

# Approximations and Bounds for $(n, k)$ Fork-Join Queues: A Linear Transformation Approach

Huajin Wang<sup>\*†</sup>, Jianhui Li<sup>\*</sup>, Zhihong Shen<sup>\*</sup> and Yuanchun Zhou<sup>\*</sup>

<sup>\*</sup>Computer Network Information Center, Chinese Academy of Sciences, China

Email: {wanghj, lijh, bluejoe, zyc}@cnic.cn

<sup>†</sup>University of Chinese Academy of Sciences, China

**Abstract**— $(n, k)$  fork-join queues are prevalent in popular distributed systems, erasure coding based cloud storages, and modern network protocols like multipath routing, estimating the sojourn time of such queues is thus critical for the performance measurement and resource plan of computer clusters. However, the estimating keeps to be a well-known open challenge for years, and only rough bounds for a limited range of load factors have been given. This paper developed a closed-form linear transformation technique for jointly-identical random variables: An order statistic can be represented by a linear combination of maxima. This brand-new technique is then used to transform the sojourn time of non-purging  $(n, k)$  fork-join queues into a linear combination of the sojourn times of basic  $(k, k), (k + 1, k + 1), \dots, (n, n)$  fork-join queues. Consequently, existing approximations for basic fork-join queues can be bridged to the approximations for non-purging  $(n, k)$  fork-join queues. The uncovered approximations are then used to improve the upper bounds for purging  $(n, k)$  fork-join queues. Simulation experiments show that this linear transformation approach is practiced well for moderate  $n$  and relatively large  $k$ .

## I. INTRODUCTION

The performance of fork-join queues is a highly focused research topic for many years for the ubiquitousness of fork-join queues in both real-life workflows and computing systems. In a fork-join queueing system, a job is *forked* into  $n$  sub-tasks when it arrives at a control node, and each sub-task is sent to a single node to be conquered. Results of finished sub-tasks are summarized at a central *join* node. When the job arrival rate  $\lambda$  is high, a sub-task may have to wait for service in the sub-queue of its hosting node in a first-come-first-serving order. A basic fork-join queue considers a job is done after all results of the job have been received at the join node (see Fig. 1 (a)).

In Big Data era, more and more mainstream computing infrastructures become distributively deployed, and inevitably recruit fork-join queues to facilitate the storing and processing of large-scale datasets. For example: 1) Cassandra [1] and Dynamo [2], two popular key-value data stores, use fork-join queues to concurrently perform read and write operations on all the replicas of the target key-value pairs; 2) The client of an  $(n, k)$  MDS (maximum distance separable) erasure coding based cloud storage system only needs to retrieve any  $k$  out of all  $n$  blocks of a file to reconstruct the file; 3) The data transmission process of multipath routing protocols can

generally be simplified as a multi-stage fork-join queueing process.

Latency is commonly a critical concern in building and optimizing Big Data clusters. For example, in Amazon's cloud platform, services commonly have latency requirements which are in general measured at the 99.9<sup>th</sup> percentile of the distribution [2]. The Dynamo storage system must be capable of meeting such stringent SLAs. In this scenario, basic fork-join queues may cause serious performance issues when the number of data replicas are large, since they require all the sub-tasks of a job to be finished before making the job's response. By contrast,  $(n, k)$  fork-join queues, as named in [3], only require the job's any  $k$  out of  $n$  sub-tasks to be finished, and thus have performance advantages in such scenarios. For example, a write request in Cassandra can either be responded when a quorum of replicas have been successfully written, or just get responded once the fast answer from all touched replicas is acknowledged when there is a need to pursue high throughputs.

As depicted in Fig. 1, there are mainly two versions of  $(n, k)$  fork-join queues: The purging one removes all the remaining sub-tasks of a job from both sub-queues and service stations once it receives the job's  $k^{\text{th}}$  answer. The file retrieval process from an MDS coded cloud storage system is such an example. As a contrast, the non-purging one keeps queuing and executing remaining sub-tasks. For example, a write operation in Cassandra needs to update all the replicas of the target key-value pair, while can response to user as soon as a quorum of replicas have been successfully written.

a) *The State-of-the-Art Research on Basic Fork-Join Queues*: The popularity of fork-join systems has drawn great attentions from database/OS/networking communities to the performance analyses of fork-join queues for a rather long period of time. Unfortunately, there is still no exact closed-form solution of the sojourn time of the job in  $n \geq 3$  basic fork-join queues. The difficulty lies in the fact that the sojourn times of a job's sub-tasks are not independent, as their hosting sub-queues share the same sub-task arrival process. Since most of existing exact analysis techniques are developed for independent and identical (iid) random variables, it is very hard to trace the sojourn time distribution for fork-join queues.

For  $n \geq 3$  fork-join queues under Poisson job arrival process and with iid exponential service time distributions, Nelson

This work was supported by the National Key Research Program of China (2016YFB1000600, 2016YFB0501900). Corresponding author is Jianhui Li.




Fig. 1. Fork-Join Queues:  $A^{(1)}$  Is a Sub-Task of Job  $A$

et al. [4] proposed an initiative approximation technique which is based on the fact that the sojourn times  $X_{1..k}$  of sub-tasks  $1, 2, \dots, k$  are associated variables, whose maximum can be bounded by the maximum of their iid equivalents [5]:  $P(X_{1..n} \leq t) \geq \prod_{i=1}^n P(X_i^{IID} \leq t)$ . According to that, the upper bounds and closed-form approximations of the sojourn time were given in this work. Simulation experiments in [6] showed that Nelson's approximation is still the most reliable one, compared to following works such as [7] and [8].

*b) The State-of-the-Art Research and Open Challenges on  $(n, k)$  Fork-Join Queues:* Despite the popularity of  $(n, k)$  fork-join queues in Big Data systems and many other fields, there are even no practical approximations on the sojourn time of  $(n, k)$  fork-join queues: Unlike the maximum, the  $k^{th}$  order statistic cannot be bounded by using associated variables' property, which makes the sojourn time of  $(n, k)$  fork-join queues more hard to analyze, compared to basic fork-join queues.




Fig. 2. A  $(3, 2)$  Split-Merge Queue

Currently, there are only exact quantity analyses for purging  $(n, 1)$  fork-join queues [9], [10], because such a queue is equivalent to a single queue with  $n$  times the sub-queue's service rate. For general purging  $(n, k)$  fork-join queues, there are only rough bounds have been given: Joshi et al. [3], [11] resort to the split-merge queue model (see Fig. 2) to find proper upper and lower bounds. Compared to purging  $(n, k)$  fork-join queues, all empty sub-queues in the split-merge model are blocked and cannot serve subsequent tasks until  $k$  sub-tasks of the current job are completed, which makes the split-merge model much easier to trace. However, these split-merge based bounds tend to be extremely loose when increasing  $k$  or the load factor  $\rho$ , as we depict in Section IV.

Since non-purging  $(n, k)$  fork-join queues cannot be reduced to the split-merge model, they are more difficult to analyze, even including  $(n, 1)$  queues. Recently, Fidler et al. [12] gave non-asymptotic statistical bounds on the sojourn times for non-purging fork-join queues. However, no reasonable approximations have been proposed.

*c) Methodology and Contributions:* This paper aims at fixing the lack of proper approximations for non-purging  $(n, k)$  fork-join queues and tackling the uncontrollability of bounds

for purging  $(n, k)$  fork-join queues. To achieve these objectives, we trace fork-join queues in a fundamental way: The linear relationship between  $(n, k)$  fork-join queues and their basic  $(k, k), (k + 1, k + 1), \dots, (n, n)$  equivalents is depicted for the first time; This relationship is then used to bridge the existing approximations for basic fork-join queues to the approximations and bounds for  $(n, k)$  fork-join queues.

Our innovations and contributions are highlighted as follows:

- A brand-new closed-form linear transformation technique for jointly-identical random variables, by which order statistics can be transformed into a closed-form linear combination of maxima. Besides, there is no need to assume the independence of variables.
- The first reasonable and practical method to approximate the expected sojourn time of non-purging  $(n, k)$  fork-join queues with general service time distributions. This method relies on the cooperation between the linear transformation technique and the existing approximations for basic fork-join queues.
- Improvements over the upper bounds on the expected sojourn time of purging  $(n, k)$  fork-join queues, which are gained by resorting the bounds to that of the non-purging equivalent  $(n, k)$  fork-join queues.

This paper is organized as follows: The linear transformation technique is developed in Section II; This technique is then employed in Section III to find proper approximations for non-purging  $(n, k)$  fork-join queues; The flaws of existing bounds for purging  $(n, k)$  fork-join queues and our improvements over upper bounds are depicted in Section IV; In Section V, we discuss the limitation of this linear transformation technique; Related works are reviewed in Section VI; We conclude this work and point out some promising future research directions in Section VII.

## II. PRELIMINARIES: LINEAR TRANSFORMATIONS OF ORDER STATISTICS

In this section, we consider a family of rvs (random variables)  $X_1, X_2, \dots, X_n$  (denoted as  $X_{1..n}$ ) defined on a probability space, and let  $X_{(n,k)}$  denotes their  $k^{th}$  order statistic,  $P_k$  denotes the possibility  $P(X_1 \leq t, X_2 \leq t, \dots, X_k \leq t)$  and  $P_{n,k}$  denotes the possibility  $P(X_1 \leq t, X_2 \leq t, \dots, X_k \leq t, X_{k+1} > t, X_{k+2} > t, \dots, X_n > t)$ . Obviously,  $P_k$  is the distribution of the maximum of  $X_{1..k}$ .

**Definition 1** (Jointly-Identical). For  $n$  identically distributed rvs  $X_{1..n}$  and  $\forall k \in [1 \dots n]$ , if any  $k$  arbitrarily chosen rvs

keep the same joint probability distribution, these  $n$  identical rvs are named as jointly-identical rvs.

**Lemma 1.** For  $n$  jointly-identical rvs  $X_{1..n}$ ,

$$P_{n,k} = \sum_{i=k}^n A_i^{n,k} P_i, 1 \leq \forall k \leq n,$$

where the const coefficient  $A_i^{n,k}$  can be calculated by the following recurrence:

$$A_i^{n,k} = \begin{cases} 1 & i = k, \\ -\sum_{j=1}^{i-k} \binom{n-i+j}{j} A_{i-j}^{n,k} & k+1 \leq i \leq n. \end{cases}$$

*Proof.* Let  $P_{n-k|k}$  denotes  $P(X_{k+1} > t, X_{k+2} > t, \dots, X_n > t | X_1 \leq t, X_2 \leq t, \dots, X_k \leq t)$  and  $P_{n-i, i-k|k}$  denotes  $P(X_{i+1} > t, X_{i+2} > t, \dots, X_n > t, X_{k+1} \leq t, X_{k+2} \leq t, \dots, X_i \leq t | X_1 \leq t, X_2 \leq t, \dots, X_k \leq t)$ ,  $k+1 \leq i \leq n$ . Certainly, we have

$$P_{n-k|k} = \frac{P_{n,k}}{P_k}, P_{n-i, i-k|k} = \frac{P_{n,i}}{P_k}. \quad (1)$$

As  $X_{1..n}$  are jointly-identical rvs, the following equation holds:

$$P_{n-k|k} = 1 - \sum_{i=k+1}^n \binom{n-k}{i-k} P_{n-i, i-k|k}. \quad (2)$$

By insertion of Eq. 1 into Eq. 2, the following recurrence holds:

$$P_{n,k} = P_k - \sum_{i=k+1}^n \binom{n-k}{i-k} P_{n,i} \quad (3)$$

Expanse Eq. 3, we complete the proof of Lemma 1.  $\square$

#### A. Linear Transformation of Order Statistics

**Theorem 1** (LT of Order Statistics). For  $n$  jointly-identical rvs  $X_{1..n}$ , there exists a linear transformation from maxima to order statistics:

$$\begin{bmatrix} X_{(n,1)} \\ X_{(n,2)} \\ \vdots \\ X_{(n,n)} \end{bmatrix} = \begin{bmatrix} W_1^{n,1} & W_2^{n,1} & \dots & W_n^{n,1} \\ 0 & W_2^{n,2} & \dots & W_n^{n,2} \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \dots & W_n^{n,n} \end{bmatrix} \begin{bmatrix} X_{(1,1)} \\ X_{(2,2)} \\ \vdots \\ X_{(n,n)} \end{bmatrix}.$$

Namely,  $X_{(n,k)} = \sum_{i=k}^n W_i^{n,k} X_{(i,i)}$ ,  $1 \leq \forall k \leq n$ , where the const coefficient

$$W_i^{n,k} = \sum_{j=k}^i \binom{n}{j} A_i^{n,j} \quad (4)$$

*Proof.* Let  $F_{n,k} \equiv P(X_{(n,k)} \leq t)$  be the probability distribution of the  $k^{th}$  order statistic. Equivalently, we need to prove

$$F_{n,k} = \sum_{i=k}^n W_i^{n,k} P(X_{(i,i)} \leq t) = \sum_{i=k}^n W_i^{n,k} P_i. \quad (5)$$

As  $F_{n,k} = \sum_{i=k}^n \binom{n}{i} P_{n,i}$  and  $P_{n,k} = \sum_{i=k}^n A_i^{n,k} P_i$  (Lemma 1), we derive the following recurrence:

$$F_{n,k} = \begin{cases} P_n & k = n, \\ F_{n,k+1} + \binom{n}{k} \sum_{i=k}^n A_i^{n,k} P_i & 1 \leq k \leq n-1. \end{cases} \quad (6)$$

Expanse Eq. 6, we complete the proof of Eq. 5.  $\square$

*Remark.* There is no need to assume the independence of the identical rvs  $X_{1..n}$ . There may exist other linear transformations of order statistics than the one given by Theorem 1.

**Definition 2** ( $W$  Coefficient). For any possible linear transformations from maxima to order statistics, the corresponding const coefficient  $W_i^{n,k}$  of  $X_{(i,i)}$  is named as a  $W$  coefficient.

*Remark.* The calculation of the  $W$  coefficient given by Eq. 4 is not straightforward, as it consists of many items. We use a simple solver (see Appendix) to give  $W$  coefficient values.

#### B. Linear Transformation of Expectations

Let  $\mathbb{E}_{n,k} \equiv E[X_{(n,k)}]$  be the expectation of the  $k^{th}$  order statistic of rvs  $X_{1..n}$ . Specially, we use  $\mathbb{E}_n$  to denote  $\mathbb{E}_{n,n}$ . Then the following theorem holds.

**Theorem 2** (LT of Expectations). For  $n$  jointly-identical rvs  $X_{1..n}$ , there exists a linear transformation from the expectations of maxima to the expectations of order statistics:

$$\mathbb{E}_{n,k} = \sum_{i=k}^n W_i^{n,k} \mathbb{E}_i, 1 \leq \forall k \leq n. \quad (7)$$

*Proof.*  $E[X_{(n,k)}] = E[\sum_{i=k}^n W_i^{n,k} X_{(i,i)}] = \sum_{i=k}^n W_i^{n,k} \mathbb{E}_i$   $\square$

### III. APPROXIMATIONS FOR NON-PURGING $(n, k)$ FORK-JOIN QUEUES

We consider a homogenous cluster consisting of  $n$  nodes, where each node  $i$  ( $1 \leq i \leq n$ ) has the same service time distribution when processing sub-tasks of the same job. Each node owns a first-come-first-serving sub-queue  $q_i$  ( $1 \leq i \leq n$ ) with the assumption of unlimited queue capacity. These  $n$  sub-queues constitute a homogenous fork-join queue.

Each incoming job consists of  $n$  tasks. Let  $t_i^j$  be the sojourn time of the  $j^{th}$  job's sub-task assigned to node  $i$ . Then, the stable sojourn time of a sub-task in the sub-queue  $q_i$  is  $t_i \equiv \lim_{j \rightarrow \infty} t_i^j$ , and the sojourn time of a job in the  $(n, k)$  fork-join queue is the  $k^{th}$  order statistic  $t_{(n,k)}$  consequently.

**Lemma 2.** For an  $(n, n)$  fork-join queue, the sub-queues' stable sojourn times  $t_{1..n}$  constitute a family of jointly-identical rvs.

*Proof.* Recall that all the sub-queues have an unlimited capacity, then the sub-task sojourn time distribution of a sub-queue depends only on the job arrival process and the sub-queue's service time distribution. As these sub-queues are under the same job arrival process and have the same service time distribution, the sub-queues' stable sojourn times  $t_{1..n}$  are identical rvs.

By the constitution methodology of fork-join queues, an  $(n, n)$  fork-join queue is symmetrical, which means its sub-queues are interchangeable and thus any  $k$  arbitrarily chosen sub-queues keep the same joint probability distribution. By definition the stable sojourn times  $t_{1..n}$  are jointly-identical rvs.  $\square$

**Definition 3** ( $(\lambda, \mu)$ -Equivalent Queues). Those basic fork-join queues, purging/non-purging  $(n, k)$  fork-join queues and split-merge queues that are under the same job arrival process and with the same sub-queue's service time distribution are called  $(\lambda, \mu)$ -equivalent queues to each other, where  $\lambda$  and  $\mu$  are the job arrival rate and the sub-queue's service rate respectively.

#### A. Approximations for General Fork-Join Queues

This paper uses the term `general` queues to denote the fork-join queues with identically and generally distributed sub-queues' service times and job inter-arrival times.

**Theorem 3** (LT of Sojourn Time). *The sojourn time of a general non-purging  $(n, k)$  fork-join queue can be represented by a linear combination of the sojourn times of the  $(\lambda, \mu)$ -equivalent basic fork-join queues:*

$$t_{(n,k)} = \sum_{i=k}^n W_i^{n,k} t_{(i,i)} \quad (8)$$

where  $W_i^{n,k}$  is the corresponding  $W$  coefficient.

*Proof.* The only difference between the  $(\lambda, \mu)$ -equivalent non-purging  $(n, k)$  queue and basic  $(n, n)$  queue lies in the job departure process, which has no influence over the job arrival process, the sub-queue's service time distribution and therefore the distribution of the sub-queue length. Consequently, the target non-purging  $(n, k)$  queue and its  $(\lambda, \mu)$ -equivalent  $(n, n)$  fork-join queue have the same family of jointly-identical rvs  $t_{1..n}$ , and therefore the same order statistic  $t_{(n,k)}$ .

According to Lemma 2 and Theorem 1,  $t_{(n,k)} = \sum_{i=k}^n W_i^{n,k} t_{(i,i)}^{sub}$ , where  $t_{(i,i)}^{sub}$  is the maximum stable sojourn time of  $i$  arbitrarily chosen sub-queues from the target non-purging  $(n, k)$  queue. By the constitution methodology of fork-join queues, the  $i$  chosen sub-queues can constitute an  $(i, i)$  queue which is  $(\lambda, \mu)$ -equivalent to the target non-purging  $(n, k)$  queue. Therefore,  $t_{(i,i)} = t_{(i,i)}^{sub}$  and  $t_{(n,k)} = \sum_{i=k}^n W_i^{n,k} t_{(i,i)}^{sub} = \sum_{i=k}^n W_i^{n,k} t_{(i,i)}$  hold.  $\square$

**Theorem 4** (LT of Expected Sojourn Time). *The expected sojourn time  $\mathbb{N}\mathbb{T}_{n,k} \equiv E[t_{(n,k)}]$  of a general non-purging  $(n, k)$  fork-join queue can be represented by a linear combination of the expected sojourn times of the  $(\lambda, \mu)$ -equivalent basic fork-join queues:*

$$\mathbb{N}\mathbb{T}_{n,k} = \sum_{i=k}^n W_i^{n,k} \mathbb{T}_i \quad (9)$$

where  $\mathbb{T}_i$  is the expected sojourn time of the  $(\lambda, \mu)$ -equivalent basic  $(i, i)$  fork-join queue and  $W_i^{n,k}$  is the corresponding  $W$  coefficient.

*Proof.*  $E[t_{(n,k)}] = E[\sum_{i=k}^n W_i^{n,k} t_{(i,i)}] = \sum_{i=k}^n W_i^{n,k} \mathbb{T}_i$   $\square$

*Remark.* The independence of the service times of the sub-queues is not required for Theorem 3 and 4 to hold. To put Theorem 4 into practices, we need to find the existing computing methods of  $\mathbb{T}_i$ .

#### B. Approximations for Exponential Fork-Join Queues

This paper uses the term `exponential` queues to denote the fork-join queues under a Poisson job arrival process and with iid exponential sub-queues' service time distributions.

For exponential  $(n, n)$  fork-join queues, there exist two reliable approximation methods: Nelson's approximation [4] and Varma's approximation [7]. Accordingly, we can propose two approximation methods for exponential non-purging  $(n, k)$  fork-join queues: The Nelson-LT approximation based on Nelson's approximation and the Varma-LT approximation based Varma's approximation.

1) *The Nelson-LT Approximation:* For exponential  $(n, n)$  fork-join queues, Nelson et al. [4] proposed the following approximations:

$$\begin{aligned} \mathbb{T}_1 &= \frac{1}{\mu(1-\rho)}, & \mathbb{T}_2 &= \frac{12-\rho}{8} \mathbb{T}_1 \\ \mathbb{T}_n &\simeq \left[ \frac{H_n}{H_2} + \frac{4}{11} \left( 1 - \frac{H_n}{H_2} \right) \rho \right] \mathbb{T}_2, n \geq 2, \end{aligned}$$

where  $\lambda$  and  $\mu$  are respectively the job arrival rate and the sub-queue's service rate,  $\rho \equiv \frac{\lambda}{\mu}$  is called the load factor of the queue,  $\mathbb{T}_n$  is the expected sojourn time of  $(n, n)$  basic fork-join queue, and  $H_n = \sum_{i=1}^n \frac{1}{i}$  is called the harmonic number. Consequently, our approximations can be specified in the following theorem.

**Theorem 5** (Nelson-LT Approximation). *According to Nelson's approximation, the expected sojourn time  $\mathbb{N}\mathbb{T}_{n,k}$  of an exponential non-purging  $(n, k)$  fork-join queue can be approximated as follow:*

$$\mathbb{N}\mathbb{T}_{n,k} \simeq \begin{cases} \frac{n}{\mu(1-\rho)} + \frac{12-\rho}{88\mu(1-\rho)} \times \sum_{i=2}^n W_i^{n,1} \left[ \frac{11H_i + 4\rho(H_2 - H_i)}{H_2} \right] & k = 1, \\ \frac{12-\rho}{88\mu(1-\rho)} \sum_{i=k}^n W_i^{n,k} \left[ \frac{11H_i + 4\rho(H_2 - H_i)}{H_2} \right] & k \geq 2. \end{cases} \quad (10)$$

where  $\lambda$  and  $\mu$  are respectively the job arrival rate and the sub-queue's service rate, and  $\rho \equiv \frac{\lambda}{\mu}$  is the load factor of the target queue. Specially, we replace any negatively approximated  $\mathbb{N}\mathbb{T}_{n,k}$  with 0.

We examine above linear-transformation approximations against the mean sojourn times of jobs sampled from various simulated exponential non-purging fork-join queues (details of simulations in this paper can be found in Appendix). The value of  $W$  coefficients used in Eq. 10 are given by Eq. 4. The results depicted in Fig. 3 confirmed the validity of our technique under a moderate value of  $n$  ( $n \leq 50$ ) and a relatively large value

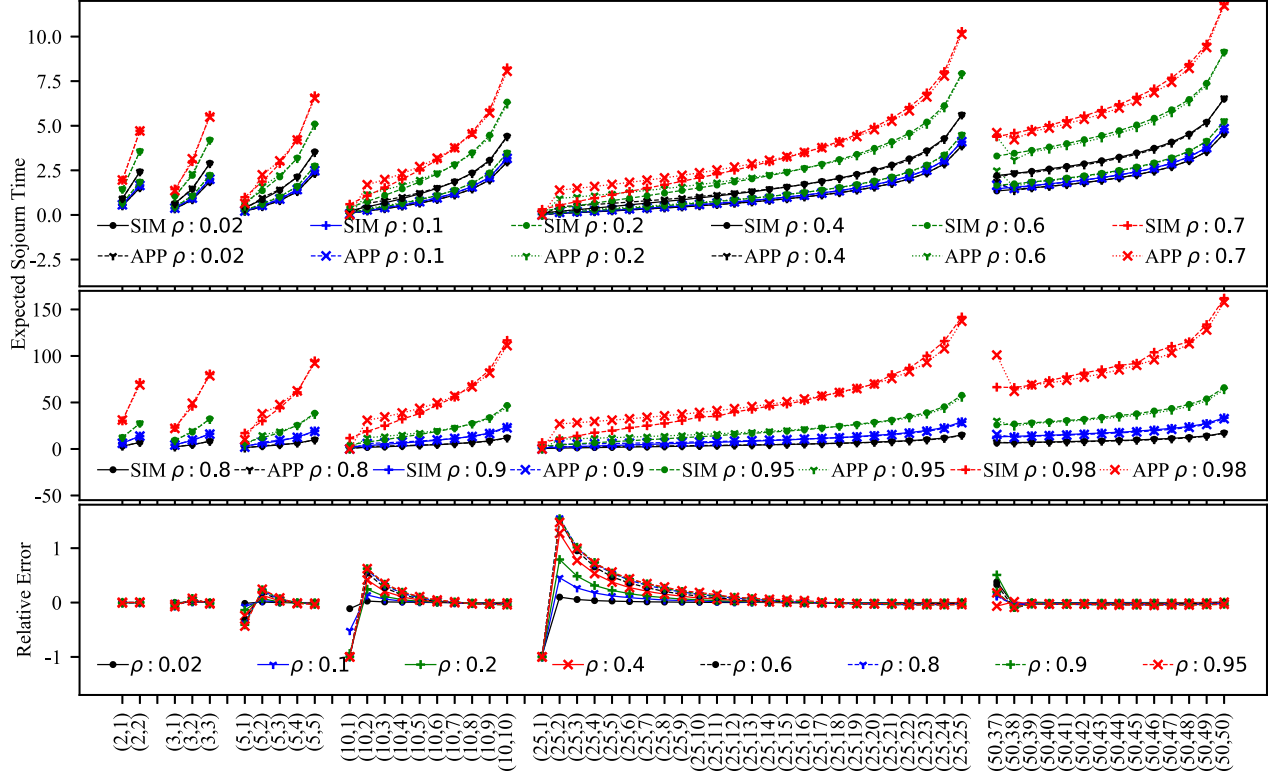


Fig. 3. The Nelson-LT Approximations for Exponential Non-Purging  $(n, k)$  Fork-Join Queues ( $\mu = 1$ , SIM: Simulation, APP: approximation)

of  $k$  (compared to  $n$ ). The relative errors are calculated as  $\frac{APP}{SIM} - 1$ .

We notice that when  $k$  is relatively small, the approximation tends to be uncontrollable, which can be due to the fact that the smaller  $k$  is, the more items in Eq. 10 are summed, and consequently, the more relative errors introduced by Nelson's approximations are accumulated.

These results also confirmed the high-precision merit of Nelson's approximations, since  $W$  coefficients tend to be very large with the increase of  $n$ , for example  $W_{16}^{25,9} = 13146544125$ . As a result, the relative error introduced by Nelson's approximation has to be amplified by the large value of the corresponding  $W$  coefficient.

2) *The Varma-LT Approximation:* For exponential non-purging  $(n, n)$  fork-join queues, Varma et al. [7] gave another well-known approximation method based on the so-called light traffic interpolation technique. The expected sojourn time is approximated as

$$\mathbb{T}_n \simeq \left[ H_n + (V_n - H_n) \frac{\lambda}{\mu} \right] \frac{1}{\lambda - \mu}, 0 \leq \lambda < \mu$$

where  $V_n = \sum_{r=1}^n \binom{n}{r} (-1)^{r-1} \sum_{m=1}^r \binom{r}{m} \frac{(m-1)!}{r^{m+1}}$ .

As our linear transformation technique is orthogonal to the concrete approximation methods for basic fork-join queues, we replace Nelson's approximation with the above Varma's approximation to try to avoid the approximations' uncon-

trollability appeared in Theorem 5. Consequently, the new approximations can be specified in the following theorem.

**Theorem 6** (Varma-LT Approximation). *According to Varma's approximation, the expected sojourn time of an exponential non-purging  $(n, k)$  fork-join queue can be approximated as follow:*

$$\mathbb{NT}_{n,k} \simeq \sum_{i=k}^n W_i^{n,k} \left[ H_i + (V_i - H_i) \frac{\lambda}{\mu} \right] \frac{1}{\lambda - \mu} \quad (11)$$

where  $\lambda$  and  $\mu$  are respectively the job arrival rate and the sub-queue's service rate.

We examine Eq. 11 against the mean sojourn times of jobs sampled from various simulated non-purging fork-join queues. The employed  $W$  coefficients are calculated from Eq. 4. The results depicted in Fig. 4 showed that the new approximations are fairly good when  $n \leq 10$  and  $\rho$  is not too extreme ( $\rho \leq 0.9$ ): The relative error is generally less than 10%, which is much more controllable than the approximations given by Theorem 5.

However, as the Varma's approximation itself becomes out of control when  $n \geq 55$ , Theorem 5 is more valuable in general cases.

#### IV. BOUNDS FOR PURGING $(n, k)$ FORK-JOIN QUEUES

Unlike in non-purging  $(n, k)$  fork-join queues, the sojourn time distribution of a sub-task in purging  $(n, k)$  fork-join

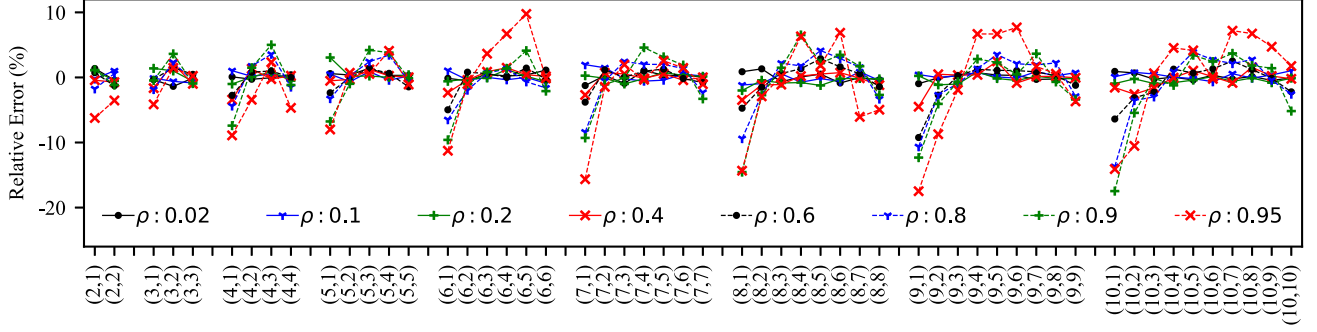


Fig. 4. The Relative Errors of the Varma-LT Approximations for Exponential Non-Purging  $(n, k)$  Fork-Join Queues

queues changes when either  $n$  or  $k$  varies, and thus differs from the sojourn time distribution of a sub-task in the  $(\lambda, \mu)$ -equivalent basic fork-join queues. As a result, we cannot build similar linear-transformation approximations for purging queues. However, the expected sojourn times of non-purging queues can serve as the upper bounds of the  $(\lambda, \mu)$ -equivalent purging queues' expected sojourn times.

#### A. The Naive Upper Bounds

**Theorem 7** (Naive Upper Bounds). *The expected sojourn time  $\mathbb{P}\mathbb{T}_{n,k}$  of a purging  $(n, k)$  fork-join queue can be upper bounded as follow:*

$$\mathbb{P}\mathbb{T}_{n,k} \leq \sum_{i=k}^n W_i^{n,k} \mathbb{T}_i \quad (12)$$

where  $\mathbb{T}_i$  is the expected sojourn time of the  $(\lambda, \mu)$ -equivalent basic fork-join queue.

*Proof.* The right side of Eq. 12 is the expected sojourn time of the  $(\lambda, \mu)$ -equivalent non-purging  $(n, k)$  fork-join queue. As the expected sub-queue length of a stable purging  $(n, k)$  fork-join queue is no longer than that of the  $(\lambda, \mu)$ -equivalent stable non-purging queue, the expected sojourn time of the purging  $(n, k)$  fork-join queue is thus no larger than that of its non-purging  $(\lambda, \mu)$ -equivalent queue.  $\square$

*a) Comparing with Existing Stat-of-the-Art Upper Bounds:* For purging  $(n, k)$  fork-join queues with iid service time distributions and under a Poisson job arrival process, Existing state-of-the-art upper bounds on the expected sojourn time are the *split-merge upper bounds* given by Joshi et al. [11]:

$$\mathbb{P}\mathbb{T}_{n,k} \leq E[X_{(n,k)}] + \frac{\lambda E[X_{(n,k)}^2]}{2(1 - \lambda E[X_{(n,k)}])} \quad (13)$$

where  $\lambda$  is the job arrival rate and  $X_{(n,k)}$  is the  $k^{th}$  order statistic of the iid service time rvs  $X_{1..n}$ .

The right side of Eq. 13 is the expected sojourn time of the  $(\lambda, \mu)$ -equivalent  $(n, k)$  split-merge queue.

**Corollary 1.** *Split-merge upper bounds become much looser than naive upper bounds when  $E[X_{(1,1)}] < E[X_{(n,k)}]$  and  $\lambda \rightarrow [\frac{1}{E[X_{(n,k)}]}]^-$ .*

*Proof.* When  $\lambda \rightarrow [\frac{1}{E[X_{(n,k)}]}]^-$ , the bounds given by Eq. 13 approach  $+\infty$ . Moreover, the bounds become meaningless when  $\lambda \geq \frac{1}{E[X_{(n,k)}]}$ . As a contrast, our naive bounds are finite meaningful values, as long as  $\lambda < [\mu \equiv \frac{1}{E[X_{(1,1)}]}]$ . Besides,  $\mu > \frac{1}{E[X_{(n,k)}]}$ .  $\square$

**Remark.** Apparently there is a range of load factors  $\rho \in [\frac{E[X_{(1,1)}]}{E[X_{(n,k)}]}, 1)$  cannot be bounded by Eq. 13, and the larger  $k$  is, the smaller bounded-able  $\rho$  range becomes, while the naive bounds are applicable as long as  $\rho < 1$ .

**Corollary 2.** *Naive upper bounds become much looser than split-merge upper bounds when  $k \rightarrow 1$ .*

*Proof.* When  $k \rightarrow 1$ , more and more sub-tasks are purged from both the sub-queues and the service stations when the  $k^{th}$  finished sub-task is acknowledged by the purging  $(n, k)$  queue, as a result of which, the expected sub-queue length becomes shorter and shorter than that of the  $(\lambda, \mu)$ -equivalent non-purging  $(n, k)$  queue. On the contrary, the expected sub-queue length of the target purging  $(n, k)$  fork-join queue becomes closer and closer to that of the  $(\lambda, \mu)$ -equivalent  $(n, k)$  split-merge queue. At last, the purging  $(n, 1)$  fork-join queue equates to the  $(n, 1)$  split-merge queue, which gives us the following exact closed-form solution:  $\mathbb{P}\mathbb{T}_{n,1} = E[X_{(n,1)}] + \frac{\lambda E[X_{(n,1)}^2]}{2(1 - \lambda E[X_{(n,1)}])}$ .  $\square$

**Remark.** On the other side of Corollary 2, when  $k \rightarrow n$ , the expected sub-queue length of the purging  $(n, k)$  fork-join queue becomes closer and closer to that of the  $(\lambda, \mu)$ -equivalent non-purging queues. At last, the purging  $(n, n)$  fork-join queue equates to the  $(\lambda, \mu)$ -equivalent non-purging  $(n, n)$  fork-join queue.

#### B. The Refined Upper Bounds

**Theorem 8** (Refined Upper Bounds). *For a purging  $(n, k)$  fork-join queue with iid service time distributions and under a*

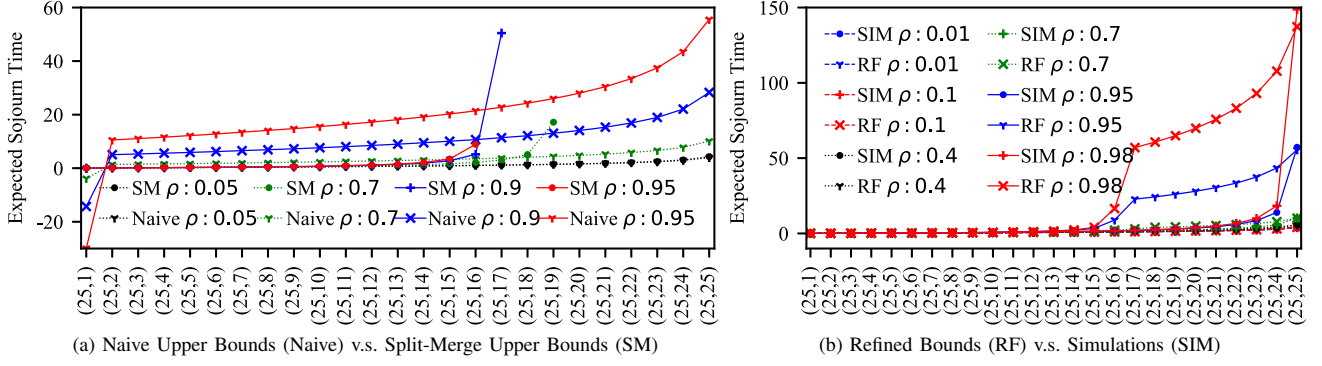


Fig. 5. Upper Bounds for Exponential Purging (25, 1), (25, 2), ..., (25, 25) Fork-Join Queues ( $\mu = 1$ )

Poisson job arrival process, the expected sojourn time  $\mathbb{PT}_{n,k}$  can be upper bounded as follows:

$$\mathbb{PT}_{n,k} \leq \begin{cases} \sum_{i=k}^n W_i^{n,k} \mathbb{T}_i & \lambda \geq \frac{1}{E[X_{(n,k)}]}, \\ \min \left( \sum_{i=k}^n W_i^{n,k} \mathbb{T}_i, \frac{E[X_{(n,k)}] + \frac{\lambda E[X_{(n,k)}^2]}{2(1-\lambda E[X_{(n,k)})}}}{\lambda} \right) & \text{otherwise.} \end{cases} \quad (14)$$

where  $\mathbb{T}_i$  is the expected sojourn time of the  $(\lambda, \mu)$ -equivalent  $(i, i)$  fork-join queue,  $\lambda$  is the job arrival rate, and  $X_{(n,k)}$  is the  $k^{th}$  order statistic of the iid service time rvs  $X_{1..n}$ .

*Proof.* According to Corollary 1 and 2, and excluding the split-merge bounds when  $\lambda \geq \frac{1}{E[X_{(n,k)}]}$ , we derive Eq. 14.  $\square$

*Remark.* Although Eq. 14 has extended the bounded-able range of  $\rho$  from  $[0, \frac{E[X_{(1,1)}]}{E[X_{(n,k)}]}]$  to  $[0, 1)$ , there is still an untamed range of  $\rho$ , since purging  $(n, k)$  fork-join queues may still keep stable even when  $\rho \geq 1$ .

a) *Upper Bounds for Exponential Queues:* Specially, we give the refined upper bounds for exponential purging  $(n, k)$  fork-join queues.

**Theorem 9** (Refined Upper Bounds for Exponential Purging Queues). *For exponential purging  $(n, k)$  fork-join queues, the expected sojourn time  $\mathbb{PT}_{n,k}$  can be upper bounded as follows:*

$$\mathbb{PT}_{n,k} \leq \begin{cases} \underbrace{\frac{12-\rho}{88\mu(1-\rho)} \sum_{i=k}^n W_i^{n,k} \left[ \frac{11H_i + 4\rho(H_2 - H_i)}{H_2} \right]}_{\text{Naive}} & \text{when } k \geq 2 \text{ and } \rho \geq \frac{1}{H_n - H_{n-k}}, \\ \min \left( \underbrace{\frac{H_n - H_{n-k}}{\mu}}_{\text{Split-Merge}}, \underbrace{\frac{\rho[(H_{n^2} - H_{(n-k)^2}) + (H_n - H_{(n-k)})^2]}{2\mu[1-\rho(H_n - H_{n-k})]}}_{\text{Naive}} \right) & \text{otherwise.} \end{cases} \quad (15)$$

where  $\lambda$  and  $\mu$  are respectively the job arrival rate and the sub-queue's service rate,  $\rho \equiv \frac{\lambda}{\mu}$ , and  $H_{n^2} = \sum_{i=1}^n \frac{1}{i^2}$ .

*Proof.* The split-merge part of the Eq. 15 is already given by Theorem 2 of [3]. According to Theorem 5 and 8, we derive Eq. 15.  $\square$

We make numerical comparisons between naive upper bounds and split-merge upper bounds, and examine refined upper bounds against the mean sojourn times of jobs sampled from various simulated purging fork-join queues. The value of  $W$  coefficients used in Eq. 15 are given by Eq. 4. We find that:

- The split-merge upper bounds become extremely pessimistic with the increase of  $k$ , but tend to be much tighter than naive bounds when  $k$  is small (see Fig. 5 (a)). These results are consistent with Corollary 1 and 2.
- There is still plenty of room of improving the upper bounds when  $k$  is relatively large (see Fig. 5 (b)).

### C. Lower Bounds

To complete our work, we review and compare the state-of-the-art lower bounds for purging  $(n, k)$  fork-join queues.

For purging  $(n, k)$  fork-join queues with iid service time distributions and under a Poisson job arrival process, the state-of-the-art lower bounds are the *split-merge lower bounds* given in [11]:

$$\mathbb{PT}_{n,k} \geq E[X_{(n,k)}] + \frac{\lambda E[X_{(n,1)}^2]}{2(1 - \lambda E[X_{(n,1)}])} \quad (16)$$

where  $\lambda$  is the job arrival rate and  $X_{(n,k)}$  is the  $k^{th}$  order statistic of the iid service time rvs  $X_{1..n}$ .

For exponential purging  $(n, k)$  fork-join queues, there is another staging analysis based lower bound [3]:

$$\mathbb{PT}_{n,k} \geq \frac{H_n - H_{n-k}}{\mu} + \frac{\rho(H_{n(n-\rho)} - H_{(n-k)(n-k-\rho)})}{\mu} \quad (17)$$

where  $\lambda$  and  $\mu$  are respectively the job arrival rate and the sub-queue's service rate,  $\rho \equiv \frac{\lambda}{\mu}$ , and  $H_{n(n-\rho)} = \sum_{i=1}^n \frac{1}{i(i-\rho)}$ .

This *staging lower bound* is adapted from the staging lower bound for basic fork-join queues proposed in [8], which requires a memory-less property of the service time distribution. Accordingly, this bound cannot be applied to general purging queues.






Fig. 6. Bounds v.s. Simulations for Exponential Purging (25, 1), (25, 2), ..., (25, 25) Fork-Join Queues ( $\mu = 1, \rho = 0.7$ )

**Theorem 10.** *The staging lower bounds are tighter than the split-merge lower bounds.*

*Proof.* The exact form of Eq. 16 for exponential queues can be transformed into:

$$\mathbb{P}\mathbb{T}_{n,k} \geq \frac{H_n - H_{n-k}}{\mu} + \frac{\rho(H_{n(n-\rho)} - H_{(n-1)(n-1-\rho)})}{\mu}.$$

As  $H_{(n-k)(n-k-\rho)} < H_{(n-1)(n-1-\rho)}$  when  $k > 1$ , we derive Theorem 10.  $\square$

We examine bounds for exponential purging queues against simulations. Fig. 6 depicts the large gap between the upper bounds and the lower bounds when  $k$  is relatively large, due to which, we can hardly find reasonable approximations of the expected sojourn time of purging  $(n, k)$  fork-join queues.

## V. DISCUSSION

Currently, there is an unnegligible limitation when put the new proposed linear transformation technique into practices: The value of  $W$  coefficient given by Eq. 4 increases explosively with the increase of  $n$ , for example  $W_{37}^{40,37} = \binom{40}{37} = 9880$ ,  $W_{37}^{50,37} = \binom{50}{37} = 354860518600$ , and  $W_{37}^{100,37} = \binom{100}{37} = 3.42002954749393 \times 10^{27}$ , as a result of which, the original negligible relative error of  $\mathbb{T}_i$  in Theorem 4 will be amplified into a huge deviation of the approximated  $\mathbb{NT}_{n,k}$ . Consequently, the practicability of the linear transformation technique depends on whether we can find high-precision approximated or simulated  $\mathbb{T}_i$ . For example, when we use simulated  $\mathbb{T}_i$  to estimate  $\mathbb{NT}_{n,k}$ , the results are far from acceptable (see Table I), and also far behind the Nelson-LT approximations (see Fig. 3). These surprising results can be due to the fluctuation of the simulated  $\mathbb{T}_i$  (see Fig. 7). Fig. 3 has depicted that the accuracy of the Nelson-LT approximation is acceptable when  $n \leq 50$  and  $k$  is relatively large (for example,  $k > 37$  when  $n = 50$ ). However the approximations are similarly unacceptable when  $k$  is relatively small (see Table II).

The fundamental solution of this problem is to scale down  $W$  coefficients, a promising research direction needs further efforts.

Anyway, this linear transformation technique is capable of estimating/bounding the performance of most of practical non-purging/purging  $(n, k)$  fork-join queueing systems, where




Fig. 7. Simulations (SIM) v.s. Nelson's Approximations for Basic (1, 1), (2, 2), ..., (20, 20) Fork-Join Queues ( $\mu=1$ )

TABLE I  
APPROXIMATED  $\mathbb{NT}_{n,k}$  BASED ON SIMULATED  $\mathbb{T}_i$  ( $\mu = 1$ )

$(n, k)$	$\rho: 0.05$	$\rho: 0.4$	$\rho: 0.8$
(20,17)	25.17029838	-37.93974618	70.76161205
(20,18)	0.412468275	6.238904164	3.024177347
(20,19)	2.772567927	3.796009747	11.47066833

TABLE II  
APPROXIMATED  $\mathbb{NT}_{n,k}$  BASED ON NELSON'S  $\mathbb{T}_i$  ( $\mu = 1$ )

$(n, k)$	$\rho: 0.2$	$\rho: 0.4$	$\rho: 0.8$
(50,34)	-317.7265625	-203.25	-580.0859375
(50,35)	23	-7.599609375	-18.87890625
(50,36)	-4.269042969	4.208007813	9.935546875

the replication factor ( $n$ ) rarely exceeds 10, a result of cost-effective tradeoff. For example, the replication factor of either Dynamo or Cassandra is commonly 3. Under such configurations, a write operation in Dynamo/Cassandra will be forked into 3 copies exactly. For such  $n \leq 10$  cases, we have proposed the fairly good Varma-LT approximations (Theorem 6).

From another perspective, the linear transformation technique can be used to check simulators' precision and to find better closed-form approximations for basic fork-join queues.

## VI. RELATED WORKS

*a) Order Statistics:* Bertsimas et al. [13] gave some tight bounds on the expectation of the  $k^{th}$  order statistic given the first and second moment information on  $n$  real-valued rvs. We gave exact linear transformations for  $k^{th}$  order statistic instead of bounds. Shi et al. [14] proposed a dynamic programming algorithm to compute the order statistics of a family of correlated rvs whose distributions are not required to be identical. This algorithm relies on the existence of computing methods for the distributions of both minimum and maximum of target rvs. As a contrast, our work is more formal and easier to practice for the reveal of the closed-form linear transformation, and only relies on the existence of computing methods for the maximum's distribution.

*b) Basic Fork-Join Queues:* For  $n = 2$  fork-join queues under a Poisson job arrival process: 1) Flatto et al. [15] gave the queue length distribution for exponential queues in stable state; 2) Baccelli [16] extended Flatto's work to queues with general service time distributions; 3) Nelson et al. [4] proposed



the exact closed-form solution of the expected sojourn time for exponential queues in stable state.

For  $n \geq 3$  exponential fork-join queues, the most influential approximation work was proposed by Nelson et al. in 1988 [4], which is based on the fact that the sojourn times  $X_{1..k}$  of sub-tasks  $1, 2, \dots, k$  are associated rvs [5], whose maximum can be bounded by the maximum of their iid equivalents. The lower bound is obtained by neglecting queueing effects. The approximation is a linear mixture of the upper and lower bounds. Parameters of the mixture are learned based on the mean sojourn times of jobs sampled from simulated basic fork-join queues. Varki et al. [8] improved the lower bound by using a staging analysis technique [17] based on the memory-less property of the exponential service time distribution, and use the mean value of Nelson's upper bound and the staging lower bound as the approximation. According to experiments in [6], Nelson's approximation is still the most reliable one for exponential queues, compared to following works including [7] and [8].

Varma et al. [7] extended Nelson's approximation to general service time distributions using a light traffic interpolation technique. Thomasian et al. [18] employed linear regression over the statistics of simulated fork-join jobs to find the parameters of their approximation equation for the expected sojourn time. However any change in service time distributions will require for re-simulations and re-regressions. Recently, Rizk et al. [19] proposed the first computable bounds on waiting and sojourn time of fork-join queues with general service time distributions by using martingales. However the upper bound is looser than Nelson's when it comes to the exponential service time distribution. Fidler et al. [12] considered the multi-stage nature of many fork-join queue networks, and proposed their end-to-end delay bounds.

We refer readers to [20] for a more comprehensive survey on fork-join queuing systems. To conclude, our work is orthogonal to existing approximation methods for basic fork-join queues.

*c) Purging  $(n, k)$  Fork-Join Queues:* There are some exact quantity analyses [9], [10] for purging  $(n, 1)$  fork-join queues, as it is equivalent to a single queue with  $n$  times the service rate. Gardner et al. [9] gave comprehensive research on purging  $(n, 1)$  fork-join queues, with considerations on multi-class jobs, interferences from un-forked jobs and heterogeneous service time distributions. Lee et al. [10] take the purging overheads into consideration, since the cancellation of running jobs typically incurs unnegligible delays in practice.

For purging  $(n, k > 1)$  fork-join queues, there are even no applicable approximations currently. Joshi et al. [3] extended the staging analysis to exponential  $(n, k)$  fork-join queues to find the lower bounds. Bounds for queues with general service time distributions are given by [3] and [11], by resorting the fork-join queue to the split-merge queue model, where all empty sub-queues are blocked until any  $k$  sub-tasks of the current job are completed. As depicted in Fig. 5 (a), the proposed upper bounds tend to be very rough when increasing  $k$  or the load factor  $\rho$ .

*d) Non-Purging  $(n, k)$  Fork-Join Queues:* A typical use case of non-purging  $(n, k)$  fork-join queues is the writing process in Cassandra [21]. Fidler et al. [12] gave non-asymptotic statistical bounds on the sojourn time of non-purging  $(n, k)$  fork-join queues. As a contrast, we give proper approximations instead of bounds.

## VII. CONCLUSION AND FUTURE WORK

Despite the popularity of  $(n, k)$  fork-join queues, there were no practical existing approximations of their expected sojourn times. Only some rough bounds have been given, which tend to be extremely loose when increasing  $k$  or the load factor  $\rho$ . This paper gave the first applicable approximation method for non-purging  $(n, k)$  fork-join queues and tackled the uncontrollability of the bounds for purging  $(n, k)$  fork-join queues:

- A brand-new closed-form linear transformation technique is developed for jointly-identical rvs, which provides a bridge to reduce the sojourn time approximation problem of non-purging  $(n, k)$  fork-join queues to that of basic fork-join queues.
- Improvements over upper bounds on the expected sojourn time of purging  $(n, k)$  fork-join queues are also gained by resorting the purging queues to their non-purging  $(\lambda, \mu)$ -equivalents.

Above innovations are examined by simulation experiments and numerically compared to the stat-of-the-arts. Results show that this linear transformation approach is practiced well for exponential  $(n, k)$  fork-join queues with moderate  $n$  and relatively large  $k$ . However, as currently found  $W$  coefficients (coefficients of the linear combination) increase explosively with the increase of  $n$ , there is an uncontrollable deviation in new proposed approximations when  $n$  is large and  $k$  is relatively small. Fortunately, approximations for real-life fork-join systems are unlikely influenced by this problem.

In the future, more efforts should be put into: Scaling down  $W$  coefficients, improving the approximations for basic fork-join queues with the help of the linear transformation technique, and evaluating the performance of real-life  $(n, k)$  fork-join systems as complement to existing experimental methods [22], [23].

## APPENDIX

The simulator employed by this work is Forkulator-p: <https://github.com/excelwang/forkulator-p>, which is modified from Forkulator [24] with additional features of simulating purging queues. For each  $(n, k)$  pair, the sub-queue's service rate  $\mu$  is set to 1, simulated jobs are sampled at a rate of 1%, and the mean sojourn times are calculated on 10000 samples.

The solver for currently found  $W$  coefficients and some pre-calculated values can be accessed from: <https://github.com/excelwang/WCoefficients>. We give some frequently used  $W_i^{n,k}$  coefficients bellow.

$W_1^{3,1} : 3$	$W_4^{6,3} : -45$	$W_7^{8,1} : 8$	$W_6^{9,2} : 420$	$W_3^{10,2} : -240$
$W_2^{3,1} : -3$	$W_5^{6,3} : 36$	$W_8^{8,1} : -1$	$W_7^{9,2} : -216$	$W_4^{10,2} : 630$
$W_3^{3,1} : 1$	$W_6^{6,3} : -10$	$W_2^{8,2} : 28$	$W_8^{9,2} : 63$	$W_5^{10,2} : -1008$
$W_2^{3,2} : 3$	$W_4^{6,4} : 15$	$W_3^{8,2} : -112$	$W_9^{9,2} : -8$	$W_6^{10,2} : 1050$
$W_3^{3,2} : -2$	$W_5^{6,4} : -24$	$W_4^{8,2} : 210$	$W_3^{9,3} : 84$	$W_7^{10,2} : -720$
$W_3^{3,3} : 1$	$W_6^{6,4} : 10$	$W_5^{8,2} : -224$	$W_4^{9,3} : -378$	$W_8^{10,2} : 315$
$W_1^{4,1} : 4$	$W_5^{6,5} : 6$	$W_6^{8,2} : 140$	$W_5^{9,3} : 756$	$W_9^{10,2} : -80$
$W_2^{4,1} : -6$	$W_6^{6,5} : -5$	$W_7^{8,2} : -48$	$W_6^{9,3} : -840$	$W_{10}^{10,2} : 9$
$W_3^{4,1} : 4$	$W_6^{6,6} : 1$	$W_8^{8,2} : 7$	$W_7^{9,3} : 540$	$W_3^{10,3} : 120$
$W_4^{4,1} : -1$	$W_1^{7,1} : 7$	$W_3^{8,3} : 56$	$W_8^{9,3} : -189$	$W_4^{10,3} : -630$
$W_2^{4,2} : 6$	$W_2^{7,1} : -21$	$W_4^{8,3} : -210$	$W_9^{9,3} : 28$	$W_5^{10,3} : 1512$
$W_3^{4,2} : -8$	$W_3^{7,1} : 35$	$W_5^{8,3} : 336$	$W_4^{9,4} : 126$	$W_6^{10,3} : -2100$
$W_4^{4,2} : 3$	$W_4^{7,1} : -35$	$W_6^{8,3} : -280$	$W_5^{9,4} : -504$	$W_7^{10,3} : 1800$
$W_3^{4,3} : 4$	$W_5^{7,1} : 21$	$W_7^{8,3} : 120$	$W_6^{9,4} : 840$	$W_8^{10,3} : -945$
$W_4^{4,3} : -3$	$W_6^{7,1} : -7$	$W_8^{8,3} : -21$	$W_7^{9,4} : -720$	$W_9^{10,3} : 280$
$W_4^{4,4} : 1$	$W_7^{7,1} : 1$	$W_8^{8,4} : 35$	$W_8^{9,4} : 315$	$W_{10}^{10,3} : -36$
$W_1^{5,1} : 5$	$W_2^{7,2} : 21$	$W_4^{8,4} : 70$	$W_9^{9,4} : -56$	$W_4^{10,4} : 210$
$W_2^{5,1} : -10$	$W_3^{7,2} : -70$	$W_5^{8,4} : -224$	$W_8^{9,5} : -315$	$W_5^{10,4} : -1008$
$W_3^{5,1} : 10$	$W_4^{7,2} : 105$	$W_6^{8,4} : 280$	$W_9^{9,5} : 70$	$W_6^{10,4} : 2100$
$W_4^{5,1} : -5$	$W_5^{7,2} : -84$	$W_7^{8,4} : -160$	$W_8^{9,5} : 126$	$W_7^{10,4} : -2400$
$W_5^{5,1} : 1$	$W_6^{7,2} : 35$	$W_8^{8,5} : -35$	$W_6^{9,5} : -420$	$W_8^{10,4} : 1575$
$W_2^{5,2} : 10$	$W_7^{7,2} : -6$	$W_5^{8,5} : 56$	$W_7^{9,5} : 540$	$W_9^{10,4} : -560$
$W_3^{5,2} : -20$	$W_3^{7,3} : 35$	$W_6^{8,5} : -140$	$W_8^{9,6} : 189$	$W_{10}^{10,4} : 84$
$W_4^{5,2} : 15$	$W_4^{7,3} : -105$	$W_7^{8,5} : 120$	$W_9^{9,6} : -56$	$W_5^{10,5} : 252$
$W_5^{5,2} : -4$	$W_5^{7,3} : 126$	$W_8^{8,6} : 21$	$W_6^{9,6} : 84$	$W_6^{10,5} : -1050$
$W_3^{5,3} : 10$	$W_6^{7,3} : -70$	$W_8^{8,6} : 28$	$W_7^{9,6} : -216$	$W_7^{10,5} : 1800$
$W_4^{5,3} : -15$	$W_7^{7,3} : 15$	$W_7^{8,6} : -48$	$W_8^{9,7} : -63$	$W_8^{10,5} : -1575$
$W_5^{5,3} : 6$	$W_4^{7,4} : 35$	$W_8^{8,7} : -7$	$W_9^{9,7} : 28$	$W_9^{10,5} : 700$
$W_4^{5,4} : 5$	$W_5^{7,4} : -84$	$W_7^{8,7} : 8$	$W_7^{9,7} : 36$	$W_{10}^{10,5} : -126$
$W_5^{5,4} : -4$	$W_6^{7,4} : 70$	$W_8^{8,8} : 1$	$W_8^{9,8} : 9$	$W_8^{10,6} : 945$
$W_5^{5,5} : 1$	$W_7^{7,4} : -20$	$W_1^{9,1} : 9$	$W_9^{9,8} : -8$	$W_9^{10,6} : -560$
$W_1^{6,1} : 6$	$W_5^{7,5} : 21$	$W_2^{9,1} : -36$	$W_9^{9,9} : 1$	$W_{10}^{10,6} : 126$
$W_2^{6,1} : -15$	$W_6^{7,5} : -35$	$W_3^{9,1} : 84$	$W_1^{10,1} : 10$	$W_6^{10,6} : 210$
$W_3^{6,1} : 20$	$W_7^{7,5} : 15$	$W_4^{9,1} : -126$	$W_2^{10,1} : -45$	$W_7^{10,6} : -720$
$W_4^{6,1} : -15$	$W_6^{7,6} : 7$	$W_5^{9,1} : 126$	$W_3^{10,1} : 120$	$W_8^{10,7} : -315$
$W_5^{6,1} : 6$	$W_7^{7,6} : -6$	$W_6^{9,1} : -84$	$W_4^{10,1} : -210$	$W_9^{10,7} : 280$
$W_6^{6,1} : -1$	$W_7^{7,7} : 1$	$W_7^{9,1} : 36$	$W_5^{10,1} : 252$	$W_{10}^{10,7} : -84$
$W_2^{6,2} : 15$	$W_1^{8,1} : 8$	$W_8^{9,1} : -9$	$W_6^{10,1} : -210$	$W_7^{10,7} : 120$
$W_3^{6,2} : -40$	$W_2^{8,1} : -28$	$W_9^{9,1} : 1$	$W_7^{10,1} : 120$	$W_8^{10,8} : 45$
$W_4^{6,2} : 45$	$W_3^{8,1} : 56$	$W_2^{9,2} : 36$	$W_8^{10,1} : -45$	$W_9^{10,8} : -80$
$W_5^{6,2} : -24$	$W_4^{8,1} : -70$	$W_3^{9,2} : -168$	$W_9^{10,1} : 10$	$W_{10}^{10,8} : 36$
$W_6^{6,2} : 5$	$W_5^{8,1} : 56$	$W_4^{9,2} : 378$	$W_{10}^{10,1} : -1$	$W_9^{10,9} : 10$
$W_3^{6,3} : 20$	$W_6^{8,1} : -28$	$W_5^{9,2} : -504$	$W_2^{10,2} : 45$	$W_{10}^{10,9} : -9$

## REFERENCES

- [1] A. Lakshman and P. Malik, "Cassandra: A decentralized structured storage system." *Operating Systems Review*, vol. 44, no. 2, pp. 35–40, 2010.
- [2] G. DeCandia, D. Hastorun, M. Jampani, G. Kakulapati, A. Lakshman, A. Pilchin, S. Sivasubramanian, P. Voshall, and W. Vogels, "Dynamo: Amazon's highly available key-value store," in *ACM SIGOPS symposium*. New York, New York, USA: ACM Press, 2007, pp. 205–220.
- [3] G. Joshi, Y. Liu, and E. Soljanin, "Coding for fast content download," in *2012 50th Annual Allerton Conference on Communication, Control, and Computing (Allerton)*. IEEE, 2012, pp. 326–333.
- [4] R. D. Nelson and A. N. Tantawi, "Approximate analysis of fork/join synchronization in parallel queues." *IEEE Trans. Computers*, vol. 37, no. 6, pp. 739–743, 1988.
- [5] J. D. Esary, F. Proschan, and D. W. Walkup, "Association of random variables, with applications," *The Annals of Mathematical Statistics*, vol. 38, no. 5, pp. 1466–1474, Oct. 1967.
- [6] A. Lebrecht and W. Knottenbelt, "Response time approximations in fork-join queues," 2007.
- [7] S. Varma and A. M. Makowski, "Interpolation approximations for symmetric fork-join queues." *Perform. Eval.*, vol. 20, no. 1, pp. 245–265, 1994.
- [8] E. Varki, "Response time analysis of parallel computer and storage systems." *IEEE Trans. Parallel Distrib. Syst.*, 2001.
- [9] K. Gardner, S. Zbarsky, S. Doroudi, M. Harchol-Balter, and E. Hytiä, "Reducing latency via redundant requests: Exact analysis." *SIGMETRICS*, pp. 347–360, 2015.
- [10] K. Lee, R. Pedarsani, and K. Ramchandran, "On scheduling redundant requests with cancellation overheads." *IEEE/ACM Transactions on Networking*, vol. 25, no. 2, pp. 1279–1290, 2017.
- [11] G. Joshi, E. Soljanin, and G. W. Wornell, "Efficient redundancy techniques for latency reduction in cloud systems." *ACM Transactions on Modeling and Performance Evaluation of Computing Systems (TOMPECS)*, vol. 2, no. 2, pp. 1–30, 2017.
- [12] M. Fidler and Y. Jiang, "Non-asymptotic delay bounds for (k, l) fork-join systems and multi-stage fork-join networks." *INFOCOM*, 2016.
- [13] D. Bertsimas, K. Natarajan, and C.-P. Teo, "Tight bounds on expected order statistics," *Probability in the Engineering and Informational Sciences*, vol. 20, no. 04, p. 441, Sep. 2006.
- [14] Y. Shi, J. Xiong, V. Zolotov, and C. Visweswariah, "Order statistics for correlated random variables and its application to at-speed testing." *ACM Transactions on Design Automation of Electronic Systems (TODAES)*, vol. 18, no. 3, pp. 42–20, Jul. 2013.
- [15] L. Flatto and S. Hahn, "Two Parallel Queues Created by Arrivals with Two Demands I," *SIAM Journal on Applied Mathematics*, vol. 44, no. 5, pp. 1041–1053, 1984.
- [16] F. Baccelli, "Two parallel queues created by arrivals with two demands: The M/G/2 symmetrical case," Ph.D. dissertation, INRIA, 1985.
- [17] K. S. Trivedi, *Probability and Statistics with Reliability, Queuing, and Computer Science Applications*. Prentice Hall PTR Upper Saddle River, NJ, USA, 1982.
- [18] A. Thomasian and A. N. Tantawi, "Approximate solutions for m/g/1 fork/join synchronization." *Winter Simulation Conference*, pp. 361–368, 1994.
- [19] A. Rizk, F. Poloczek, and F. Ciucu, "Computable bounds in fork-join queueing systems." *SIGMETRICS*, pp. 335–346, 2015.
- [20] A. Thomasian, "Analysis of fork/join and related queueing systems." *ACM Comput. Surv.*, vol. 47, no. 2, pp. 1–71, 2014.
- [21] X. Huang, J. Wang, J. Bai, G. Ding, and M. Long, "Inherent replica inconsistency in cassandra," in *2014 IEEE International Congress on Big Data (BigData Congress)*. IEEE, 2014, pp. 740–747.
- [22] H. Wang, J. Li, H. Zhang, and Y. Zhou, "Benchmarking replication and consistency strategies in cloud serving databases: Hbase and cassandra." in *BPOE@ASPLOS/VLDB, LNCS*, vol. 8807, no. Chapter 6. Springer, Cham, 2014, pp. 71–82.
- [23] Kuhlenskamp, Jörn, Klems, Markus, and Röss, Oliver, "Benchmarking scalability and elasticity of distributed database systems." *PVLDB*, 2014.
- [24] M. Fidler, B. Walker, and Y. Jiang, "Non-asymptotic delay bounds for multi-server systems with synchronization constraints," *arxiv: 1610.06309*, Oct. 2016.