Gaussian Approximations for A Fork-Join Network with Non-Exchangeable Synchronization

Hongyuan Lu, Guodong Pang

1. Motivation and Model

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, patient flow analysis in healthcare, parallel computing, military deployment operations, and law enforcement systems. 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. 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. Here we focus on the second type of synchronization constraint.

We are primarily motivated to study fork-join networks with NES from patient flow analysis in hospitals. For example, as a prerequisite for a doctor examination, all the test results for the same patient must be ready, and they cannot be mixed among different patients. Mixing one patient's blood test result with another patient's cardiology result can lead to severe medical consequences. Fork-join networks with NES also have applications in parallel computing. For example, a computation job can be split into several tasks processed in parallel (possibly only at some stages) and joined subsequently once these parallel tasks tagged with the same job are completed. Another example is MapReduce scheduling. In the map phase, jobs are split into parallel tasks and these tasks tagged with the same job are synchronized to be processed in the reduce phase (with additional complicated interactions between those two phases).

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. 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 consider a fundamental fork-join network model with a single class of jobs and NES, where each 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 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(a) depicts such a model. Here we assume that the parallel service time vectors are iid with a continuous joint distribution. In this model, in addition to

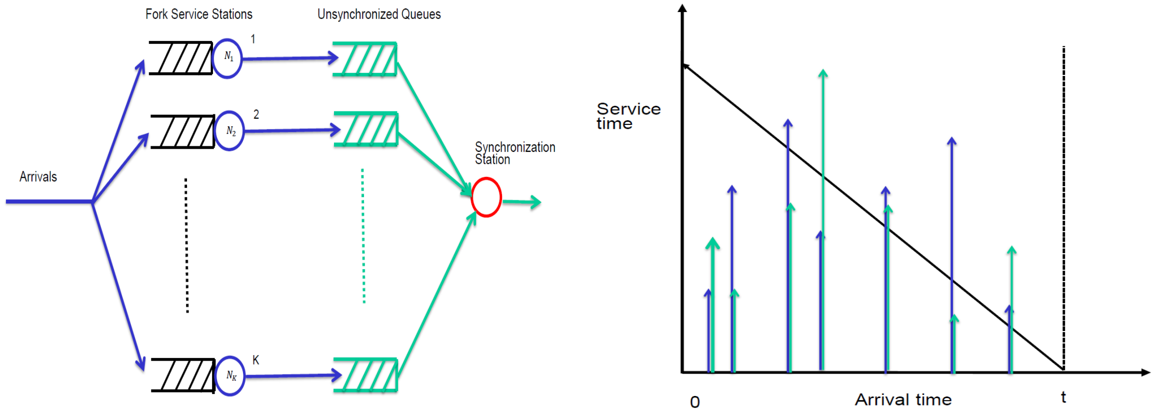

(a) (b)

Figure 1: A fundamental fork-join network and a graphical representation of its system dynamics

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. 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. In this work, we start with the situation when all the service stations operate in the quality-driven (QD) many-server regime. Asymptotically, this is equivalent to a model which has infinite numbers of servers at all service stations.

2. Methods and Results

To describe the system dynamics, we start with a graphical representation as shown in Figure 1(b) for a system of two parallel tasks. At each job's arrival epoch, we mark the arrival time on the x-axis and the service times of all parallel tasks on the 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 random measure theory, similarly as in the "physics" of $M/GI/\infty$ queues. 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. 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.

We develop a new approach to show a functional central limit theorem (FCLT) for 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, under general assumptions on the arrival and service processes. Specifically, we represent the aforementioned processes as functionals of a sequential empirical process driven by the sequence of service vectors for each job's parallel tasks. As a consequence, 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.

When the limiting arrival process is Brownian motion, we show that the limiting processes are a multidimensional Gaussian process, and thus characterize the joint transient and stationary distributions of these processes as multivariate Gaussian distribution. We study the impact of the correlation among the service vector on the system. It is found that, the positive correlation does not affect the mean and the variance of the number of tasks in parallel service stations, while it decreases the mean and the variance of the number of unsynchronized tasks in waiting buffers, the covariances of the number of unsynchronized tasks between any two waiting buffers and the covariances between the number tasks in parallel service stations and the number of unsynchronized tasks in waiting buffers. We also characterize the difference between the steady-state mean value of the waiting buffer for synchronization in our model with that in the ES system.

3. A Numerical Example

We provide a numerical example with two parallel tasks, comparing our approximations with simulations. We let the arrival process be renewal with an arrival rate 100 and the SCV 5. The service times of the 1st and 2nd parallel tasks have a bivariate Marshall-Olkin H_2 distribution with means 1 and 0.9, respectively. We consider both independent and correlated (with correlation coefficient 0.5) parallel service times. In Table 1, we show the approximation and simulation values for the mean, variance and covariance of the number of tasks in parallel service stations X and the size of waiting buffer for synchronization Y.

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