{a}(t)$ on $[0, \infty)$ with $\bar{a}(0) = 0$ such that*

$$\bar{A}^n(t) := n^{-1} A^n(t) \Rightarrow \bar{a}(t) \quad \textit{in} \quad \mathbb{D} \quad \textit{as} \quad n \to \infty. \tag{3.6}$$

We also make the following assumption on the numbers of servers.

Assumption 3. *For $k = 1, ..., K$, $\bar{N}_k^n := N_k^n / n \to N_k > 0$ as $n \to \infty$.*

Under the empty initial condition, we can write the processes $X_k^n(t)$, $Y_k^n(t)$, $k = 1, ..., K$, and $S^n(t)$ as

$$X_k^n(t) = \sum_{i=1}^{A^n(t)} \mathbf{1}(\tau_i^n + w_k^{n,i} + \eta_k^i > t), \quad t \geq 0, \tag{3.7}$$

$$Y^n_k(t) = \sum_{i=1}^{A^n(t)} \mathbf{1}(\tau^n_i + w^{n,i}_k + \eta^i_k \le t,\ \tau^n_i + w^{n,i}_{k'} + \eta^i_{k'} > t,\ \text{for some } k' \neq k), \quad t \ge 0, \tag{3.8}$$

$$S^n(t) = \sum_{i=1}^{A^n(t)} \mathbf{1}(\tau^n_i + w^{n,i}_k + \eta^i_k \le t,\ \forall k = 1, ..., K), \quad t \ge 0, \tag{3.9}$$

where $w^{n,i}_k$ is the waiting time of the i^{th} arrival at station k, $i \in \mathbb{N}$.

In addition, for $k = 1, ..., K$, let $E^n_k(t)$ be the number of tasks that have entered service at station k by time t, $t \ge 0$, and set $E^n_k := \{E^n_k(t) : t \ge 0\}$. Denote $\boldsymbol{E}^n := (E^n_1, ..., E^n_K)$. For each service station $k = 1, .., K$, we also have the balance equation

$$E^n_k(t) = A^n(t) - Q^n_k(t) = A^n(t) - (X^n_k(t) - N^n_k)^+, \quad t \ge 0. \tag{3.10}$$

Define the fluid-scaled processes $\bar{\mathcal{X}}^n := n^{-1}\mathcal{X}^n$ for $\mathcal{X}^n = \boldsymbol{X}^n, \boldsymbol{Y}^n, S^n, \boldsymbol{E}^n, \boldsymbol{Q}^n, \boldsymbol{B}^n, \boldsymbol{D}^n$. We now state the FLLN for the fluid-scaled processes.

Theorem 3.1. *Under Assumptions 1-3,*

$$(\bar{A}^n, \bar{\boldsymbol{X}}^n, \bar{\boldsymbol{Y}}^n, \bar{S}^n, \bar{\boldsymbol{E}}^n, \bar{\boldsymbol{Q}}^n, \bar{\boldsymbol{B}}^n, \bar{\boldsymbol{D}}^n) \Rightarrow (\bar{a}, \bar{\boldsymbol{X}}, \bar{\boldsymbol{Y}}, \bar{S}, \bar{\boldsymbol{E}}, \bar{\boldsymbol{Q}}, \bar{\boldsymbol{B}}, \bar{\boldsymbol{D}}) \tag{3.11}$$

in $\mathbb{D}^{6K+2}$ *as* $n \to \infty$*, where the limits are all deterministic functions:* $\bar{a}$ *is the limit in* (3.6)*,* $(\bar{\boldsymbol{E}}, \bar{\boldsymbol{X}}, \bar{\boldsymbol{Y}}, \bar{S})$ *is the unique solution to the following: for* $t \ge 0$ *and* $k = 1, ..., K$*,*

$$\bar{X}_k(t) = \int_0^t F^c_k(t-s)d\bar{a}(s) + \int_0^t (\bar{X}_k(t-s) - N_k)^+ dF_k(s), \tag{3.12}$$

$$\bar{E}_k(t) = \bar{a}(t) - (\bar{X}_k(t) - N_k)^+, \tag{3.13}$$

$$\bar{S}(t) = \int_0^t \cdots \int_0^t \left(\min_{k=1,...,K} \left\{ \bar{E}_k(t - s_k) \right\} \right) dF(s_1, ..., s_K), \tag{3.14}$$

$$\bar{Y}_k(t) = \int_0^t F_k(t-s)d\bar{a}(s) - \int_0^t (\bar{X}_k(t-s) - N_k)^+ dF_k(s) - \bar{S}(t), \tag{3.15}$$

and the limits $\bar{\boldsymbol{Q}}$*,* $\bar{\boldsymbol{B}}$ *and* $\bar{\boldsymbol{D}}$ *satisfy*

$$\bar{Q}_k(t) = (\bar{X}_k(t) - N_k)^+, \quad \bar{B}_k(t) = \bar{X}_k(t) \wedge N_k, \quad \bar{D}_k(t) = \bar{a}(t) - \bar{X}_k(t). \tag{3.16}$$

It is easy to check that for each $k = 1, ..., K$, the limit $\bar{X}_k(t)$ also satisfies the following equation:

$$\bar{X}_k(t) = \bar{a}(t) - \int_0^t \bar{E}_k(t-s)dF_k(s), \quad t \ge 0. \tag{3.17}$$

When $\bar{a}(t) = \int_0^t \lambda(s)ds$ and the service times are exponential (independent or dependent), where $\lambda(\cdot)$ is a positive function, for each $k = 1, ..., K$, the fluid limit $\bar{X}_k$ in (3.12) and (3.17) becomes an ordinary differential equation (ODE) [38], but the fluid limit $\bar{Y}_k$ in (3.15) does not have an ODE representation. We remark that the fluid limit $\bar{X}_k$ for each $k = 1, ..., K$ depends only on the marginal distribution F_k, while the fluid limits $\bar{Y}_k$, $k = 1, ..., K$, and $\bar{S}$ depend on the joint distribution F.

However, as the FCLT (Theorem 3.2) below shows, the limits for all these processes in the diffusion scale will depend on the joint distribution F.

When the arrival rate is constant and each service station is underloaded or critically loaded, we give a corollary on the steady states of the fluid limits. The proof follows from a direct calculation and is omitted. It is evident that correlation among service times of parallel tasks only affects the steady state of $\bar{\boldsymbol{Y}}$ but not that of $\bar{\boldsymbol{X}}$.

Corollary 3.1. *Under Assumptions 1-3, if the arrival rate is constant, $\bar{a}(t) = \lambda t$, for λ satisfying $0 < \lambda \leq N_k\mu_k$ for all $k = 1, ..., K$,*

$$(\bar{\boldsymbol{X}}(t), \bar{\boldsymbol{Y}}(t), \bar{\boldsymbol{Q}}(t), \bar{\boldsymbol{B}}(t)) \to (\bar{\boldsymbol{X}}(\infty), \bar{\boldsymbol{Y}}(\infty), \bar{\boldsymbol{Q}}(\infty), \bar{\boldsymbol{B}}(\infty)) \quad as \quad t \to \infty,$$

and

$$\frac{1}{t}(\bar{\boldsymbol{D}}(t), \bar{\boldsymbol{E}}(t), \bar{S}(t)) \to \boldsymbol{\lambda} := (\lambda, ..., \lambda) \quad as \quad t \to \infty,$$

where

$$\bar{X}_k(\infty) = \bar{B}_k(\infty) = \lambda E[\eta_k^1] = \lambda/\mu_k, \quad \bar{Y}_k(\infty) = \lambda(E[\eta_m^1] - E[\eta_k^1]), \quad \bar{Q}_k(\infty) = 0.$$

3.2.1 Numerical Examples

We give two numerical examples to show the effectiveness of fluid approximations comparing with simulations, when $K = 2$. We let the arrival process be Poisson with time-varying rate $\lambda(t) = 200 + 120\sin(t)$, $t \geq 0$. The numbers of servers in stations 1 and 2 are $N_1 = 300$ and $N_2 = 340$, respectively. In the first numerical example, the service times of the two parallel tasks are assumed to have a bivariate Marshall-Olkin exponential distribution [34]. For our first numerical example, we set the service times to be $MO(1, 0.9, \rho)$ such that the service times of the two parallel tasks have exponential marginals with means 1 and 10/9 in stations 1 and 2, respectively, and their correlation is ρ. The numerical results with $\rho = 0$ and $\rho = 0.5$ are provided in Figure 3(a), marked with "ind." and "corr.", respectively. In the second numerical example, we let the service times of the two parallel tasks have a bivariate Marshall-Olkin hyperexponential distribution [40], which is a mixture of two independent bivariate Marshall-Olkin exponential distributions. Specifically, we take a mixture of $MO(4/5, 1, \rho_1)$ with probability 0.4 and $MO(6/5, 27/32, \rho_2)$ with probability 0.6, such that the two parallel service times have hyperexponential marginals with the same means as the first example. By setting $\rho_1 = \rho_2 = 0$, we have two independent parallel service times, and by setting $\rho_1 = 0.7$ and $\rho_2 = 521/1232$, we get the correlation between the two parallel service times to be 0.5. In Figure 3(b), we show the numerical results with $\rho = 0$ ("ind.") and $\rho = 0.5$ ("corr."). To calculate the simulated values, we simulated the system up to time 20 with 500 independent replications starting with an empty system. We make two remarks from numerical results. First, the fluid approximations match very well with the simulated results. Second, the positive correlation among parallel service times does not affect $\bar{X}_k$, but reduces $\bar{Y}_k$, for $k = 1, 2$.

3.3 FCLT in the Halfin-Whitt regime

In this section, we study the fork-join network with NES in the Halfin-Whitt regime, which requires that each service station operates in a critically loaded regime asymptotically. Specifically, we

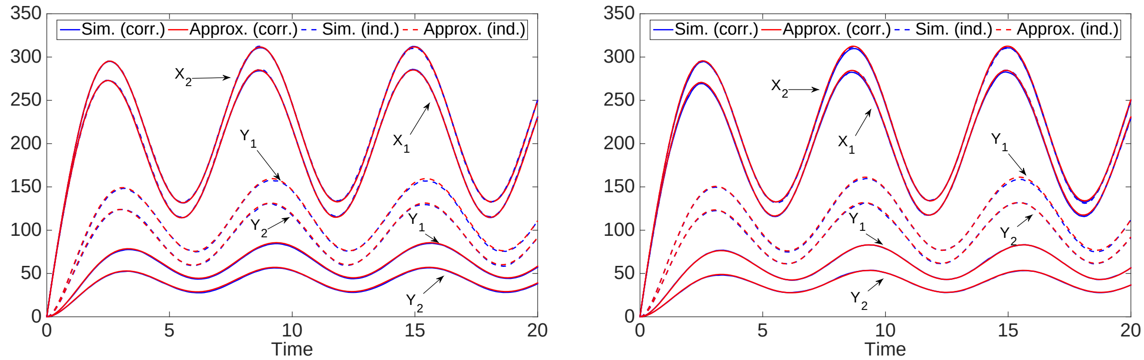

(a) Exponential service times (b) Hyperexponential service times

Figure 3: Comparison of fluid approximations with simulations.

assume the following. Let λ^n be the arrival rate of jobs such that $\bar{\lambda}^n := \lambda^n/n \to \lambda > 0$ as $n \to \infty$, and set $N_k^n := nN_k$, where $N_k \in \mathbb{N}$, and $\rho_k^n := \lambda^n/(\mu_k N_k^n)$ for each $k = 1, ..., K$.

Assumption 4. *For each $k = 1, ..., K$, $\lambda = N_k \mu_k$ and $\sqrt{n}(1 - \rho_k^n) \to \beta_k > 0$, as $n \to \infty$.*

The arrival processes $A^n = \{A^n(t) : t \geq 0\}$ satisfy an FCLT.

Assumption 5. *There exists a stochastic process $\hat{A}$ with continuous sample paths satisfying*

$$\hat{A}^n(t) := \frac{A^n(t) - \lambda^n t}{\sqrt{n}} \Rightarrow \hat{A}(t) \quad in \quad \mathbb{D} \quad as \quad n \to \infty. \tag{3.18}$$

It follows from (3.18) that we have the associated FLLN:

$$\bar{A}^n(t) \Rightarrow \lambda t \quad \text{in} \quad \mathbb{D} \quad \text{as} \quad n \to \infty. \tag{3.19}$$

We now describe the non-empty initial conditions. Due to the complexity from initial conditions, we focus on the case of $K = 2$, but our approach can be extended to $K > 2$. For convenience, we use the notation k' to denote its counterpart, i.e., $k' = 1$ ($k' = 2$, respectively) if $k = 2$ ($k = 1$, respectively), for $k = 1, 2$. At time $0-$, there are $X_k^n(0)$ tasks at service station k, and $Y_k^n(0)$ tasks in its associated waiting buffer for synchronization, for $k = 1, 2$. Let $\boldsymbol{X}^n(0) := (X_1^n(0), X_2^n(0))$ and $\boldsymbol{Y}^n(0) := (Y_1^n(0), Y_2^n(0))$. Recall the flow balance equation (3.5). At time $0-$,

$$X_k^n(0) + Y_k^n(0) = X_{k'}^n(0) + Y_{k'}^n(0), \quad k = 1, 2, \tag{3.20}$$

which is equal to the number of jobs in the system. Note that $X_k^n(0) \geq Y_{k'}^n(0)$ for each $k = 1, 2$, since tasks in the waiting buffer associated with station k' for synchronization must be in station k, either in service or in queue. Let $B_k^n(0) := \min(X_k^n(0), N_k^n)$ and $Q_k^n(0) := (X_k^n(0) - N_k^n)^+$ be the number of tasks in service (busy servers) and the queue length at station k at time $0-$, respectively, $k = 1, 2$. We also assume that $Y_{k'}^n(0) \leq B_k^n(0)$ for $k = 1, 2$. This is not a restrictive assumption, because in the Halfin-Whitt regime, waiting times for service at each station are $O(1/\sqrt{n})$ and service times are $O(1)$, and jobs that have completed tasks in one station and joined its waiting buffer for synchronization have their associated tasks receiving service in the other station with probability one asymptotically.

Let $J^n(0) := \min_{k=1,2}\{B_k^n(0) - Y_{k'}^n(0)\}$ be the number of jobs whose both tasks are in service at time $0-$. Then $Z_k^n(0) := B_k^n(0) - Y_{k'}^n(0) - J^n(0)$ represents the number of jobs in the system at time $0-$ whose task k is in service but whose task k' is in queue waiting for service, $k = 1,2$. Let $I^n(0) := Q_1^n(0) \wedge Q_2^n(0)$ be the number of jobs (if any) whose both tasks are in queue at their service stations at time $0-$. Then $R_k^n(0) := Q_k^n(0) - I^n(0)$ represents the number of jobs (if any) whose task k is waiting in queue for service while whose task k' is in service, $k = 1,2$. (Note that our assumption above implies that if a job is waiting in queue at station k, its parallel task can be either in queue or in service at station k'.) By our definition, we can see that $Z_k^n(0) = R_{k'}^n(0)$, $k = 1,2$. Set $\boldsymbol{R}^n(0) := (R_1^n(0), R_2^n(0))$ and $\boldsymbol{Z}^n(0) := (Z_1^n(0), Z_2^n(0))$. We also obtain a decomposition for $X_k^n(0)$:

$$X_k^n(0) = B_k^n(0) + Q_k^n(0) = Y_{k'}^n(0) + J^n(0) + Z_k^n(0) + I^n(0) + R_k^n(0), \quad k = 1,2. \tag{3.21}$$

We let $\{\tilde{w}_k^{n,i} : i = 1, ..., Q_k^n(0)\}$ be the sequence of remaining waiting times of the tasks in station k at time $0-$, $k = 1,2$. It is in the order of their positions in queue: $\tilde{w}_k^{n,1}$ is the remaining waiting time of the task in the front of the queue while $\tilde{w}_k^{n,Q_k^n(0)}$ is that for the task in the end of the queue at station k at time $0-$, $k = 1,2$. Let $\{\tilde{\eta}_k^i : i = 1, ..., B_k^n(0)\}$ be the sequence of remaining service times of the tasks in station k at time $0-$, for $k = 1,2$. Let $\{\eta_k^{i,Q} : i = 1, ..., Q_k^n(0)\}$ be the sequence of service times of the tasks in station k that are in queue at time $0-$, $k = 1,2$. Without abuse of notation, we use $\{\tilde{\eta}_k^{i,Y_k} : i = 1, ..., Y_{k'}^n(0)\}$, $\{\tilde{\eta}_k^{i,J} : i = 1, ..., J^n(0)\}$ and $\{\tilde{\eta}_k^{i,Z} : i = 1, ..., Z_k^n(0)\}$, which are partitioning subsets of $\{\tilde{\eta}_k^i : i = 1, ..., B_k^n(0)\}$, to represent the remaining service times of the tasks in station k at time $0-$ corresponding to the quantities $Y_{k'}^n(0)$, $J^n(0)$ and $Z_k^n(0)$, respectively, $k = 1,2$. Similarly, we use $\{\tilde{w}_k^{n,i,I} : i = 1, ..., I^n(0)\}$ and $\{\tilde{w}_k^{n,i,R} : i = 1, ..., R_k^n(0)\}$, which are partitioning subsets of $\{\tilde{w}_k^{n,i} : i = 1, ..., Q_k^n(0)\}$, to represent the remaining waiting times of the tasks in station k at time $0-$ corresponding to the quantities $I^n(0)$ and $R_k^n(0)$, respectively, $k = 1,2$. Finally, we use $\{\eta_k^{i,I} : i = 1, ..., I^n(0)\}$ and $\{\eta_k^{i,R} : i = 1, ..., R_k^n(0)\}$, which are partitioning subsets of $\{\eta_k^{i,Q} : i = 1, ..., Q_k^n(0)\}$, to represent the service times of the tasks in station k corresponding to the quantities $I^n(0)$ and $R_k^n(0)$ in queue at time $0-$, respectively, $k = 1,2$. We assume that these initial quantities are independent of the arrival process A^n and the service times of new arrivals after time 0.

We can now give a representation for the processes $\boldsymbol{X}^n$, $\boldsymbol{Y}^n$ and S^n: for $t \geq 0$ and $k = 1,2$,

$$X_k^n(t) = \sum_{i=1}^{B_k^n(0)} \mathbf{1}(\tilde{\eta}_k^i > t) + \sum_{i=1}^{Q_k^n(0)} \mathbf{1}(\tilde{w}_k^{n,i} + \eta_k^{i,Q} > t) + \sum_{i=1}^{A^n(t)} \mathbf{1}(\tau_i^n + w_k^{n,i} + \eta_k^i > t), \tag{3.22}$$

$$\begin{aligned} S^n(t) = {} & \sum_{i=1}^{Y_2^n(0)} \mathbf{1}(\tilde{\eta}_1^{i,Y_1} \leq t) + \sum_{i=1}^{Y_1^n(0)} \mathbf{1}(\tilde{\eta}_2^{i,Y_2} \leq t) + \sum_{i=1}^{J^n(0)} \mathbf{1}(\tilde{\eta}_j^{i,J} \leq t, \forall j) \\ & + \sum_{i=1}^{Z_1^n(0)} \mathbf{1}(\tilde{\eta}_1^{i,Z} \leq t, \tilde{w}_2^{n,i,R} + \eta_2^{i,R} \leq t) + \sum_{i=1}^{Z_2^n(0)} \mathbf{1}(\tilde{w}_1^{n,i,R} + \eta_1^{i,R} \leq t, \tilde{\eta}_2^{i,Z} \leq t) \\ & + \sum_{i=1}^{I^n(0)} \mathbf{1}(\tilde{w}_j^{n,i,I} + \eta_j^{i,I} \leq t, \forall j) + \sum_{i=1}^{A^n(t)} \mathbf{1}(\tau_i^n + w_j^{n,i} + \eta_j^i \leq t, \forall j), \end{aligned} \tag{3.23}$$

and

$$Y_k^n(t) = Y_k^n(0) + X_k^n(0) + A^n(t) - X_k^n(t) - S^n(t). \tag{3.24}$$

We use the convention that $\sum_{i=1}^{0} \equiv 0$ throughout the paper.

We impose the following assumptions on the initial quantities.

Assumption 6. *There exists* $(\bar{Y}_1(0), \bar{Y}_2(0)) \in \mathbb{R}_+^2$ *such that*

$$(\bar{\boldsymbol{X}}^n(0), \bar{\boldsymbol{Y}}^n(0)) := n^{-1}(\boldsymbol{X}^n(0), \boldsymbol{Y}^n(0)) \Rightarrow (\bar{\boldsymbol{X}}(0), \bar{\boldsymbol{Y}}(0)) \quad in \quad \mathbb{R}^4 \quad as \quad n \to \infty,$$

where $\bar{\boldsymbol{X}}(0) := (N_1, N_2)$ *and* $\bar{\boldsymbol{Y}}(0) := (\bar{Y}_1(0), \bar{Y}_2(0))$. *There exist random vectors* $\hat{\boldsymbol{X}}(0) := (\hat{X}_1(0), \hat{X}_2(0)) \in \mathbb{R}^2$ *and* $\hat{\boldsymbol{Y}}(0) := (\hat{Y}_1(0), \hat{Y}_2(0)) \in \mathbb{R}^2$ *such that*

$$(\hat{\boldsymbol{X}}^n(0), \hat{\boldsymbol{Y}}^n(0)) := \sqrt{n}(\bar{\boldsymbol{X}}^n(0) - \bar{\boldsymbol{X}}(0), \bar{\boldsymbol{Y}}^n(0) - \bar{\boldsymbol{Y}}(0)) \Rightarrow (\hat{\boldsymbol{X}}(0), \hat{\boldsymbol{Y}}(0)) \quad in \quad \mathbb{R}^4 \quad as \quad n \to \infty.$$

This assumption implies that the associated fluid-scaled initial quantities

$$(\bar{J}^n(0), \bar{\boldsymbol{Z}}^n(0), \bar{I}^n(0), \bar{\boldsymbol{R}}^n(0)) := n^{-1}(J^n(0), \boldsymbol{Z}^n(0), I^n(0), \boldsymbol{R}^n(0)) \Rightarrow (\bar{J}(0), \bar{\boldsymbol{Z}}(0), \bar{I}(0), \bar{\boldsymbol{R}}(0))$$

in $\mathbb{R}^6$ as $n \to \infty$, where

$$\bar{J}(0) := N_1 - \bar{Y}_2(0) = N_2 - \bar{Y}_1(0),\ \bar{\boldsymbol{Z}}(0) := (\bar{Z}_1(0), \bar{Z}_2(0)) := (0,0),\ \bar{I}(0) := 0,\ \bar{\boldsymbol{R}}(0) := (0,0).$$

Define the associated diffusion-scaled quantities $(\hat{J}^n(0), \hat{\boldsymbol{Z}}^n(0), \hat{I}^n(0), \hat{\boldsymbol{R}}^n(0))$ by

$$\hat{J}^n(0) := \frac{J^n(0) - n\bar{J}(0)}{\sqrt{n}}, \quad \hat{Z}_k^n(0) := \frac{Z_k^n(0)}{\sqrt{n}}, \quad \hat{I}^n(0) := \frac{I^n(0)}{\sqrt{n}}, \quad \hat{R}_k^n(0) := \frac{R_k^n(0)}{\sqrt{n}},\ k = 1, 2.$$

Then Assumption 6 implies that

$$\left(\hat{J}^n(0), \hat{\boldsymbol{Z}}^n(0), \hat{I}^n(0), \hat{\boldsymbol{R}}^n(0)\right) \Rightarrow \left(\hat{J}(0), \hat{\boldsymbol{Z}}(0), \hat{I}(0), \hat{\boldsymbol{R}}(0)\right) \quad \text{in} \quad \mathbb{R}^6 \quad \text{as} \quad n \to \infty,$$

where

$$\hat{J}(0) := \min_{k=1,2}\{-(\hat{X}_k(0))^- - \hat{Y}_{k'}(0)\}, \quad \hat{Z}_k(0) := -(\hat{X}_k(0))^- - \hat{Y}_{k'}(0) - \hat{J}(0),\ k = 1, 2,$$

$$\hat{I}(0) := \min_{k=1,2}(\hat{X}_k(0))^+, \quad \hat{R}_k(0) := (\hat{X}_k(0))^+ - \hat{I}(0),\ k = 1, 2.$$

Let

$$F_{k,e}(t) := \frac{1}{E[\eta_k^1]}\int_0^t F_k^c(s)ds, \quad t \geq 0,$$

be the equilibrium distribution associated with F_k, $k = 1, 2$.

Assumption 7. *For* $k = 1, 2$, $\{\tilde{\eta}_k^i : i \in \mathbb{N}\}$ *is a sequence of i.i.d. random variables with distribution* $F_{k,e}$ *and for each* $i \in \mathbb{N}$, $\tilde{\eta}_1^i$ *and* $\tilde{\eta}_2^i$ *are independent.* $\{\eta_k^{i,Q} : i \in \mathbb{N}\}$ *is a sequence of i.i.d. random variables with distribution* F_k *for each* $i \in \mathbb{N}$ *and* $k = 1, 2$. $\{(\eta_1^{i,I}, \eta_2^{i,I}) : i \in \mathbb{N}\}$ *is a sequence of i.i.d. random vectors with a joint distribution* $F(\cdot,\cdot)$. $\{(\eta_k^{i,R}, \tilde{\eta}_{k'}^{i,Z}) : i \in \mathbb{N}\}$ *is a sequence of i.i.d. random vectors with independent components,* $k = 1, 2$.

Finally, we also make an assumption for the residual waiting times $\{\tilde{w}_k^{n,i} : i = 1, ..., Q_k^n(0)\}$, $k = 1, 2$.

Assumption 8. *The residual waiting times of the tasks in queue* $\{\tilde{w}_k^{n,i} : i = 1, ..., Q_k^n(0)\}$, $k = 1, 2$, *converge to zero a.s. as* $n \to \infty$.

We define the diffusion-scaled processes $\hat{\boldsymbol{X}}^n := (\hat{X}_1^n, \hat{X}_2^n)$, $\hat{\boldsymbol{Y}}^n := (\hat{Y}_1^n, \hat{Y}_2^n)$ and $\hat{S}^n$ by

$$\hat{X}_k^n(t) := \frac{X_k^n(t) - N_k^n}{\sqrt{n}}, \quad \hat{Y}_k^n(t) := \frac{Y_k^n(t) - \tilde{Y}_k^n(t)}{\sqrt{n}}, \quad \hat{S}^n(t) := \frac{S^n(t) - \tilde{S}^n(t)}{\sqrt{n}}, \quad t \geq 0, \tag{3.25}$$

for $k = 1, 2$, where

$$\tilde{S}^n(t) := n\bar{S}^0(t) + \bar{\lambda}^n \int_0^t \int_0^t ((t - s_1) \wedge (t - s_2))\, dF(s_1, s_2), \tag{3.26}$$

$$\bar{S}^0(t) := \bar{Y}_2(0) F_{1,e}(t) + \bar{Y}_1(0) F_{2,e}(t) + \bar{J}(0) F_{1,e}(t) F_{2,e}(t), \tag{3.27}$$

$$\tilde{Y}_k^n(t) := n\bar{Y}_k(0) + \lambda^n t - \tilde{S}^n(t). \tag{3.28}$$

From the balance equation for Y_k^n in (3.24), we can rewrite $\hat{Y}_k^n$ as

$$\hat{Y}_k^n(t) = \hat{Y}_k^n(0) + \hat{X}_k^n(0) + \hat{A}^n(t) - \hat{X}_k^n(t) - \hat{S}^n(t), \quad t \geq 0, \quad k = 1, 2. \tag{3.29}$$

Recall $E_k^n(t)$ is defined as the cumulative number of tasks entering service by time $t \geq 0$ at station k, $k = 1, 2$, assuming the system starts empty in §3.2. Without abuse of notation, in §3.3 related to the FCLT, we let $E_k^n(t)$ be the number of *new arrivals* after time 0 whose task k has entered service by time $t \geq 0$ at station k, $k = 1, 2$.

Define the diffusion-scaled processes $(\hat{\boldsymbol{E}}^n, \hat{\boldsymbol{Q}}^n, \hat{\boldsymbol{B}}^n, \hat{\boldsymbol{D}}^n)$, $\hat{\boldsymbol{E}}^n := (\hat{E}_1^n, \hat{E}_2^n)$, $\hat{\boldsymbol{Q}}^n := (\hat{Q}_1^n, \hat{Q}_2^n)$, $\hat{\boldsymbol{B}}^n := (\hat{B}_1^n, \hat{B}_2^n)$ and $\hat{\boldsymbol{D}}^n := (\hat{D}_1^n, \hat{D}_2^n)$, by

$$\begin{aligned} &\hat{E}_k^n(t) := \frac{E_k^n(t) - \lambda^n t}{\sqrt{n}}, \quad \hat{Q}_k^n(t) := (\hat{X}_k^n(t))^+, \ \hat{B}_k^n(t) := -(\hat{X}_k^n(t))^-, \\ &\hat{D}_k^n(t) := \hat{X}_k^n(0) + \hat{A}^n(t) - \hat{X}_k^n(t), \quad t \geq 0, \quad k = 1, 2. \end{aligned} \tag{3.30}$$

For $s_1, s_2 \geq 0$, let

$$\begin{aligned} \hat{\mathcal{E}}^n(s_1, s_2) \ &:= \ \frac{1}{\sqrt{n}} ((E_1^n(s_1) \wedge E_2^n(s_2)) - \lambda^n (s_1 \wedge s_2)) \\ &= \ (\hat{E}_1^n(s_1) + (\lambda^n/\sqrt{n})(s_1 - s_1 \wedge s_2)) \wedge (\hat{E}_2^n(s_2) + (\lambda^n/\sqrt{n})(s_2 - s_1 \wedge s_2)). \end{aligned} \tag{3.31}$$

Before we present the FCLT for the fork-join network with NES in the Halfin-Whitt regime, we provide some preliminaries for the limit processes. The limit processes will be functionals of a generalized multiparameter Kiefer process, as a limit of the multiparameter sequential empirical process driven by the service time vectors of new arrivals. Define the multiparameter sequential empirical processes $\hat{K}^n := \{\hat{K}^n(t_1, t_2, \boldsymbol{x}) : t_1 \geq 0, t_2 \geq 0, \boldsymbol{x} \in \mathbb{R}_+^2\}$ by

$$\hat{K}^n(t_1, t_2, \boldsymbol{x}) := \frac{1}{\sqrt{n}} \sum_{i=1}^{\lfloor nt_1 \rfloor \wedge \lfloor nt_2 \rfloor} (\mathbf{1}(\boldsymbol{\eta}^i \leq \boldsymbol{x}) - F(\boldsymbol{x})). \tag{3.32}$$

We prove the convergence of $\hat{K}^n$ in the space $\mathbb{D}([0,\infty)^2, \mathbb{D}([0,\infty)^2, \mathbb{R}))$ endowed with a generalized Skorohod J_1 topology defined in [18] in Proposition 3.1.

Proposition 3.1. *Under Assumption 1,*

$$\hat{K}^n(t_1, t_2, \boldsymbol{x}) \Rightarrow \hat{K}(t_1, t_2, \boldsymbol{x}) \quad in \quad \mathbb{D}([0,\infty)^2, \mathbb{D}([0,\infty)^2, \mathbb{R})) \quad as \quad n \to \infty, \tag{3.33}$$

where $\hat{K}(t_1, t_2, \boldsymbol{x})$ is a continuous Gaussian random field, called a generalized multiparameter Kiefer process, with mean $E[\hat{K}(t_1, t_2, \boldsymbol{x})] = 0$ and covariance function

$$Cov(\hat{K}(s_1, s_2, \boldsymbol{x}), \hat{K}(t_1, t_2, \boldsymbol{y})) = (s_1 \wedge s_2 \wedge t_1 \wedge t_2)(F(\boldsymbol{x} \wedge \boldsymbol{y}) - F(\boldsymbol{x})F(\boldsymbol{y})), \tag{3.34}$$

for $s_k, t_k \geq 0$, $k = 1, 2$, and $\boldsymbol{x}, \boldsymbol{y} \in \mathbb{R}_+^2$.

We define the processes $\hat{W}_k := \{\hat{W}_k(t) : t \geq 0\}$, $\hat{W}_k^c := \{\hat{W}_k^c(t) : t \geq 0\}$ and $\hat{W} := \{\hat{W}(t) : t \geq 0\}$ as integral functionals of $\hat{K}$: for $t \geq 0$, $k = 1, 2$,

$$\hat{W}_k(t) := \int_0^t \int_0^t \int_{\mathbb{R}_+^2} \mathbf{1}(s_k + x_k \leq t) d\hat{K}(\lambda s_1, \lambda s_2, \boldsymbol{x}), \tag{3.35}$$

$$\hat{W}(t) := \int_0^t \int_0^t \int_{\mathbb{R}_+^2} \mathbf{1}(s_j + x_j \leq t, \forall j) d\hat{K}(\lambda s_1, \lambda s_2, \boldsymbol{x}), \tag{3.36}$$

and

$$\hat{W}_k^c(t) := \hat{W}_k(t) - \hat{W}(t) = \int_0^t \int_0^t \int_{\mathbb{R}_+^2} \mathbf{1}(s_k + x_k \leq t, s_{k'} + x_{k'} > t) d\hat{K}(\lambda s_1, \lambda s_2, \boldsymbol{x}), \tag{3.37}$$

where the integrals are defined in the sense of mean-square limits (see the precise definition in §**??**).

Proposition 3.2. *The processes $\hat{W}_k$, $\hat{W}_k^c$ and $\hat{W}$ are well-defined continuous Gaussian processes with mean zero, and for $0 \leq s < t$ and $k = 1, 2$,*

$$E[(\hat{W}_k(t) - \hat{W}_k(s))^2] = \lambda \int_0^t (F_k(t-u) - F_k(s-u))(1 - F_k(t-u) + F_k(s-u)) du,$$

$$\begin{aligned} E[(\hat{W}(t) - \hat{W}(s))^2] = \lambda \int_0^t \int_0^t &[\Delta F((s-s_1, s-s_2); (t-s_1, t-s_2))] \\ &\times [1 - \Delta F((s-s_1, s-s_2); (t-s_1, t-s_2))] d(s_1 \wedge s_2), \end{aligned} \tag{3.38}$$

$$\begin{aligned} E[(\hat{W}_k^c(t) - \hat{W}_k^c(s))^2] &= E[(\hat{W}_k(t) - \hat{W}_k(s))^2] + E[(\hat{W}(t) - \hat{W}(s))^2] \\ &\quad - 2\lambda \int_0^t \int_0^t [F(t-s_1, t-s_2) - F_{k,k'}(s-s_k, t-s_{k'}) \\ &\quad + (F_k(t-s_k) - F_k(s-s_k))(F(s-s_1, s-s_2) - F(t-s_1, t-s_2))] d(s_1 \wedge s_2), \end{aligned}$$

and covariance functions

$$Cov(\hat{W}_k(t), \hat{W}_{k'}(t)) = \lambda \int_0^t \int_0^t [F(t-s_1, t-s_2) - F_k(t-s_k)F_{k'}(t-s_{k'})] d(s_1 \wedge s_2),$$

$$Cov(\hat{W}_k(t), \hat{W}^c_{k'}(t)) = \lambda \int_0^t \int_0^t [F_k(t-s_k)F(t-s_1, t-s_2) - F_k(t-s_k)F_{k'}(t-s_{k'})] d(s_1 \wedge s_2),$$

$$Cov(\hat{W}_k(t), \hat{W}(t)) = \lambda \int_0^t \int_0^t [F(t-s_1, t-s_2) - F_k(t-s_k)F(t-s_1, t-s_2)] d(s_1 \wedge s_2),$$

$$Cov(\hat{W}^c_k(t), \hat{W}(t)) = \lambda \int_0^t \int_0^t [(F(t-s_1, t-s_2))^2 - F_k(t-s_k)F(t-s_1, t-s_2)] d(s_1 \wedge s_2),$$

where $F_{k,k'}(x,y) := P(\eta^i_k \le x, \eta^i_{k'} \le y)$ *for* $x, y \in \mathbb{R}_+$, *and*

$$\Delta F(\boldsymbol{x};\boldsymbol{y}) := F(y_1,y_2) - F(x_1,y_2) - F(y_1,x_2) + F(x_1,x_2), \quad \boldsymbol{x},\boldsymbol{y} \in \mathbb{R}^2_+,\ \boldsymbol{x} \le \boldsymbol{y}.$$

In addition, let $\hat{U} := \{\hat{U}(\boldsymbol{t}) : \boldsymbol{t} \in \mathbb{R}^2_+\}$ be a continuous two-parameter Gaussian process with mean zero and covariance function:

$$Cov(\hat{U}(\boldsymbol{s}), \hat{U}(\boldsymbol{t})) = (F_{1,e}(s_1 \wedge t_1)F_{2,e}(s_2 \wedge t_2) - F_{1,e}(s_1)F_{2,e}(s_2)F_{1,e}(t_1)F_{2,e}(t_2)), \tag{3.39}$$

for $\boldsymbol{s} := (s_1, s_2) \in \mathbb{R}^2_+$ and $\boldsymbol{t} := (t_1, t_2) \in \mathbb{R}^2_+$. Define $\hat{U}_k := \{\hat{U}_k(t) : t \ge 0\}$, for $k = 1, 2$, by

$$\hat{U}_1(t) := \hat{U}(t, \infty), \quad \hat{U}_2(t) := \hat{U}(\infty, t), \quad t \ge 0, \tag{3.40}$$

and without abuse of notation, we denote $\hat{U}(t) = \hat{U}(t,t)$, $t \ge 0$. Note that the processes $\hat{W}_k$, $\hat{W}^c_k$ and $\hat{W}$ are independent with $\hat{U}$, as well as $\hat{U}_k$, $k = 1, 2$.

We are now ready to state the FCLT.

Theorem 3.2. *Under Assumptions 1 and 4-8,*

$$\left(\hat{A}^n, \hat{\boldsymbol{X}}^n, \hat{\boldsymbol{Y}}^n, \hat{S}^n, \hat{\boldsymbol{E}}^n, \hat{\boldsymbol{Q}}^n, \hat{\boldsymbol{B}}^n, \hat{\boldsymbol{D}}^n\right) \Rightarrow \left(\hat{A}, \hat{\boldsymbol{X}}, \hat{\boldsymbol{Y}}, \hat{S}, \hat{\boldsymbol{E}}, \hat{\boldsymbol{Q}}, \hat{\boldsymbol{B}}, \hat{\boldsymbol{D}}\right) \tag{3.41}$$

in $\mathbb{D}^{14}$ *as* $n \to \infty$, *where* $\hat{A}$ *is in* (3.18), $\hat{\boldsymbol{X}}$, $\hat{\boldsymbol{Y}}$ *and* $\hat{S}$ *are the unique solutions to the following set of stochastic integral equations: for* $t \ge 0$ *and* $k = 1, 2$,

$$\begin{aligned}
\hat{X}_k(t) &= \hat{X}^0_k(t) - N_k\beta_k F_{k,e}(t) - \bar{J}(0)^{1/2}\hat{U}_k(t) - \bar{Y}_{k'}(0)^{1/2}\hat{B}_{0,k}(F_{k,e}(t)) \\
&\quad + \int_0^t (\hat{X}_k(t-s))^+ dF_k(s) + \int_0^t F^c_k(t-s) d\hat{A}(s) - \hat{W}_k(t),
\end{aligned} \tag{3.42}$$

$$\begin{aligned}
\hat{Y}_k(t) &= \hat{Y}^0_k(t) + N_k\beta_k F_{k,e}(t) - \bar{Y}_k(0)^{1/2}\hat{B}_{0,k'}(F_{k',e}(t)) + \bar{J}(0)^{1/2}(\hat{U}_k(t) - \hat{U}(t)) \\
&\quad - \int_0^t (\hat{X}_k(t-s))^+ dF_k(s) + \int_0^t F_k(t-s) d\hat{A}(s) + \hat{W}^c_k(t) - \hat{\Psi}(t),
\end{aligned} \tag{3.43}$$

$$\hat{S}(t) = \hat{S}^0(t) + \bar{Y}_2(0)^{1/2}\hat{B}_{0,1}(F_{1,e}(t)) + \bar{Y}_1(0)^{1/2}\hat{B}_{0,2}(F_{2,e}(t)) + \bar{J}(0)^{1/2}\hat{U}(t) + \hat{W}(t) + \hat{\Psi}(t), \tag{3.44}$$

and $\hat{\boldsymbol{E}}^n$, $\hat{\boldsymbol{Q}}^n$, $\hat{\boldsymbol{B}}^n$ *and* $\hat{\boldsymbol{D}}^n$ *are given as follows:*

$$\hat{E}_k(t) = \hat{A}(t) - (\hat{X}_k(t))^+, \quad \hat{D}_k(t) = \hat{X}_k(0) + \hat{A}(t) - \hat{X}_k(t), \tag{3.45}$$

$$\hat{Q}_k(t) = (\hat{X}_k(t))^+, \quad \hat{B}_k(t) = -(\hat{X}_k(t))^-,$$

where

$$\hat{X}_k^0(t) := \hat{X}_k(0)F_{k,e}^c(t) + (\hat{X}_k(0))^+(F_k^c(t) - F_{k,e}^c(t)), \tag{3.46}$$

$$\hat{S}^0(t) := \sum_{k=1}^{2}(\hat{Y}_{k'}(0)F_{k,e}(t) + \hat{Z}_{k'}(0)F_k(t)F_{k',e}(t)) + \hat{J}(0)F_{1,e}(t)F_{2,e}(t) + \hat{I}(0)F_m(t), \tag{3.47}$$

$$\hat{Y}_k^0(t) := \hat{Y}_k(0) + \hat{X}_k(0)F_{k,e}(t) + (\hat{X}_k(0))^+(F_k(t) - F_{k,e}(t)) - \hat{S}^0(t), \tag{3.48}$$

the processes $\hat{B}_{0,k} := \{\hat{B}_{0,k}(t) : t \geq 0\}$, $k = 1, 2$, are independent standard Brownian bridges, the process $\hat{U}$ is a continuous two-parameter Gaussian process defined above with the processes $\hat{U}_1$ and $\hat{U}_2$ defined in (3.40)*, and the processes $\hat{W}_k$, $\hat{W}_k^c$ and $\hat{W}$ are defined in* (3.35)*,* (3.37) *and* (3.36)*, and $\hat{B}_{0,k}$ is independent of $\hat{U}$ and $\hat{W}_k$, $\hat{W}_k^c$ and $\hat{W}$, and the process $\hat{\Psi} := \{\hat{\Psi}(t) : t \geq 0\}$ defined by*

$$\hat{\Psi}(t) := \int_0^t \int_0^t \hat{\mathcal{E}}(t - s_1, t - s_2) dF(s_1, s_2), \tag{3.49}$$

is a well-defined continuous process, where, for $s_1, s_2 \geq 0$,

$$\hat{\mathcal{E}}(s_1, s_2) := \hat{E}_1(s_1)\mathbf{1}(s_1 < s_2) + \hat{E}_2(s_2)\mathbf{1}(s_2 < s_1) + (\hat{E}_1(s_1) \wedge \hat{E}_2(s_2))\mathbf{1}(s_1 = s_2). \tag{3.50}$$

We remark that the limit processes $\hat{X}_k$, $k = 1, 2$, have the same structure as the unique solution to an integral convolution equation, as shown in Reed [45], but are also different because they are both driven by the same generalized multiparameter Kiefer process $\hat{K}$ defined in Proposition 3.1. These two limiting processes $\hat{X}_k$, $k = 1, 2$, are correlated because of the correlated service times of the parallel tasks of each job, which is captured by the process $\hat{K}$, as well as the same arrival limit process $\hat{A}$. In fact, these two processes $\hat{K}$ and $\hat{A}$ as well as the limits associated with the initial quantities are the driving stochastic components of all the limit processes in (3.42)-(3.45).

4 Concluding Remarks and Future Work

We remark on the main ideas of the proofs for the limit theorems due to space constraint. The main difficulty in the study of many-server fork-join networks with NES is the resequencing of arrival orders after service completion at each service station. Tasks of distinct jobs must be differentiated and tracked in order to describe the waiting buffer dynamics for synchronization. To mathematically describe the system dynamics, we develop a new approach using multiparameter sequential empirical processes driven by service vectors for parallel tasks of each job, as depicted in Figure 2. This approach is used to establish FLLNs and FCLTs for the waiting buffer processes for synchronization and the service processes jointly in the fundamental fork-join network where all service stations are operating in the many-server heavy-traffic regimes.

As a prerequisite, we first establish a new FCLT for multiparameter sequential empirical processes driven by random vectors (Theorem 2.1). To prove Theorem 2.1, we employ the standard approach of establishing convergence of finite-dimensional distributions and tightness [**?**, 20, 55]. The convergence of finite-dimensional distributions follows from the strong convergence result of multiparameter empirical processes in [42]. To prove tightness, we present a new decomposition property for multiparameter sequential empirical processes, which have a multiparameter martingale [23, 19], and a second term of finite variation. We apply properties of multiparameter martingales [23, 19] and strong approximations of random walks by Brownian motions (see section 3.5 in [30]) to show the tightness of those two decomposed terms, respectively. This decomposition also plays a very important role in proving the tightness of the number of tasks in each waiting buffer for synchronization, the number of tasks in each parallel service station and the number of synchronized jobs. Specifically, the aforementioned processes can be decomposed into a linear combination of three terms: an integral functional of the arrival process and two other terms from the decomposition of the multiparameter sequential empirical process. We apply Aldous' tightness criteria (see, e.g., Lemma 3.7 in [28]) and another tightness criteria for processes with proper decompositions satisfying certain conditions (Lemma VI.3.32 in [20]) to verify the tightness property of the two terms related to the sequential empirical process driven by the service vector, respectively.

The proofs of the limit theorems in the QD regime can be regarded as generalizations of those for $G/GI/\infty$ queues in [28]. However, since all the processes, $\boldsymbol{X}$, $\boldsymbol{Y}$ and S, are represented via the multiparameter sequential empirical processes driven by the service vectors, many technical challenges must be addressed in the multiparameter setting, for example, using multiparameter L_2 martingales, and mean-square limits of (integral functionals of) multiparameter processes defined on $\mathbb{R}^k$ ($k \geq 2$). One important advantage of our new approach is that all the diffusion-scale limit processes for $\boldsymbol{X}$, $\boldsymbol{Y}$ and S are all functionals of two independent processes - the arrival limit and the multiparameter generalized Kiefer process driven by the service vector (Theorem 2.3). From that, the characterization of the joint transient and stationary distributions of these processes is made possible (Theorem 2.4).

The proofs in the QED regime are based on the important observations that the system dynamics of $G/GI/n$ queues can be represented via the corresponding $G/GI/\infty$ service dynamics [45], and that waiting times in the QED regime are $O(1/\sqrt{n})$ while service times are $O(1)$. For the fork-join network, we represent the dynamics of $\boldsymbol{X}$, $\boldsymbol{Y}$ and S via that in the corresponding infinite-server fork-join network where the entering service times in the model are regarded as the "arrival" times for the corresponding infinite-server fork-join network, as shown in Figure 2(b). The observation that the entering service times in the parallel stations have a difference of order $O(1/\sqrt{n})$ is key to prove the joint convergence of the aforementioned processes. On the other hand, since we have to simultaneously handle the waiting times of all parallel tasks and work with multiparameter sequential empirical processes, we must develop new techniques to prove tightness, including establishing new properties for multiparameter L_2 martingales, and identifying a new multivariate integral mapping to apply the continuous mapping theorem.

We believe that a general framework has been developed to study fork-join networks with NES in the many-server heavy-traffic regimes (QD and QED). It can be potentially used to study performance evaluation, capacity allocation, and control problems in multi-class fork-join networks under NES with multi-stage processing. We want to find optimal scheduling and routing policies such that delays for synchronization as well as delays for service can be minimized, particularly, reducing delays for synchronization to be of a smaller order than service. We also want to find optimal staffing policies to stabilize delays for synchronization in addition to delays for service when arrival rates are time inhomogeneous. Our methods can be extended to investigate reliability of many-server fork-join networks under NES in random environments (e.g., service disruptions). Fork-join networks with NES are more likely to suffer from service disruptions due to the structural complexity of parallel and sequential task processing. Component-level unreliability can be much more amplified by its large scale. We will extend our approach to investigate the impact of service disruptions in one or multiple service stations upon system congestion, particularly, delays for synchronization and throughput.

References

[1] M. Armony, S. Israelit, A. Mandelbaum, Y. Marmor, Y. Tseytlin and G. Yom-Tov. (2011) Patient flow in hospitals: a data-based queueing-science perspective. *Submitted to Stochastic Systems,* 20.

[2] R. Atar, A. Mandelbaum and A. Zviran. (2012) Control of fork-join networks in heavy traffic. *Communication, Control, and Computing (Allerton), 2012 50th Annual Allerton Conference.*

[3] F. Baccelli, S. Foss and M. Lelarge. (2004) Asymptotics of sub exponential max-plus networks: the stochastic event graph case. *Queueing Systems*, 46(1-2), 75–96.

[4] F. Baccelli, A. M. Makowski and A. Shwartz. (1989) The fork-join queue and related systems with synchronization constraints: stochastic ordering and computable bounds. *Advances in Applied Probability*, 629–660.

[5] P. Billingsley. (2009) *Convergence of Probability Measures.* John Wiley & Sons.

[6] S. Bouzebda and T. Zari. (2008) Strong approximation of empirical copula process by Kiefer process. *arXiv preprint arXiv:* 0811.3330.

[7] P. Brémaud. (1981) *Point Processes and Queues: Martingale Dynamics,* New York: Springer.

[8] M. Csörgö and L. Horváth. (1988) A note on strong approximations of multivariate empirical processes. *Stochastic Processes and Their Applications*, 28(1), 101–109.

[9] M. Csörgö and P. Révész. (1975) A strong approximation of the multivariate empirical process. *Studia Sci. Math. Hungar,* 10, 427-434.

[10] H. Dai. (2011) Exact Monte Carlo simulation for fork-join networks. *Advances in Applied Probability,* 43(2), 484-503.

[11] J. Dean and S. Ghemawat. (2008) MapReduce: simplified data processing on large clusters. *Communications of the ACM.* Vol. 51, No. 1, 107–113.

[12] A. B. Dieker and M. Lelarge. (2006) Tails for (max, plus) recursions under subexponentiality. *Queueing Systems,* 53(4), 213-230.

[13] P. Doukhan, J. D. Fermanian and G. Lang. (2009) An empirical central limit theorem with applications to copulas under weak dependence. *Statistical Inference for Stochastic Processes,* 12(1), 65–87.

[14] S. G. Eick, W. A. Massey and W. Whitt. (1993) The physics of the M_t/G/∞ queue. *Operations Research,* 41(4), 731–742.

[15] S. N. Ethier and T. G. Kurtz. (1986) *Markov Processes: Characterization and Convergence.* John Wiley & Sons.

[16] J. Gallien and L. M. Wein. (2001) A simple and effective component procurement policy for stochastic assembly systems. *Queueing systems,* 38(2), 221-248.

[17] I. Gurvich and A. Ward. (2014) On the dynamic control of matching queues. Forthcoming in *Stochastic Systems.*

[18] B. G. Ivanoff. (1980) The function space D $([0,\infty)^q$, E). *Canadian Journal of Statistics.* Vol. 8, No. 2, 179–191.

[19] B. G. Ivanoff and E. Merzbach. (1999) *Set-Indexed Martingales.* CRC Press.

[20] J. Jacod and A. N. Shiryaev. (1987) *Limit Theorems for Stochastic Processes.* Berlin: Springer-Verlag.

[21] A. Jean-Marie and L. Gün. (1993) Parallel queues with resequencing. *Journal of the ACM*, 40(5), 1188-1208.

[22] L. Jiang and R. E. Giachetti. (2008) A queueing network model to analyze the impact of parallelization of care on patient cycle time. *Health Care Management Science*, 11(3), 248–261.

[23] D. Khoshnevisan. (2002) *Multiparameter Processes: An Introduction to Random Fields.* Springer.

[24] L. J. Klementowski. (1978) *PERT/CPM and Supplementary Analytical Techniques. An Analysis of Aerospace Usage.* Ph.D. Thesis.

[25] S. S. Ko and R. F. Serfozo. (2004) Response times in M/M/s fork-join networks. *Advances in Applied Probability,* 854-871.

[26] S. S. Ko and R. F. Serfozo. (2008) Sojourn times in G/M/1 fork-join networks. *Naval Research Logistics*, 55(5), 432-443.

[27] P. Konstantopoulos and J. Walrand. (1989) Stationary and stability of fork-join networks. *Journal of Applied Probability*, 604-614.

[28] E. V. Krichagina and A. A. Phalskii. (1997) A heavy-traffic analysis of a closed queueing system with a GI/∞ service center. *Queueing Systems,* 25(1-4), 235-280.

[29] R. C. Larson, M. F. Cahn and M. C. Shell. (1993) Improving the New York City arrest-to-arraignment system. *Interfaces*, 23(1), 76-96.

[30] G. F. Lawler and V. Limic. (2010) *Random Walk: A Modern Introduction.* Cambridge University Press.

[31] M. R. Leadbetter, G. Lindgren and H. Rootzén. (1983) *Extremes and Related Properties of Random Sequences and Processes.* Springer-Verlag.

[32] M. Lin, L. Zhang, A. Wierman and J. Tan. (2013) Joint optimization of overlapping phases in MapReduce. *Performance Evaluation,* 70(10), 720-735.

[33] A. Mandelbaum and P. Momcilovic. (2012) Queues with many servers and impatient customers. *Mathematics of Operations Research,* 37(1), 41-65.

[34] A. W. Marshall and I. Olkin. (1967) A multivariate exponential distribution. *Journal of the American Statistical Association.* 62.317: 30-44.

[35] G. Neuhaus. (1971) On weak convergence of stochastic processes with multidimensional time parameter. *The Annals of Mathematical Statistics*, 42(4) 1285-1295.

[36] V. Nguyen. (1993) Processing networks with parallel and sequential tasks: heavy traffic analysis and Brownian Limits. *The Annals of Applied Probability*, 3(1), 28-55.

[37] V. Nguyen. (1994) The trouble with diversity: fork-join networks with heterogeneous customer population. *The Annals of Applied Probability*, 4(1), 1-25.

[38] G. Pang, R. Talreja and W. Whitt. (2007) Martingale proofs of many-server heavy-traffic limits for Markovian queues. *Probability Surveys.* 4, 193–267.

[39] G. Pang and W. Whitt. (2010) Two-parameter heavy-traffic limits for infinite-server queues. *Queueing Systems*, 65(4), 325-364.

[40] G. Pang and W. Whitt. (2012) Infinite-server queues with batch arrivals and dependent service times. *Probability in Engineering and Informational Sciences.* 26.02: 197-220.

[41] G. Pang and W. Whitt. (2013) Two-parameter heavy-traffic limits for infinite-server queues with dependent service times. *Queueing Systems*, 73(2), 119-146.

[42] W. Philipp and L. Pinzur. (1980) Almost sure approximation theorems for the multivariate empirical process. *Zeitschrift für Wahrscheinlichkeitstheorie und verwandte Gebiete,* 54(1), 1-13.

[43] B. Prabhakar, N. Bambos and T. S. Mountford. (2000) The synchronization of Poisson processes and queueing networks with service and synchronization nodes. *Advances in Applied Probability*, 824–843.

[44] A. A. Puhalskii and J. E. Reed. (2010) On many-server queues in heavy traffic. *The Annals of Applied Probability*, 20(1), 129-195.

[45] J. E. Reed. (2009) The $G/GI/N$ queue in the Halfin-Whitt regime. *The Annals of Applied Probability*, 19(6), 2211-2269.

[46] M. Sklar. (1959) Fonctions de répartition à n dimensions et leurs marges. *Université Paris* 8.

[47] M. S. Squillante, Y. Zhang, A. Sivasubramaniam and N. Gautam. (2008) Generalized parallel-server fork-join queues with dynamic task scheduling. *Annals of Operations Research*, 160(1), 227–255.

[48] M. L. Straf. (1972) Weak convergence of stochastic processes with several parameters. *Proc. Sixth Berkeley Symp. Math. Statist. Probab.*, 2, 187-221.

[49] M. Takahashi, H. Ōsawa and T. Fujisawa. (2000) On a synchronization queue with two finite buffers. *Queueing Systems*, 36(1-3), 107-123.

[50] R. Talreja and W. Whitt. (2009) Heavy-traffic limits for waiting times in many-server queues with abandonment. *The Annals of Applied Probability*, 2137-2175.

[51] J. Tan, X. Meng and L. Zhang. (2012) Delay tails in MapReduce scheduling. *ACM SIGMETRICS Performance Evaluation Review*, 40(1), 5-16.

[52] X. Tan and C. Knessl. (1996) A fork-join queueing model: diffusion approximation, integral representations and asymptotics. *Queueing Systems*, 22(3-4), 287-322.

[53] S. Varma. (1990) *Heavy and Light Traffic Approximations for Queues with Synchronization Constraints.* Ph.D. Thesis.

[54] W. Wang, K. Zhu, L. Ying, J. Tan and L. Zhang. (2013) Map task scheduling in MapReduce with data locality: throughput and heavy-traffic optimality. *Proceedings of IEEE INFOCOM.* 1609–1617.

[55] W. Whitt. (2002) *Stochastic-Process Limits: An Introduction to Stochastic-Process Limits and Their Application to Queues.* Springer.

[56] C. J. Willits and D. C. Dietz. (2001) Nested fork-join queueing network model for analysis of airfield operations. *Journal of Aircraft*, 38(5), 848-855.

[57] I. Zaied. (2010) *The Offered Load in Fork-Join Networks: Calculations and Applications to Service Engineering of Emergency Department.* M. Sc. Research Proposal.

[58] A. Zviran. (2008) *Fork-Join Networks in Heavy Traffic: Diffusion Approximations and Control.* M. Sc. Research Proposal, Technion.