MATHEMATICS OF OPERATIONS RESEARCH
Vol. 00, No. 0, Xxxxx 0000, pp. 000–000
ISSN 0364-765X | EISSN 1526-5471 | 00 | 0000 | 0001

INFORMS
DOI 10.1287/xxxx.0000.0000

Stationary waiting time in parallel queues with synchronization

Mariana Olvera-Cravioto
Department of Statistics and Operations Research, University of North Carolina at Chapel Hill, molvera@email.unc.edu

Octavio Ruiz-Lacedelli
Department of Industrial Engineering and Operations Research, Columbia University, or2200@columbia.edu

Motivated by database locking problems in today's massive computing systems, we analyze a queueing network with many servers in parallel (files) to which jobs (writing access requests) arrive according to a Poisson process. Each job requests simultaneous access to a random number of files in the database, and will lock them for a random period of time. Alternatively, one can think of a queueing system where jobs are split into several fragments that are then randomly routed to specific servers in the network to be served in a synchronized fashion. We assume that the system operates in a FCFS basis. The synchronization and service discipline create blocking and idleness among the servers, which leads to a strict stability condition compared to other distributed queueing models. We analyze the stationary waiting time distribution of jobs under a many servers limit and provide exact tail asymptotics; these asymptotics generalize the celebrated Cramér-Lundberg approximation for the single-server queue.

Key words: Queueing networks with synchronization, many servers queues, Cramér-Lundberg approximation, high-order Lindley equation, database locking, reader-writer queues, weighted branching processes, distributional fixed-point equations, large deviations
MSC2000 subject classification: 60K22, 60G10, 60F10
OR/MS subject classification: Primary: Queues - Networks; secondary: Probability - Stochastic model applications, Queues - Limit theorems,
History:

1. Introduction

Consider a large database with a total of n files. User requests arrive to the system according to a Poisson process, and each of them requires access to a random number of files from the database. Once a user is given access to its requested files, it will lock them for a random period of time, during which it may alter their content. Any other users requesting access to any file currently being locked will have to wait until all its requested files are available. Note that the locking mechanism is essential for maintaining consistency in the database. A variety of such models have been analyzed in the existing literature, with most of the work focusing on loss systems with no queueing [27, 26], or systems with two classes of users, readers and writers, each requesting access to one file at a time, but only the writers locking the files they view. An exception is given in [12], where writers can request simultaneous access to all n files. In the reader-writer queueing context, our model corresponds to a writer only system with no constraints in the number of files being requested.

Motivated by the rapid growth in the size of database systems, we analyze a queueing model for a large network of parallel servers (the files). Throughout the paper we use the generic term server to represent a computing unit, e.g., a file in a database or a processor in a computer network. Jobs arrive to the network at random times and are split at the time of arrival into a number of pieces. These pieces are then immediately assigned to randomly selected servers, where they join the corresponding queues. In terms of the database application, the random assignment is equivalent to having all subsets of k files having the same probability of being requested. In a computer network framework, the random assignment means that no centralized information regarding the workload at the different servers is needed. The service requirements of each of the fragments of a job are allowed to be random, dependent among themselves and on the total number of fragments in the job. In the database locking application, one can think of the service requirements (time the files will remain locked) being all the same, and becoming stochastically larger if many files are requested simultaneously; however, our model also allows files to be released as soon as they are no longer needed. The main distinctive features of this model are: 1) all the fragments of a job must begin their service at the same time, i.e., in a synchronized fashion; 2) jobs are processed in a first-come-first-serve (FCFS) basis, i.e., each of the individual queues at the servers follow a FCFS service discipline. We refer to these two characteristics as the *synchronization* and *fairness* requirements. Figure 1 depicts our model.

The fairness requirement is common in many queueing systems where jobs originate from different users, while the job synchronization is a distinctive characteristic of this model that allows us to incorporate the need to simultaneously make changes to various files or exchange information among the different fragments of a job during their processing. The synchronization of all the requested files is an essential feature of database locking systems. The fairness and synchronization requirements create, nonetheless, blocking and idleness that are not present in other distributed systems, e.g., in multi server queues where the different fragments of a job can be processed independently and can therefore be thought of as batch arrivals.

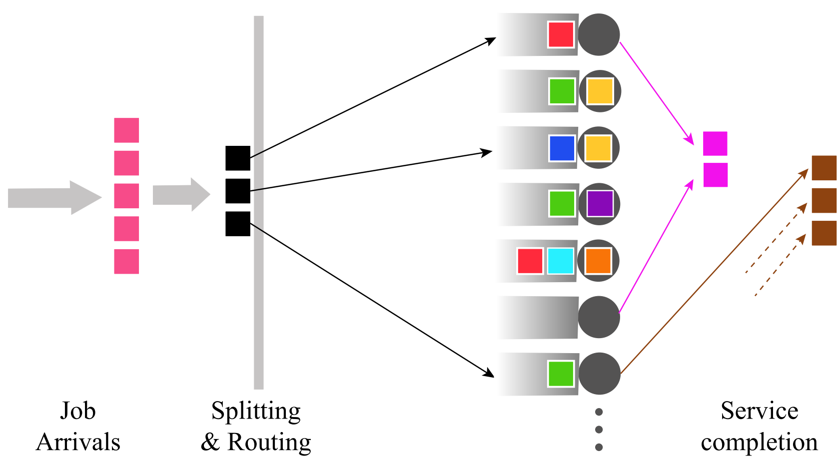

Figure 1. Queueing model for a database system where multiple files can be requested at a time. In this figure, the yellow, purple and orange jobs are being processed, while the pink and brown jobs have completed their service; two of the three brown pieces were processed at files/servers not depicted in the diagram. Note that the last file/server at the bottom will remain idle until the yellow and purple jobs complete their processing, and the first file/server at the top will need to wait for both the orange and light blue jobs to be done before starting to process the red job. The blue job in queue at the third server from the top has only one piece and can begin its processing as soon as the yellow job leaves.

We analyze in this paper the stationary waiting time of jobs (excluding service) in an asymptotic regime where the arrival rate of jobs and the number of servers grow to infinity but the distributions of the job sizes and of the service requirements of the individual fragments remain essentially constant. By waiting time we mean the amount of time a job has to wait until all its fragments can start being processed. For simplicity, we will refer to this type of limit as a "many server asymptotic regime", not to be confused with the Halfin-Whitt regime used in multiserver queueing systems [35], since our model does not lead to a heavy traffic regime and is in fact a lightly loaded system. In particular, after establishing sufficient conditions for the stability of the finite system, we show that the limiting stationary waiting time W of jobs is given by the all-time maximum in a branching random walk, and can be written in terms of the so-called endogenous solution[1] to a stochastic fixed-point equation of the form:

$$V \stackrel{\mathcal{D}}{=} \max\left\{ Y, \max_{1\leq i\leq N} (X_i + V_i) \right\}, \tag{1}$$

where the $\{V_i\}_{i\geq 1}$ are i.i.d. copies of V, independent of the vector $(Y, N, X_1, X_2, \dots)$, with $(Y, X_1, X_2, \dots)$ real-valued and N a nonnegative integer; $\stackrel{\mathcal{D}}{=}$ stands for equality in distribution. The random variable N will correspond to the number of fragments of a job, and the interpretation of Y and $\{X_i\}$ will be given later. Interestingly, whether W itself admits a representation like the one above or not depends on whether the service requirements of the fragments are allowed to depend on the total number of fragments N or not. When we have such independence we will have that W has the distribution of the endogenous solution to

$$W \stackrel{\mathcal{D}}{=} \left(\max_{1\leq i\leq N} (\chi_i - \tau_i + W_i) \right)^+, \tag{2}$$

where the $\{W_i\}_{i\in\mathbb{N}}$ are i.i.d. copies of W, independent of $(N, \chi_1, \tau_1, \chi_2, \tau_2, \dots)$, N is the number of fragments of a job, τ_i is the limiting interarrival time between piece i and the job immediately in front of it at its assigned queue, and χ_i is the service time of the fragment of the job in front of the ith piece; $x^+ = \max\{0, x\}$. Note that for $N \equiv 1$, (2) reduces to the classical Lindley equation, satisfied by the GI/GI/1 queue (in our model we assume that the $\{\tau_i\}$ are exponentially distributed). Recursion (2) was termed "high-order Lindley equation" and studied in the context of queues with synchronization in [20], although only for deterministic N.

Moreover, by applying the main result in [18] for the maximum of the branching random walk, we provide the exact asymptotics for the tail distribution of W. To explain the significance of this result it is worth considering first a single-server queue with renewal arrivals and i.i.d. service requirements, for which it is well known that the stationary waiting time distribution, $W^{\mathrm{GI/GI/1}}$, satisfies

$$P\left(W^{GI/GI/1} > x\right) = P\left(\max_{k\geq 0} S_k > x \right),$$

where $S_k = X_1 + \cdots + X_k$ is a random walk with i.i.d. increments satisfying $E[X_1] < 0$. Using the ladder heights of $\{S_k\}$ and renewal theory yields the celebrated Cramér-Lundberg approximation

$$P\left(W^{\mathrm{GI/GI/1}} > x\right) \sim H_{\mathrm{GI/GI/1}}\, e^{-\theta x}, \qquad x \to \infty,$$

where $0 < H_{\mathrm{GI/GI/1}} < \infty$ is a constant that can be written in terms of the limiting excess of the renewal process defined by the ladder heights, and $\theta > 0$, known as the Cramér-Lundberg root, solves $E[e^{\theta X_1}] = 1$ and satisfies $0 < E[X_1 e^{\theta X_1}] < \infty$ (see [7], Chapter XIII, for more details). Throughout the paper, $f(x) \sim g(x)$ as $x \to \infty$ stands for $\lim_{x\to\infty} f(x)/g(x) = 1$.

[1] The term "endogenous solution" refers to the fact that it can be explicitly written in terms of the random variables defining a marked Galton-Watson process; more details are given in Section 3.2.

Back to the asymptotic behavior of W in this paper, Theorem 3.4 in [18] states that for $\theta > 0$ satisfying the conditions

$$E\left[\sum_{i=1}^{N} e^{\theta X_i}\right] = 1 \qquad \text{and} \qquad 0 < E\left[\sum_{i=1}^{N} X_i e^{\theta X_i}\right] < \infty,$$

we have that the endogenous solution to (1) satisfies

$$P(V > x) \sim He^{-\theta x} \tag{3}$$

as $x \to \infty$ for some constant $0 < H < \infty$. This asymptotic will in turn imply that W satisfies

$$P(W > x) \sim HE[N]e^{-\theta x}$$

as $x \to \infty$. In other words, the queueing system with parallel servers and synchronization requirements in this paper naturally generalizes the single-server queue (since $N \equiv 1$ leads to (2)), as well as its Cramér-Lundberg approximation.

1.1. Related Literature As mentioned in the introduction, the inspiration for our model lies on large database management systems. In particular, on database locking models. Database locking models date back to the 1980's, with some of the early work [27, 26] focusing on systems where a user request that, upon arrival, finds that at least one of its requested files is locked, is rejected and leaves the system. In this context, the quantity of interest is the expected number of rejected requests. Similar loss models have also been used to analyze virtual path allocation problems in communication networks, where incoming calls request a specific set of links in the network (virtual path) to establish a communication channel, and if any of the requested links is unavailable at that time the call is lost (see, e.g., [34, 21]).

Also related to the setup in this paper, is the literature on reader-writer queueing models. In a reader-writer system there are two types of customers, readers and writers. Writers lock the files they are accessing while readers can access files that are currently being accessed by other readers. Both types of customers are forced to wait in line if a file is being locked by a writer who arrived earlier. Reader-writer systems have been analyzed in [29, 22, 23], however, in all of those cases readers and writers request access to a single file at a time, which corresponds to $N \equiv 1$ in our model. The work in [22, 23] covers several combinations of priorities among the readers and writers that allow the exact analysis of various performance measures. Even closer to our model is the work done in [12], where the authors consider a reader-writer system where readers request access to a single file at a time, while writers require locking access to all n files simultaneously, which is unrealistic in very large databases. In the setting of reader-writer systems, our model corresponds to a writer-only system where writers can request any subset of the n files and will lock them for a random period of time. Most of the existing modeling approaches for reader-writer systems are fairly stylized, and sacrifice some realistic features on behalf of theoretical tractability. Ours fully generalizes the number of files that can be selected at a time as well as the length of time that they need to be blocked, however, it ignores the possibility that certain files may be requested more often than others.

Other closely related literature includes the work on distributed queueing models with synchronized service. We start with the model considered in [14], which studies a queueing system where each job requires a synchronous execution on a random number of parallel servers. Some applications mentioned there are: the deployment of fire engines in firefighting, jury selection, and the staffing of surgeons and medical personnel in emergency surgery. The main difference in [14] from our setup, besides the restriction to i.i.d. exponentially distributed service times, is that we assign the pieces of a job to specific servers at the time of arrival, while the model in [14] waits until the required number of servers is free and then assigns the pieces to these servers. Some of the ideas used in the proof of the main result in this paper are borrowed from [20], where the authors considered a queueing system with m different types of servers and n identical servers of each type (for a total of $m \times n$ servers), and where each arriving job requires service from exactly one server of each type, i.e., each job needs m parallel servers, and is assigned upon arrival to one of the n possible choices for each type. Theorem 2 in [20] shows that the steady-state distribution of the waiting time converges weakly, as $n \to \infty$, to the endogenous solution of the high-order Lindley equation (2) with $N \equiv m$. Besides allowing N to be random and the service requirements of the fragments of a job to be dependent, this paper shows not only the weak convergence of the steady-state waiting time, but also of all its moments. The proof technique used in this paper, which is based on a coupling between a graph and a weighted branching process using the Wasserstein distance, is new, and is responsible for the stronger mode of convergence. A third related distributed queueing model is the one considered in [9], which can be thought of as a stylized version of our queueing network where server assignment is not done uniformly at random, but rather according to a distribution on specific subsets of servers, e.g., blocks of adjacent or closely located servers. The setting there corresponds to all the fragments of a job having identical service requirements, but provides interesting insights into the existence of stationary distributions for different server assignment rules.

Finally, we mention that our model is also related to fork-join queues with synchronization, in particular, to the so-called split-merge queue [15, 25, 24, 32]. The main difference between a split-merge queue and the model described in this paper is that in the former the synchronization occurs once all the pieces of a job have completed their service. More precisely, job fragments are allowed to start their processing as soon as their assigned server becomes available, but will continue blocking it, even after having completed their service, until all other pieces of the same job have completed theirs.

The remainder of the paper is organized as follows. Section 2 contains the mathematical description of our queueing model; Section 3 describes the analysis of the stationary waiting time of jobs in the network, with the main result of this paper in Section 3.1 and the tail asymptotics of the limiting waiting time in Section 3.2. Finally, the proofs of the two theorems are given in the Appendix.

2. Model description

We consider a sequence of queueing networks indexed by their number of servers, n. The n servers are identical and operate in parallel. Arrivals to the nth network occur according to a Poisson process with rate λn for some parameter $\lambda > 0$. Each job, upon arrival to the network, is split into a random number of fragments. The size of a job, i.e., the number of pieces into which it is split, is determined by some distribution $f_n(k)$, $k = 1, 2, \ldots, n$; this condition on the support of f_n ensures that each piece can be routed to a different server. Once a job has been split, say into k pieces, its fragments are routed randomly to k different servers in the network (i.e., with all $n!/(n-k)!$ possible assignments equally likely), forming a queue at their assigned servers. An equivalent way of describing the arrival of jobs into the nth network is to use the thinning property of the Poisson process and think of independent Poisson processes, each generating jobs of size k, $k = 1, 2, \ldots, n$, at rate $f_n(k)\lambda n$.

The service times of the different fragments of a job are assumed to have a general distribution, although the stability condition for the model will implicitly impose that they have finite exponential moments. Moreover, they are not assumed to be identically distributed and are allowed to be dependent with each other and with the number of pieces. More precisely, a typical job has $\hat{N}$ pieces having service requirements $(\xi^{(1)}, \ldots, \xi^{(\hat{N})})$, where $\hat{N}$ has marginal probability mass function f_n, and the vector $(\hat{N}, \xi^{(1)}, \ldots, \xi^{(\hat{N})})$ is arbitrarily dependent.

Since the random routing eliminates information about the order of the pieces, the relevant distribution that will appear in the analysis of the model is that of a randomly chosen fragment. In particular, we will use $(\hat{N}, \hat{\chi})$ to denote a vector distributed according to:

$$P(\hat{N} = k, \hat{\chi} \leq x) = f_n(k) \sum_{i=1}^{k} \frac{1}{k} P(\xi^{(i)} \leq x | \hat{N} = k) \triangleq f_n(k) G_k(x).$$

The sizes of jobs and of the service requirements of their fragments are assumed to be independent of the arrival process.

In order to model the synchronization and fairness characteristics of the network, we will assign to each job a tag (not to be confused with the label that will be introduced later). More precisely, a job having k pieces receives a tag of the form $(s_1, s_2, \ldots, s_k)$, $s_i \in \{1, 2, \ldots, n\}$ for all i, $s_i \neq s_j$ for $i \neq j$, representing the different servers to which its fragments are sent for processing.

Definition 1. We say that a job having tag $\mathbf{r} = (r_1, r_2, \ldots, r_l)$ is a *predecessor* of a job having tag $\mathbf{s} = (s_1, s_2, \ldots, s_k)$ if it arrived before the job having tag $\mathbf{s}$ and they have at least one server in common (i.e., $r_i = s_j$ for some $1 \leq i \leq l$ and $1 \leq j \leq k$). We use the term *immediate predecessor* if there are no fragments between them in queue at the server(s) they have in common. A job having $\hat{N}$ fragments has up to $\hat{N}$ immediate predecessors.

In terms of this definition, the *synchronization* rule is that the job having tag $\mathbf{s} = (s_1, s_2, \ldots, s_k)$ cannot begin its service, which is to be done in parallel by servers $s_1, s_2, \ldots, s_k$, until all its immediate predecessors have completed their service. The *fairness* rule says that if the job with tag $\mathbf{r}$ is a predecessor of the job with tag $\mathbf{s}$, then it will begin its service before the job with tag $\mathbf{s}$ does.

Formally, we can think of the nth system as a superposition of $\sum_{k=1}^{n} \frac{n!}{(n-k)!}$ independent marked point processes. Each of these processes generates jobs of size k with server assignments $(s_1, \ldots, s_k)$, according to a Poisson process with rate $f_n(k)\lambda n/(n!/(n-k)!)$, and having i.i.d. marks $\left\{(\xi_i^{(1)}, \ldots, \xi_i^{(k)})\right\}_{i \geq 1}$ with common distribution

$$B_k(x_1, \ldots, x_k) \triangleq P(\xi^{(1)} \leq x_1, \ldots \xi^{(k)} \leq x_k | \hat{N} = k),$$

corresponding to the service requirements of the fragments. To establish stability, we follow the standard queueing theory technique of assuming that at time $t_0 < 0$ there are no jobs in the network, and then look at the waiting time of the first job to arrive after time zero. We will refer to this job as the "tagged" job, and we will prove that its waiting time, after having taken the limit $t_0 \to -\infty$, is finite almost surely. The stability will then follow from Loynes' lemma and Palm theory (see [8] for more details on this general technique). Note that we will implicitly be using the Palm probability associated to the superposition of the marked Poisson processes described above and the probability P, which by PASTA will be equivalent to studying the stationary distribution of the process under P.

To analyze the waiting time of the tagged job we look at a graph containing all the information of which jobs need to complete their service before the tagged job can initiate its own. We now describe how to construct such a graph, called in [20] a predecessor graph.

2.1. The predecessor graph To construct the predecessor graph we look at time in reverse, starting from the time the tagged job arrived, say $T_1 \geq 0$, and ending at time t_0. The tagged job, which we will label $\emptyset$, is split into a random number of pieces, say $\hat{N}_\emptyset = \hat{N}_\emptyset(n)$, where $\hat{N}_\emptyset$ is distributed according to f_n. Each of these pieces will be routed to one of the n servers in the network, where it will either find the server empty or join a queue. Suppose that the tagged job needs to be processed by servers $(s_1, \ldots, s_{\hat{N}_\emptyset})$, and recall from Definition 1 that any job that has a fragment requiring service at any of the servers s_i, $1 \leq i \leq \hat{N}_\emptyset$, with no other jobs in between, is an immediate predecessor of the tagged job. To construct the first set of edges in the graph we draw an edge from the tagged job to its immediate predecessors. In general, each of the $\hat{N}_\emptyset$ fragments of the tagged job has an immediate predecessor, unless its assigned server has never been assigned any other jobs since time t_0 (the time at which the system starts empty). Hence, the number of edges we draw is smaller or equal than $\hat{N}_\emptyset$. Moreover, each edge is assigned a vector of the form $(\hat{\tau}_i, \hat{\chi}_i)$, $1 \le i \le \hat{N}_\emptyset$, where $\hat{\tau}_i$ is the interarrival time between the tagged job and its ith immediate predecessor, and $\hat{\chi}_i$ is the service requirement of the fragment of the immediate predecessor that is in front of the corresponding fragment of the tagged job. Also, if a job is an immediate predecessor to more than one fragment of the tagged job, say it requires service at servers s_i and s_j, then $\hat{\tau}_i = \hat{\tau}_j$, although we may still have $\hat{\chi}_i \neq \hat{\chi}_j$ with $\hat{\chi}_i, \hat{\chi}_j$ possibly dependent.

Iteratively, once we have identified all the immediate predecessors of the tagged job we repeat the process described above with each one of them. We will call the predecessor graph $\mathcal{G}_n(t_0)$, since it will depend on both the number of servers n and the time t_0 at which the system starts empty. We point out that the predecessor graph indicates the arrival times of all predecessors of the tagged job, but it does not indicate whether those predecessors are still in the system at time T_1 or not, that is, some predecessors may have completed their service long before time T_1, and it is even possible that the tagged job encounters, upon arrival, all its assigned servers empty. The same is true for all other jobs in the graph.

Since $\mathcal{G}_n(t_0)$ will resemble a tree, it will be useful to use tree notation to refer to the predecessors of the tagged job. More precisely, let $\mathbb{N}_+ = \{1, 2, 3, \dots\}$ be the set of positive integers and let $U = \bigcup_{r=0}^\infty (\mathbb{N}_+)^r$ be the set of all finite sequences $\mathbf{i} = (i_1, i_2, \dots, i_r) \in U$, where by convention $\mathbb{N}_+^0 = \{\emptyset\}$ contains the null sequence $\emptyset$. To ease the exposition, for a sequence $\mathbf{i} = (i_1, i_2, \dots, i_k) \in U$ we write $\mathbf{i}|t = (i_1, i_2, \dots, i_t)$, provided $k \ge t$, and $\mathbf{i}|0 = \emptyset$ to denote the index truncation at level t, $k \ge 0$. To simplify the notation, for $\mathbf{i} \in \mathbb{N}_+$ we simply use $\mathbf{i} = i_1$, that is, without the parenthesis. Also, for $\mathbf{i} = (i_1, \dots, i_k)$ we will use $(\mathbf{i}, j) = (i_1, \dots, i_k, j)$ to denote the index concatenation operation, if $\mathbf{i} = \emptyset$, then $(\mathbf{i}, j) = j$.

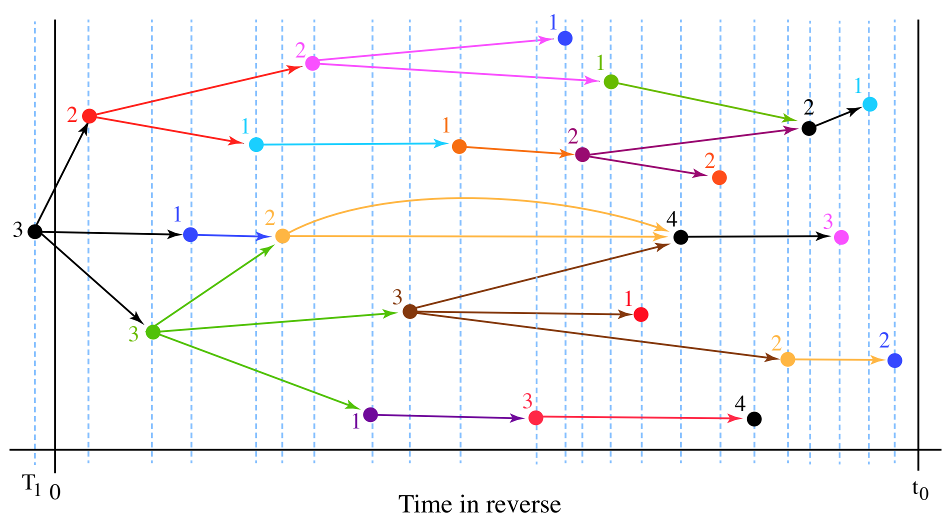

Figure 2. The predecessor graph $\mathcal{G}_n(t_0)$. The numbers in each node indicate the size of the job (number of fragments). Some nodes have fewer outbound edges than the size of the job, meaning that the corresponding fragment was the first to be processed at its assigned server since time t_0. Nodes with multiple inbound edges correspond to jobs that are immediate predecessors to more than one job in the graph. Vertical lines indicate the time of arrival of each job; service requirements for the fragments of a job can be thought of as "edge attributes" and cannot be read from the graph. This graph is consistent with Figure 1 by letting the tagged job be the three piece black one.

Now recall that $\emptyset$ denotes the tagged job and label its immediate predecessors i, with $1 \le i \le \hat{N}_\emptyset$. The jobs in the next level of predecessors will have labels of the form (i_1, i_2), and in general, any job in the predecessor graph will have a label of the form $\mathbf{i} = (i_1, i_2, \ldots, i_k)$, $k \geq 1$. With this notation, $\hat{N}_{\mathbf{i}}$ denotes the number of pieces that the job with label $\mathbf{i}$ in the graph is split into, each $\hat{\tau}_{(\mathbf{i},j)}$ will denote the interarrival time between job $\mathbf{i}$ and its jth immediate predecessor (a job with label $(\mathbf{i}, j)$), and $\hat{\chi}_{(\mathbf{i},j)}$ denotes the service requirement of the fragment who must receive service immediately before the fragment represented by $(\mathbf{i}, j)$ can start being processed at its assigned server. Note that the tag of a job, which contains the specific server assignments, allows us to identify the immediate predecessors of a given job, but it plays no role afterwards (and is not recoverable from the labels). Therefore, we will use the labels, not tags, to identify jobs in the predecessor graph. See Figure 2.

We point out that in case a job is an immediate predecessor to more than one job in the graph (or to more than one fragment of the same job), the corresponding edges will merge into the common predecessor. Moreover, in this case, the common predecessor is assigned more than one label (e.g., if a job is an immediate predecessor to both jobs $\mathbf{i} = (i_1, \ldots, i_k)$ and $\mathbf{j} = (j_1, \ldots, j_l)$, then such job can be identified by two different labels, one of the form $(\mathbf{i}, s)$ and another of the form $(\mathbf{j}, t)$). Furthermore, the merged paths and the subgraph they define from that point onwards will have multiple labels as well. See Figure 3.

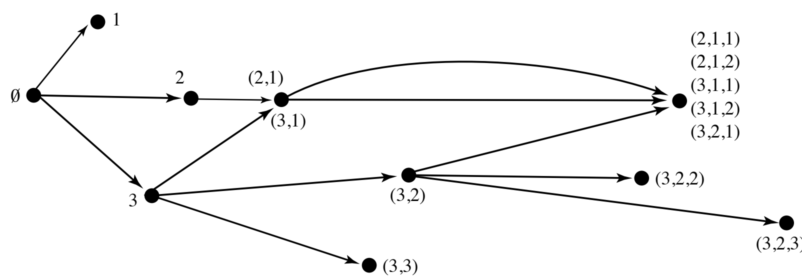

FIGURE 3. Multiple labels due to common immediate predecessors. Excerpt from the predecessor graph in Figure 2 showing the labeling of the jobs.

REMARK 1. (i) We allow multiple labels for jobs that are immediate predecessors to more than one job (or more than one fragment of the same job), since each path leading to a job represents a distinct sequence of fragments, with its own waiting time.

(ii) Note that the predecessor graph includes all the predecessors of the tagged job, regardless of whether they were present in the network at the time the tagged job arrived or not. In fact, the predecessor graph, as depicted in Figure 2 does not tell us anything about whether the jobs being depicted have completed their service by time T_1 or not.

(iii) As pointed out earlier, the concept of the predecessor graph was introduced in [20], although in less detail than here, e.g., the labeling of jobs is not rigorous and there is no mention of multiple labels. In our case, the new coupling technique that we use, which yields a stronger mode of convergence in the main theorem and allows us to take $\hat{N}$ to be random and $(\hat{N}, \xi^{(1)}, \ldots, \xi^{(\hat{N})})$ arbitrarily dependent, requires the more careful treatment described above.

3. Analysis of the steady-state waiting time To analyze the waiting time of the tagged job we now derive a representation in terms of a branching random walk. To this end, let $W_{\mathbf{i}}^{(n,t_0)}$ denote the waiting time of the job having label $\mathbf{i}$ in the network with n servers and that starts empty at time t_0. We also define $\hat{A}_0 = \{\emptyset\}$, and $\hat{A}_r = \{(i_1, \ldots, i_r) \in \mathcal{G}_n(t_0)\}$ for $r \geq 1$, to be the set of labels in the predecessor graph at graph distance r from the tagged job, i.e., labels whose corresponding job is connected to the tagged job by a directed path of length r. For $\mathbf{i} \in \hat{A}_k$ and $r \geq 1$ let

$$\mathcal{B}_{\mathbf{i},r} = \left\{ \mathbf{j} \in \hat{A}_{k+r} : \mathbf{j} = (\mathbf{i}, i_{k+1}, \ldots, i_{k+r}) \right\}$$

be the set of labels at distance r from $\mathbf{i}$.

We then have that the tagged job's waiting time is given by

$$W_\emptyset^{(n,t_0)} = \max\left\{0,\ \max_{i\in\mathcal{B}_{\emptyset,1}} \left(\hat{\chi}_i - \hat{\tau}_i + W_i^{(n,t_0)}\right)\right\}, \tag{4}$$

with the boundary condition that the first job to arrive after time t_0, and any other job that arrives thereafter and is the first one to use its assigned servers, will have a waiting time of zero (recall that the system is empty at time t_0).

To analyze (4) define $\hat{X}_\emptyset \equiv 0$ and $\hat{X}_\mathbf{i} = \hat{\chi}_\mathbf{i} - \hat{\tau}_\mathbf{i}$ for $\mathbf{i} = (i_1, \ldots, i_k)$. Next, let

$$\kappa = \max\{r \in \mathbb{N}_+ : |\mathcal{B}_{\emptyset,r}| > 0\},$$

where $|A|$ denotes the cardinality of set A. Note that κ is a random variable and corresponds to the maximum length of any directed path in $\mathcal{G}_n(t_0)$. Furthermore, for any $\mathbf{i} \in \hat{A}_\kappa$ we have $W_\mathbf{i}^{(n,t_0)} = 0$, and therefore, for any $\mathbf{i} \in \hat{A}_{\kappa-1}$,

$$W_\mathbf{i}^{(n,t_0)} = \max\left\{0, \max_{\mathbf{j}\in\mathcal{B}_{\mathbf{i},1}} \hat{X}_\mathbf{j}\right\}.$$

Similarly, iterating (4) we obtain for $\mathbf{i} \in \hat{A}_{\kappa-2}$,

$$\begin{aligned} W_\mathbf{i}^{(n,t_0)} &= \max\left\{0,\ \max_{\mathbf{j}\in\mathcal{B}_{\mathbf{i},1}} \left(\hat{X}_\mathbf{j} + W_\mathbf{j}^{(n,t_0)}\right)\right\} \\ &= \max\left\{0,\ \max_{\mathbf{j}\in\mathcal{B}_{\mathbf{i},1}} \hat{X}_\mathbf{j},\ \max_{\mathbf{j}\in\mathcal{B}_{\mathbf{i},2}} \left(\hat{X}_{\mathbf{j}|\kappa-1} + \hat{X}_\mathbf{j}\right)\right\}. \end{aligned}$$

In general, after iterating (4) κ times we obtain

$$W_\emptyset^{(n,t_0)} = \max\left\{0, \max_{j\in\mathcal{B}_{\emptyset,1}} \hat{X}_j, \ldots, \max_{\mathbf{j}\in\mathcal{B}_{\emptyset,\kappa}} \left(\hat{X}_{\mathbf{j}|1} + \hat{X}_{\mathbf{j}|2} + \cdots + \hat{X}_\mathbf{j}\right)\right\}. \tag{5}$$

Having now a recursive equation for the waiting time, we need to identify conditions under which this queueing system will be stable, and then describe the stationary distribution of the waiting time. The key idea for solving both problems is that the predecessor graph is very close to being a tree; more precisely, the only thing that prevents it from being a tree is the occasional arrival of a job that is an immediate predecessor to two or more jobs in $\mathcal{G}_n(t_0)$ (and the dependence it introduces). It turns out that under the scaling we consider in our model (arrival rate equal to λn), the probability of this occurring within the timeframe needed for the tagged job to start its service is very small. Once we show that this is the case, taking the limit as $t_0 \to -\infty$ will yield the stability of the network.

To describe the stationary distribution of the waiting time we first observe that, provided the first time that two paths in the predecessor graph merge occurs after the tagged job has initiated its service, we have that the $\hat{\tau}$'s will be i.i.d. exponential random variables with some rate λ_n^* and the vectors $\{(\hat{N}_\mathbf{i}, \hat{\chi}_\mathbf{i})\}$ will be i.i.d. (since they will all belong to different jobs). We will assume that $(\hat{N}, \hat{\chi}, \hat{\tau}) \Rightarrow (N, \chi, \tau)$ as $n \to \infty$, where $\Rightarrow$ denotes convergence in distribution. It follows that under the same conditions that guarantee the stability of the network we would have that, after taking the limit as $t_0 \to -\infty$ and the number of servers $n \to \infty$, $W_\emptyset^{(n,t_0)}$ would have to converge to the random variable

$$W \triangleq \bigvee_{r=0}^{\infty} \bigvee_{\mathbf{i}\in A_r} S_\mathbf{i},$$

where $S_\emptyset = 0$, $S_{\mathbf{i}} = X_{\mathbf{i}|1} + X_{\mathbf{i}|2} + \cdots + X_{\mathbf{i}}$ for $\mathbf{i} \neq \emptyset$, $X_{\mathbf{i}} = \chi_{\mathbf{i}} - \tau_{\mathbf{i}}$, and A_r is the rth generation of a Galton-Watson process whose offspring distribution is that of N. Note that we have replaced the sets $\mathcal{B}_{\emptyset,r}$ with A_r, since in stationarity all the fragments of a job have an immediate predecessor.

Moreover, we will show that W can be written as

$$W = \max_{1 \le i \le N} V_i^+,$$

where the $\{V_i\}$ are i.i.d., independent of N, and solve the stochastic fixed-point equation

$$V \stackrel{\mathcal{D}}{=} \max\left\{\chi - \tau,\ \max_{1 \le i \le N} (\chi - \tau + V_i)\right\} \tag{6}$$

with the $\{V_i\}$ i.i.d. copies of V, independent of (N, χ, τ). In particular, if N and χ are independent, W itself solves the stochastic fixed point equation

$$W \stackrel{\mathcal{D}}{=} \max\left\{0,\ \max_{1 \le i \le N} (\chi_i - \tau_i + W_i)\right\}, \tag{7}$$

where $\{W_i\}$ are i.i.d. copies of W, independent of $(N, \chi_1, \tau_1, \chi_2, \tau_2, \dots)$. Both (6) and (7) are known in the literature [20, 10, 18] as high-order Lindley equations, but their reflection points ($\chi - \tau$ in (6) and 0 in (7)) differ. Also, while W is always nonnegative, V may not be.

It turns out that (6) and (7) may both have multiple solutions [10], unlike the standard Lindley equation for $N \equiv 1$. It is the structure of (5) that will allow us to identify the correct one. As we will see in the following sections, the appropriate solution is the so-called endogenous one, which is also the minimal one in the usual stochastic order sense.

To identify the rate of τ recall that in the time reversed setting, for the system with n servers, we can think of independent Poisson processes each generating jobs with a tag of the form $(s_1, s_2, \dots, s_k)$, for $1 \le k \le n$, and $s_i \in \{1, 2, \dots, n\}$ for all i. There are a total of $n!/(n-k)!$ tags of length k, and each of those has an associated Poisson process with rate

$$\lambda_k = \frac{f_n(k)\lambda n}{n!/(n-k)!}.$$

Moreover, a piece of a job requiring service at server s_i can have as an immediate predecessor any job of any size requiring service at server s_i. In particular, there are a total of $\binom{n-1}{k-1}k!$ possible predecessors of size k, and therefore, the interarrival time between the piece of the job requiring service from server s_i and its predecessor is exponentially distributed with rate

$$\lambda_n^* = \sum_{k=1}^n \lambda_k \binom{n-1}{k-1} k! = \sum_{k=1}^n \frac{f_n(k)\lambda n}{\binom{n}{k}} \binom{n-1}{k-1} = \lambda \sum_{k=1}^n k f_n(k). \tag{8}$$

It follows that, assuming f_n is uniformly integrable, τ must be exponentially distributed with rate

$$\lambda^* = \lambda \sum_{k=1}^\infty k f(k) = \lambda E[N].$$

The limiting behavior of the solutions to equations equations (6) and (7) is given in Section 3.2.

3.1. Main result Before we formulate the main result of this paper it is convenient to specify the conditions we need to impose on f_n, λ, and G_k. Recall that the service requirements of a job of size k, $(\xi^{(1)}, \dots \xi^{(k)})$, are allowed to be arbitrarily dependent, and may also depend on k. In this notation, G_k is the conditional distribution of the service requirement of a randomly chosen piece, i.e.,

$$G_k(x) = P(\hat{\chi} \leq x | \hat{N} = k) = \frac{1}{k} \sum_{i=1}^{k} P(\xi^{(i)} \leq x | \hat{N} = k).$$

Assumption 1. *Suppose that f_n is a probability mass function on $\{1, 2, \dots, n\}$, G_k is a distribution on $\mathbb{R}_+$ for each $k \in \mathbb{N}_+$, and $\lambda > 0$.*

i) Suppose there exists a probability mass function f on $\mathbb{N}_+$, having finite mean, such that $f_n \Rightarrow f$ as $n \to \infty$ and f_n is uniformly integrable.

ii) Suppose there exists $\beta > 0$ such that

$$E\left[N e^{\beta(\chi - \tau)}\right] = \frac{\lambda^*}{\lambda^* + \beta} E\left[N e^{\beta \chi}\right] < 1,$$

where N is distributed according to f, χ given $N = k$ is distributed according to G_k, and τ is exponentially distributed with rate $\lambda^* = \lambda E[N]$, and is independent of (N, χ).

iii) For the same $\beta > 0$ in (ii), assume

$$\lim_{n\to\infty} E[\hat{N} e^{\beta \hat{\chi}}] = \lim_{n\to\infty} \sum_{k=1}^{n} f_n(k) \int_0^\infty e^{\beta x} G_k(dx) = E[N e^{\beta \chi}].$$

To give two examples of distributions for which Assumption 1 is satisfied, let N be distributed according to f and consider

$$f_n(k) = P(\min\{N, n\} = k) \qquad \text{or} \qquad f_n(k) = P(N = k | N \leq n).$$

In both cases, provided $E[N] < \infty$, we have that f_n is uniformly integrable, since the monotone convergence theorem gives $E[\min\{N, n\}] \to E[N]$ and $E[N | N \leq n] = E[N 1(N \leq n)]/P(N \leq n) \to E[N]$.

We are now ready to formulate the main theorem.

Theorem 1. *Let $W_\emptyset^{(n,t_0)}$ denote the waiting time, excluding service, of the tagged job (the first job to arrive after time zero) when we start the system empty at time $t_0 < 0$ and the network consists of n servers. Suppose that*

$$E\left[\hat{N} e^{\beta(\hat{\chi} - \hat{\tau})}\right] < 1 \tag{9}$$

for some $\beta > 0$, where $\hat{N}$ has distribution f_n, $\hat{\chi}$ given $\hat{N} = k$ has distribution G_k and $\hat{\tau}$ is exponentially distributed with rate λ_n^* and is independent of $(\hat{N}, \hat{\chi})$. Then, for any fixed number of servers n,

$$\lim_{t_0 \to -\infty} W_\emptyset^{(n,t_0)} = W^{(n)} \quad a.s.$$

for some finite random variable $W^{(n)}$. Moreover, provided Assumption 1 is satisfied,

$$W^{(n)} \Rightarrow W,$$

as $n \to \infty$, where W can be written as

$$W = \max_{1 \leq i \leq N} V_i^+,$$

where the $\{V_i\}$ are i.i.d., independent of N, and are distributed according to the endogenous solution to the stochastic fixed-point equation

$$V \stackrel{\mathcal{D}}{=} \max\left\{\chi - \tau, \max_{1\le i\le N} (\chi - \tau + V_i)\right\},$$

with the $\{V_i\}$ i.i.d. copies of V, independent of (N, χ, τ). In particular, if N and χ are independent, W itself has the distribution of the endogenous solution to the stochastic fixed point equation

$$W \stackrel{\mathcal{D}}{=} \max\left\{0, \max_{1\le i\le N} (\chi_i - \tau_i + W_i)\right\},$$

where $\{W_i\}$ are i.i.d. copies of W, independent of $(N, \chi_1, \tau_1, \chi_2, \tau_2, \dots)$. Furthermore, we have that for any $p > 0$,

$$E\left[(W^{(n)})^p\right] \to E\left[W^p\right] < \infty, \qquad n \to \infty.$$

The key idea for the proof of the stability result is to couple the predecessor graph with a weighted branching tree (marked Galton-Watson process) [30, 16] and show that the waiting time of the tagged job is dominated by the maximum of the random walks along all paths of the tree. The identification of the limit with the endogenous solutions to the high-order Lindley equations (6) or (7) will follow from a coupling argument between the predecessor graph and a weighted branching tree, in which we will show that with high probability the jobs present in the system at the time the tagged job arrives will have no merged paths. This critical timescale at which the first merging of paths is observed also explains why the dependence among the different service requirements of a job plays no role in the limiting distribution, since a typical job will only "see" one fragment from each of its predecessors still present in the system when it arrives.

3.2. Analyzing the limit: Generalized Cramér-Lundberg approximation As stated in Theorem 1, the stationary waiting time in the system with n servers converges to a limit that can be written in terms of the endogenous solution to a high order Lindley equation (either (6) or (7)), which receives its name from the observation that it can be explicitly constructed on a marked Galton-Watson process. For completeness, we now briefly describe the construction of a weighted branching process, which is more general than the setup considered in this paper.

Let $(Q, N, C_1, C_2, \dots)$ be a vector with $N \in \mathbb{N} \cup \{\infty\}$, and $Q, \{C_i\}$ real-valued; the interpretation of Q and the $\{C_i\}$ depends on the application. Given a sequence of i.i.d. vectors $\left\{(Q_{\mathbf{i}}, N_{\mathbf{i}}, C_{(\mathbf{i},1)}, C_{(\mathbf{i},2)}, \dots)\right\}_{\mathbf{i}\in U}$ having the same distribution as the generic branching vector $(Q, N, C_1, C_2, \dots)$, we use the random variables $\{N_{\mathbf{i}}\}_{\mathbf{i}\in U}$ to determine the structure of a tree as follows. Let $A_0 = \{\emptyset\}$ and

$$A_r = \{(\mathbf{i}, i_r) \in U : \mathbf{i} \in A_{r-1}, 1 \le i_r \le N_{\mathbf{i}}\}, \quad r \ge 1, \tag{10}$$

be the set of individuals in the rth generation. Next, assign to each node $\mathbf{i}$ in the tree a weight $\Pi_{\mathbf{i}}$ according to the recursion

$$\Pi_\emptyset = 1, \qquad \Pi_{(\mathbf{i},j)} = C_{(\mathbf{i},j)} \Pi_{\mathbf{i}}.$$

Each weight $\Pi_{\mathbf{i}}$ is also usually multiplied by its corresponding value $Q_{\mathbf{i}}$ to construct solutions to non-homogeneous stochastic fixed-point equations.

In the general formulation, the vector $(Q, N, C_1, C_2, \dots)$ is allowed to be arbitrarily dependent. In the special case appearing in this paper we either have $(Q, N, C_1, C_2, \dots) = (e^{\chi-\tau}, N, e^{\chi-\tau}, e^{\chi-\tau}, \dots)$ in the general (N, χ) case, or $(Q, N, C_1, C_2, \dots) = (0, N, e^{\chi_i-\tau_i}, e^{\chi_i-\tau_i}, \dots)$ in the case that N and χ

are independent. In the latter, the $\{e^{\chi_i-\tau_i}\}_{i\geq 1}$ are i.i.d. and independent of N. For more details we refer the reader to [30, 10, 16, 17].

To make the connection between the high-order Lindley equation (6) (or (7)) and the main result in [18], let $R=e^V$, $R_i=e^{V_i}$, $Q=e^{\chi-\tau}$, and $C_i=e^{\chi-\tau}$ for all $i\geq 1$ (or $R=e^W$, $R_i=e^{W_i}$, $Q\equiv 1$, and $C_i=e^{\chi_i-\tau_i}$ in the independent case) to obtain

$$R \stackrel{\mathcal{D}}{=} Q \vee \left(\bigvee_{i=1}^{N} C_i R_i \right), \tag{11}$$

where $x\vee y$ denotes the maximum of x and y. We refer to (11) with a generic branching vector of the form $(Q,N,C_1,C_2,\ldots)$ with the $\{C_i\}$ nonnegative, and the $\{R_i\}$ i.i.d. copies of R independent of $(Q,N,C_1,C_2,\ldots)$, as the *branching maximum equation*.

It is easy to verify, as was done in [18], that the random variable

$$R \triangleq \bigvee_{r=0}^{\infty} \bigvee_{\mathbf{j}\in A_r} \Pi_{\mathbf{j}} Q_{\mathbf{j}}$$

is a solution to (11), known in the literature as the endogenous solution [1, 10]. Moreover, when $Q\geq 0$, the endogenous solution is also the minimal one in the usual stochastic order sense (see [10] and also the survey paper [1] for additional references and a wide variety of max-plus equations). Taking logarithms on both sides of (11) (with $Q\equiv 1$), we obtain that the endogenous solution to (6) is given by

$$V \triangleq \bigvee_{r=1}^{\infty} \bigvee_{(1,\mathbf{j})\in A_r} S_{(1,\mathbf{j})}, \tag{12}$$

and the one to (7) by

$$W \triangleq \bigvee_{r=0}^{\infty} \bigvee_{\mathbf{j}\in A_r} S_{\mathbf{j}}, \tag{13}$$

where $S_\emptyset = 0$, $S_{\mathbf{j}} = \log \Pi_{\mathbf{j}} = X_{\mathbf{j}|1} + X_{\mathbf{j}|2} + \cdots + X_{\mathbf{j}}$ for $\mathbf{j}\neq\emptyset$, and $X_{\mathbf{i}} = \chi_{\mathbf{i}} - \tau_{\mathbf{i}}$. Furthermore, it was shown in [18] (see Lemma 3.1) that these endogenous solutions are finite almost surely provided

$$E[Ne^{\beta X}] < 1$$

for some $\beta>0$, which we will refer to as the stability condition (see also Theorem 6 in [10]). Note that with respect to the queueing model in this paper, the stability condition implies the usual "load condition", i.e., arrival rate divided by service rate strictly smaller than one, which in this case would be $\lambda E[N\chi] = E[N\chi]/E[\tau] < 1$; the two are equivalent when $E[N]=1$.

Remark 2. The stability condition guarantees that V and W, as defined by (12) and (13), respectively, are finite almost surely. Moreover, by Theorem 4 in [10], the existence of $\beta>0$ such that $E[Ne^{\beta X}]\leq 1$ is the corresponding necessary condition. However, we do not consider in this paper the boundary condition where $E[Ne^{\theta X}]=1$ for some $\theta>0$ but $E[Ne^{\beta X}]>1$ for all $\beta\neq\theta$.

By rewriting W as

$$W = \max\left\{0, \max_{j\in A_1} X_j, \max_{\mathbf{j}\in A_2}\left(X_{\mathbf{j}|1}+X_{\mathbf{j}}\right), \max_{\mathbf{j}\in A_3}\left(X_{\mathbf{j}|1}+X_{\mathbf{j}|2}+X_{\mathbf{j}}\right), \ldots\right\},$$

the similarities with (5) become apparent. To give some additional intuition as to why (13) is the appropriate solution, it is helpful to recall the $N\equiv 1$ case, where Lindley's equation is known to have a unique solution whenever $E[X_1] < 0$. Moreover, as mentioned in the introduction, this solution can be expressed in terms of the supremum of the random walk $S_k = X_1 + \cdots + X_k$, $S_0 = 0$. A standard proof of this relation consists in iterating the recursion

$$W_{n+1} = \max\{0, X_n + W_n\}, \qquad W_0 = 0,$$

to obtain

$$W_{n+1} = \max\{0, X_n, X_n + X_{n-1}, \ldots, X_n + X_{n-1} + \cdots + X_1\} \stackrel{\mathcal{D}}{=} \max_{0 \leq k \leq n} S_k.$$

It follows by taking the limit as $n \to \infty$ on both sides that the stationary waiting time in the FCFS GI/GI/1 queue satisfies

$$W \stackrel{\mathcal{D}}{=} \max_{k \geq 0} S_k.$$

It is then to be expected that the asymptotic analysis of the waiting time in the single-server queue can also be generalized to the branching setting. This is indeed the case, as was recently shown in [18]. There, for the endogenous solution to the general branching maximum recursion (11), it was shown that if there exists $\theta > 0$ such that the root condition $E\left[\sum_{i=1}^N C_i^\theta\right] = 1$ and the derivative condition $0 < E\left[\sum_{i=1}^N C_i^\theta \log C_i\right] < \infty$ are satisfied, then

$$P(R > x) \sim H x^{-\theta}, \qquad x \to \infty, \tag{14}$$

for some constant $H > 0$. Note that for the high-order Lindley's equation (6), this condition translates into the existence of a root $\theta > 0$ such that $E[N e^{\theta X}] = 1$. The power-law asymptotics of R are a consequence of the Implicit Renewal Theorem on Trees from [16, 17], which constitutes a powerful tool for the analysis of many different types of branching recursions, e.g., the maximum recursion [6, 18], the linear recursion or smoothing transform [2, 4, 5, 16, 17, 3], the discounted tree sum [1], etc. This theorem is in turn a generalization of the Implicit Renewal Theorem of [13] for non-branching recursions, which can be used to analyze the random coefficient autoregressive process of order one and the reflected random walk, among others. The name "implicit" refers to the fact that the Renewal Theorem is applied to a random variable R (e.g., the solution to a stochastic fixed-point equation) without having knowledge of its distribution, which in turn leads to the resulting constant H in the asymptotics to be implicitly defined in terms of R itself.

We conclude this section with the theorem describing the asymptotic behavior of W.

THEOREM 2. *Let W be given by*

$$W = \bigvee_{r=0}^{\infty} \bigvee_{\mathbf{j} \in A_r} S_{\mathbf{j}},$$

where $X_{\mathbf{j}} = \chi_{\mathbf{i}} - \tau_{\mathbf{i}}$, and the vectors $\{(N_{\mathbf{i}}, \chi_{\mathbf{i}}, \tau_{\mathbf{i}})\}_{\mathbf{i} \in U}$ are i.i.d. copies of (N, χ, τ), with N distributed according to f, χ given $N = k$ distributed according to G_k, and τ exponentially distributed with rate λ^* and independent of (N, χ). Suppose that for some $\theta > 0$, $E[N e^{\theta(\chi - \tau)}] = 1$ and $0 < E\left[e^{\theta(\chi-\tau)}(\chi - \tau)\right] < \infty$. In addition, assume that for some $\epsilon > 0$, $E\left[N^{\theta \vee (1+\epsilon)} e^{\theta(\chi-\tau)}\right] < \infty$. Then,

$$P(W > x) \sim H E[N] e^{-\theta x}, \qquad x \to \infty,$$

where $0 < H < \infty$ is given by

$$H = \frac{E\left[e^{\theta(\chi-\tau)}\left(\bigvee_{i=1}^N e^{\theta V_i^+} - \sum_{i=1}^N e^{\theta V_i}\right)\right]}{\theta E\left[N e^{\theta(\chi-\tau)}(\chi - \tau)\right]},$$

with the $\{V_i\}$ *i.i.d. copies of the endogenous solution* V *to* (6)*, independent of* (N,χ,τ)*. In addition, if* N *is independent of* χ*, we have*

$$HE[N] = \frac{E\left[1 \vee \bigvee_{i=1}^N e^{\theta(\chi_i-\tau_i+W_i)} - \sum_{i=1}^N e^{\theta(\chi_i-\tau_i+W_i)}\right]}{\theta E[N] E\left[e^{\theta(\chi-\tau)}(\chi-\tau)\right]},$$

with the $\{W_i\}$ *i.i.d. copies of* W*, independent of* $(N,\chi_1,\tau_1,\ldots,\chi_N,\tau_N)$*.*

The constant H in the asymptotic tail of W can be computed via simulation, for example, by using the algorithm recently developed in [11, 28], which can be used to generate the $\{V_i\}$ (or $\{W_i\}$) appearing in the expectation. That algorithm uses a bootstrap technique for computing enough iterations of the stochastic fixed-point equation (6) (or (7)) to ensure that the samples it produces have approximately the same distribution as its endogenous solution.

4. Concluding remarks The model presented in this paper captures the complexity of a large database system where users can request any subset of files and block them while they alter their content. Compared to existing results in the database locking literature, this work is the first to incorporate the possibility of requesting an arbitrary number of files, which could also affect the duration of the locking period itself. The techniques used to analyze our model could also potentially yield new ideas on how to analyze more general reader-writer systems. Our model also adds to the existing literature on distributed queueing systems by providing an analytically tractable model that is closely related to popular split-merge networks and certain resource allocation problems.

Appendix. Proofs.

This appendix contains the proofs of Theorem 1 and Theorem 2. To ease the exposition we separate Theorem 1 into two parts, the first one concerning the existence of a stationary waiting time for a fixed number of servers and a fixed arrival rate of jobs, the second one establishing the limiting distribution of the stationary waiting time as the number of servers and the arrival rate of jobs grow to infinity.

We start by providing some intuition for the limiting distribution in Theorem 1. Recall that the waiting time of the tagged job in the system started at time $t_0<0$, $W_\emptyset^{(n,t_0)}$, satisfies

$$W_\emptyset^{(n,t_0)} = \max\left\{0, \max_{j\in\mathcal{B}_{\emptyset,1}} \hat{X}_j, \ldots, \max_{\mathbf{j}\in\mathcal{B}_{\emptyset,\kappa}} \left(\hat{X}_{\mathbf{j}|1} + \hat{X}_{\mathbf{j}|2} + \ldots \hat{X}_{\mathbf{j}}\right)\right\},$$

where $\mathcal{B}_{\emptyset,r}$ is the set of labels at distance r from the tagged job in the predecessor graph and $\kappa = \max\{r\in\mathbb{N}_+ : |\mathcal{B}_{\emptyset,r}|>0\}$. The first part of Theorem 1 will establish that

$$\lim_{t_0\to-\infty} W_\emptyset^{(n,t_0)} = W^{(n)} \triangleq \bigvee_{r=0}^\infty \bigvee_{\mathbf{j}\in\hat{A}_r} \hat{S}_{\mathbf{j}},$$

where $\hat{S}_{\mathbf{j}} = \hat{X}_{\mathbf{j}|1} + \hat{X}_{\mathbf{j}|2} + \cdots + \hat{X}_{\mathbf{j}}$, $\hat{X}_{\mathbf{i}} = \hat{\chi}_{\mathbf{i}} - \hat{\tau}_{\mathbf{i}}$, and $\hat{A}_r$ is the set of labels in the predecessor graph at graph distance r from the tagged job. Note that for finite t_0 we have $\mathcal{B}_{\emptyset,r}\subseteq\hat{A}_r$, since some jobs/fragments may not have any predecessors; however, as $t_0\to-\infty$, all jobs/fragments will. To analyze the distribution of the random variables involved, we look at the queueing system with time running backwards, so that predecessors of the tagged job "arrive" rather than "precede". In this time-reversed setting, multiple paths in the graph merge every time a job that is an immediate predecessor to multiple fragments arrives, and all jobs from that point onwards will have multiple labels. Since different labels can represent the same job, we have that some of the $\{(\hat{N}_{\mathbf{i}}, \hat{\chi}_{\mathbf{i}})\}$ can be repeated, and are therefore not independent. Similarly, the $\{\hat{\tau}_{\mathbf{i}}\}$ correspond to the interarrival times between jobs in the predecessor graph, and are therefore, in general, neither independent of each other nor of the $\{\hat{N}_{\mathbf{i}}\}$. More precisely, the marginal distribution of each of the $\hat{\tau}_{\mathbf{i}}$ is exponential with rate λ_n^*, but conditionally on knowing that $\mathbf{i}$ shares a predecessor with one or more other jobs, its rate changes and all the interarrival times corresponding to edges that merge into the same job become dependent. Finally, the $\{(\hat{N}_{\mathbf{i}}, \hat{\chi}_{\mathbf{i}})\}$ are identically distributed, but become dependent for $\mathbf{i} \neq \mathbf{j}$ when the two labels correspond to the same job.

Despite all the lack of independence described above, we can intuitively derive the limiting distribution by simply ignoring the merging of paths. In other words, the limiting distribution W is constructed on a marked Galton-Watson process defined through a sequence of i.i.d. vectors $\{(N_{\mathbf{i}}, \chi_{\mathbf{i}}, \tau_{\mathbf{i}})\}_{\mathbf{i}\in U}$ having common distribution

$$P(N = k, \chi \in dx, \tau \in dt) = \lambda^* e^{-\lambda^* t} f(k) G_k(dx), \qquad k \in \mathbb{N}_+,\, x, t \in \mathbb{R}_+.$$

To construct the tree we set $A_0 = \{\emptyset\}$ and use the $\{N_{\mathbf{i}}\}$ in the sequence to define $A_r = \{(\mathbf{i}, i_r) : \mathbf{i} \in A_{r-1},\, 1 \le i_r \le N_{\mathbf{i}}\}$ for $r \ge 1$. Next, set $X_{\mathbf{i}} = \chi_{\mathbf{i}} - \tau_{\mathbf{i}}$ and $S_{\mathbf{j}} = X_{\mathbf{j}|1} + X_{\mathbf{j}|2} + \cdots + X_{\mathbf{j}}$. The limit W can then be formally written as

$$W \triangleq \bigvee_{r=0}^{\infty} \bigvee_{\mathbf{j}\in A_r} S_{\mathbf{j}}.$$

Note that if $N_{\mathbf{i}}$ is independent of $X_{\mathbf{i}}$, then W satisfies

$$W = 0 \vee \bigvee_{r=1}^{\infty} \bigvee_{\mathbf{j}\in A_r} S_{\mathbf{j}} = 0 \vee \bigvee_{i=1}^{N_\emptyset} \left(X_i + \left(\bigvee_{r=1}^{\infty} \bigvee_{(i,\mathbf{j})\in A_r} (S_{(i,\mathbf{j})} - X_i) \right) \right) \stackrel{\mathcal{D}}{=} 0 \vee \bigvee_{i=1}^{N_\emptyset} (X_i + W_i)\,, \tag{15}$$

where the $\{W_i\}_{i\ge 1}$ are i.i.d. copies of W, independent of $(N_\emptyset, X_1, X_2, \dots)$. However, if $N_{\mathbf{i}}$ and $\chi_{\mathbf{i}}$ are dependent, then X_i and W_i in (15) fail to be independent and we need to rearrange the expression for W. To do this, define

$$V_i^{(k)} \triangleq \bigvee_{r=1}^{k} \bigvee_{(i,\mathbf{j})\in A_r} S_{(i,\mathbf{j})}, \qquad k \ge 1,$$

and note that

$$W = 0 \vee \bigvee_{r=1}^{\infty} \bigvee_{\mathbf{j}\in A_r} S_{\mathbf{j}} = 0 \vee \bigvee_{i=1}^{N_\emptyset} \bigvee_{r=1}^{\infty} \bigvee_{(i,\mathbf{j})\in A_r} S_{(i,\mathbf{j})} = 0 \vee \bigvee_{i=1}^{N_\emptyset} V_i^{(\infty)}.$$

By making this shift in the indexes we obtain that

$$\begin{aligned} V_i^{(k+1)} &= \max\left\{ X_i,\, \bigvee_{r=2}^{k+1} \bigvee_{(i,\mathbf{j})\in A_r} S_{(i,\mathbf{j})} \right\} \\ &= \max\left\{ X_i,\, X_i + \bigvee_{l=1}^{N_i} \bigvee_{r=2}^{k+1} \bigvee_{(i,l,\mathbf{j})\in A_r} (S_{(i,l,\mathbf{j})} - X_i) \right\} \\ &\triangleq \max\left\{ X_i,\, X_i + \bigvee_{l=1}^{N_i} V_{(i,l)}^{(k)} \right\}, \end{aligned} \tag{16}$$

where the $\{V_{(i,l)}^{(k)}\}_{l\ge 1}$ are i.i.d. copies of $V_i^{(k)}$, independent of (N_i, X_i). Therefore, by monotonicity, the $\{V_i^{(\infty)}\}_{i\ge 1}$ are i.i.d. copies of the attracting endogenous solution to the stochastic fixed-point equation:

$$V \stackrel{\mathcal{D}}{=} \max\left\{ X,\, X + \max_{1\le i\le N} V_i \right\} = X + \max_{1\le i\le N} V_i^+,$$

with the $\{V_i\}_{i\geq 1}$ i.i.d. copies of V, independent of (N,X). From the weighted branching processes literature (see, Lemma 3.1 in [18] with $(Q,N,C_1,C_2,\dots)=(e^X,N,e^X,e^X,\dots)$, or Theorem 6 in [10]), we have that $V<\infty$ a.s. provided there exists $\beta>0$ such that $E[Ne^{\beta X}]=E[Ne^{\beta\chi}]E[e^{-\beta\tau}]<1$. The above discussion is summarized in the following lemma.

LEMMA 1. *Suppose there exists $\beta>0$ such that $E[Ne^{\beta\chi}]E[e^{-\beta\tau}]<1$. Then W can be written as*

$$W=\max_{1\leq i\leq N} V_i^+,$$

where the $\{V_i\}_{i\geq 1}$ are i.i.d. copies of V, independent of N, and V is the attracting endogenous solution to

$$V \stackrel{\mathcal{D}}{=} \chi-\tau+\max_{1\leq i\leq N} V_i^+,$$

with the $\{V_i\}_{i\geq 1}$ i.i.d. copies of V, independent of (N,χ,τ). Moreover, if N is independent of χ, then W is the attracting endogenous solution to

$$W \stackrel{\mathcal{D}}{=} \max_{1\leq i\leq N} (\chi_i-\tau_i+W_i)^+,$$

with the $\{W_i\}_{i\geq 1}$ i.i.d. copies of W, independent of $(N,\chi_1-\tau_1,\chi_2-\tau_2,\dots)$.

Now that we have characterized the limit W, we proceed with the proof of Theorem 1. The key idea is to first approximate $W^{(n)}$ with $\bigvee_{r=0}^k \bigvee_{\mathbf{j}\in\hat{A}_r} \hat{S}_{\mathbf{j}}$ for a sufficiently large k, and then argue that the latter coincides with the first k generations of a marked Galton-Watson process. To make this precise, let $\{(\tilde{N}_{\mathbf{i}},\tilde{\chi}_{\mathbf{i}},\tilde{\tau}_{\mathbf{i}})\}_{\mathbf{i}\in U}$ be i.i.d. with common distribution

$$P(\tilde{N}=k,\tilde{\chi}\in dx,\tilde{\tau}\in dt)=\lambda_n^* e^{-\lambda_n^* t} f_n(k) G_k(dx), \qquad k\in\mathbb{N}_+,\, x,t\in\mathbb{R}_+,$$

and use them to define a branching process by setting $\tilde{A}_0=\{\emptyset\}$ and $\tilde{A}_r=\{(\mathbf{i},i_r):\mathbf{i}\in\tilde{A}_{r-1},\, 1\leq i_r\leq \tilde{N}_{\mathbf{i}}\}$ for $r\geq 1$. Next, define

$$\tilde{V}_i^{(k)}=\bigvee_{r=1}^k \bigvee_{(i,\mathbf{j})\in\tilde{A}_r} \tilde{S}_{(i,\mathbf{j})}, \qquad k\geq 1,$$

where $\tilde{S}_{\mathbf{i}}=\tilde{X}_{\mathbf{j}|1}+\tilde{X}_{\mathbf{j}|2}+\cdots+\tilde{X}_{\mathbf{j}}$ and $\tilde{X}_{\mathbf{i}}=\tilde{\chi}_{\mathbf{i}}-\tilde{\tau}_{\mathbf{i}}$. We will show that for sufficiently large k,

$$W^{(n)}\approx \bigvee_{r=0}^k \bigvee_{\mathbf{j}\in\hat{A}_r} \hat{S}_{\mathbf{j}} \approx \bigvee_{r=0}^k \bigvee_{\mathbf{j}\in\tilde{A}_r} \tilde{S}_{\mathbf{j}} \approx W.$$

We begin the proof of Theorem 1 with the stability part of the theorem. Define $\hat{U}_0=\tilde{U}_0=U_0=0$, and for $r\geq 1$,

$$\hat{U}_r=\bigvee_{\mathbf{i}\in\hat{A}_r}\hat{S}_{\mathbf{i}}, \qquad \tilde{U}_r=\bigvee_{\mathbf{i}\in\tilde{A}_r}\tilde{S}_{\mathbf{i}} \qquad \text{and} \qquad U_r=\bigvee_{\mathbf{i}\in A_r} S_{\mathbf{i}}.$$

The following result will provide upper bounds for the moments of $\hat{U}_r,\tilde{U}_r$ and U_r later on.

LEMMA 2. *Let $\hat{X}=\hat{\chi}-\hat{\tau}$, $\tilde{X}=\tilde{\chi}-\tilde{\tau}$, and $X=\chi-\tau$. Then, for any $\beta>0$, and $r\geq 1$, we have*

$$E\left[\sum_{\mathbf{j}\in\hat{A}_r} e^{\beta\hat{S}_{\mathbf{j}}}\right]=E\left[\sum_{\mathbf{j}\in\tilde{A}_r} e^{\beta\tilde{S}_{\mathbf{j}}}\right]=E[\tilde{N}]E\left[\sum_{(1,\mathbf{j})\in\tilde{A}_r} e^{\beta\tilde{S}_{(1,\mathbf{j})}}\right]=E[\tilde{N}]E[e^{\beta\tilde{X}}]\left(E[\tilde{N}e^{\beta\tilde{X}}]\right)^{r-1},$$

$$\text{and} \qquad E\left[\sum_{\mathbf{j}\in A_r} e^{\beta S_{\mathbf{j}}}\right]=E[N]E\left[\sum_{(1,\mathbf{j})\in A_r} e^{\beta S_{(1,\mathbf{j})}}\right]=E[N]E[e^{\beta X}]\left(E[Ne^{\beta X}]\right)^{r-1}.$$

Proof. We will start by computing the expectation $E\left[\sum_{\mathbf{j}\in\hat{A}_r} e^{\beta\hat{S}_\mathbf{j}}\right]$. To do this, we look at the queueing process with time flowing backwards. Note that in this time-reversed setting, predecessors "arrive" and connect to vertices already present in the predecessor graph. Define $\mathcal{G}_n \triangleq \mathcal{G}_n(-\infty)$. Recall from the observations made at the beginning of this section that the $\{\hat{\tau}_\mathbf{i}\}_{\mathbf{i}\in\mathcal{G}_n}$ are neither i.i.d. nor independent of the $\{\hat{N}_\mathbf{i}\}_{\mathbf{i}\in\mathcal{G}_n}$. More precisely, for each piece of a job requiring service at server s_i there are $\binom{n-1}{k-1}k!$ possible immediate predecessors of size k (i.e., jobs that also require service from server s_i). Since the arrival of jobs into the system is assumed to follow a Poisson process, this leads in the reversed time setting to the interarrival times between a fixed piece of a job and its unique immediate predecessor to be exponentially distributed with rate λ_n^*. The problem arises when a job is an immediate predecessor to two or more jobs in $\mathcal{G}_n$, which introduces dependence and changes the rate of the exponentials representing the interarrival times.

More precisely, consider an arrival that is predecessor to two jobs (or two pieces of the same job), j_1 and j_2, and note that there must be two different servers, say s_{i_1} and s_{i_2}, that are required by the arriving job and that are also assigned to jobs j_1 and j_2, respectively. There are only $\binom{n-2}{k-2}k!$ possible jobs of size k requiring service by servers s_{i_1} and s_{i_2}, and therefore, the rate at which such a predecessor arrives is given by

$$\lambda_n^{(2)} = \sum_{k=2}^n \lambda_k \binom{n-2}{k-2} k!.$$

In general, a job that is predecessor to jobs $j_1, j_2, \dots, j_r$ in the graph arrives at a rate

$$\lambda_n^{(r)} = \sum_{k=r}^n \lambda_k \binom{n-r}{k-r} k! \leq \sum_{k=1}^n \lambda_k \binom{n-1}{k-1} k! = \lambda_n^*.$$

As for the lack of independence between the $\{\hat{\tau}_\mathbf{i}\}_{\mathbf{i}\in\mathcal{G}_n}$, note that the interarrival times between fragments of jobs that have a common immediate predecessor are dependent. The sequence $\{\hat{\tau}_\mathbf{i}\}_{\mathbf{i}\in\mathcal{G}_n}$ is also dependent on the $\{\hat{N}_\mathbf{i}\}_{\mathbf{i}\in\mathcal{G}_n}$, since a large number of jobs awaiting for a predecessor to arrive increases the probability of an arriving job being predecessor to two or more pieces at a time. Hence, in order to compute $E\left[\sum_{\mathbf{j}\in\hat{A}_r} e^{\beta\hat{S}_\mathbf{j}}\right]$ we first eliminate most of the dependence by rewriting it as

$$E\left[\sum_{\mathbf{j}\in\hat{A}_r} e^{\beta\hat{S}_\mathbf{j}}\right] = \sum_{\mathbf{j}\in\mathbb{N}_+^r} E\left[e^{\beta\hat{S}_\mathbf{j}} 1(\mathbf{j}\in\hat{A}_r)\right],$$

and noting that

$$e^{\beta\hat{S}_\mathbf{j}} 1(\mathbf{j}\in\hat{A}_r) = e^{-\beta\sum_{l=1}^r \hat{\tau}_{\mathbf{j}|l}} \cdot e^{\beta\sum_{i=1}^r \hat{\chi}_{\mathbf{j}|i}} \cdot \prod_{k=0}^{r-1} 1(j_{k+1} \leq \hat{N}_{\mathbf{j}|k}).$$

Since all labels along a path correspond to different jobs, then the vectors $\{(\hat{N}_{\mathbf{j}|k}, \hat{\chi}_{\mathbf{j}|k}))\}_{k=0}^{r-1}$ are i.i.d. copies of $(\tilde{N}, \tilde{\chi})$. Also, along a single path, the $\{\hat{\tau}_{\mathbf{j}|k}\}_{k=1}^r$ are i.i.d. copies of $\tilde{\tau}$, since $\hat{\tau}_{\mathbf{j}|k}$ can be identified with the interarrival time between a specific fragment of the job with label $(\mathbf{j}|k-1)$ and its immediate predecessor. In other words, the dependence occurs only among different paths. To compute the required expectation, let $\hat{T}_{\mathbf{j}|r-1}$ denote the arrival time of the job whose label is $(\mathbf{j}|r-1)$, and condition on the history of the arrival process starting from time $T_1(=\hat{T}_\emptyset)$ until $\hat{T}_{\mathbf{j}|r-1}$ (recall that time flows in reverse). At time $\hat{T}_{\mathbf{j}|r-1}$, $\hat{N}_\emptyset$ as well as all of the $\{(\hat{N}_{\mathbf{j}|k}, \hat{\chi}_{\mathbf{j}|k}, \hat{\tau}_{\mathbf{j}|k})\}_{k=1}^{r-1}$ have been revealed, and the j_r-th fragment of job $(\mathbf{j}|r-1)$ is awaiting an immediate predecessor; note that it is waiting for the first job to require service at its assigned server, and it does not matter whether this immediate predecessor merges paths or not. We then have that

$$\begin{aligned} E\left[e^{\beta \hat{S}_{\mathbf{j}}} 1(\mathbf{j} \in \hat{A}_r)\right] &= E\left[e^{-\beta \sum_{l=1}^{r-1} \hat{\tau}_{\mathbf{j}|l}} \cdot e^{\beta \sum_{i=1}^{r-1} \hat{\chi}_{\mathbf{j}|i}} \cdot \prod_{k=0}^{r-1} 1(j_{k+1} \leq \hat{N}_{\mathbf{j}|k}) E\left[e^{\beta(\hat{\chi}_{\mathbf{j}} - \hat{\tau}_{\mathbf{j}})}\right]\right] \\ &= E\left[e^{\beta \tilde{X}}\right] E\left[e^{-\beta \sum_{l=1}^{r-1} \hat{\tau}_{\mathbf{j}|l}} \cdot e^{\beta \sum_{i=1}^{r-1} \hat{\chi}_{\mathbf{j}|i}} \cdot \prod_{k=0}^{r-1} 1(j_{k+1} \leq \hat{N}_{\mathbf{j}|k})\right]. \end{aligned}$$

Now condition on the history of the arrival process up to time $\hat{T}_{\mathbf{j}|r-2}$ to obtain

$$\begin{aligned} & E\left[e^{-\beta \sum_{l=1}^{r-1} \hat{\tau}_{\mathbf{j}|l}} \cdot e^{\beta \sum_{i=1}^{r-1} \hat{\chi}_{\mathbf{j}|i}} \cdot \prod_{k=0}^{r-1} 1(j_{k+1} \leq \hat{N}_{\mathbf{j}|k})\right] \\ &= E\left[e^{-\beta \sum_{l=1}^{r-2} \hat{\tau}_{\mathbf{j}|l}} \cdot e^{\beta \sum_{i=1}^{r-2} \hat{\chi}_{\mathbf{j}|i}} \cdot \prod_{k=0}^{r-2} 1(j_{k+1} \leq \hat{N}_{\mathbf{j}|k}) E\left[e^{\beta(\hat{\chi}_{\mathbf{j}|r-1} - \hat{\tau}_{\mathbf{j}|r-1})} 1(j_r \leq \hat{N}_{\mathbf{j}|r-1})\right]\right] \\ &= E\left[e^{\beta \tilde{X}} 1(j_r \leq \tilde{N})\right] E\left[e^{-\beta \sum_{l=1}^{r-2} \hat{\tau}_{\mathbf{j}|l}} \cdot e^{\beta \sum_{i=1}^{r-2} \hat{\chi}_{\mathbf{j}|i}} \cdot \prod_{k=0}^{r-2} 1(j_{k+1} \leq \hat{N}_{\mathbf{j}|k})\right] \\ &= \prod_{k=1}^{r-1} E\left[e^{\beta \tilde{X}} 1(j_k \leq \tilde{N})\right] \cdot E\left[1(j_1 \leq \tilde{N}_\emptyset)\right] \qquad \text{(after iterating } r-1 \text{ times)}. \end{aligned}$$

It follows that

$$\begin{aligned} \sum_{\mathbf{j} \in \mathbb{N}_+^r} E\left[e^{\beta \hat{S}_{\mathbf{j}}} 1(\mathbf{j} \in \hat{A}_r)\right] &= E\left[e^{\beta \tilde{X}}\right] \sum_{\mathbf{j} \in \mathbb{N}_+^r} \prod_{k=1}^{r-1} E\left[e^{\beta \tilde{X}} 1(j_k \leq \tilde{N})\right] \cdot E\left[1(j_1 \leq \tilde{N})\right] \\ &= E[\tilde{N}] \left(E\left[\tilde{N} e^{\beta \tilde{X}}\right]\right)^{r-1} E\left[e^{\beta \tilde{X}}\right]. \end{aligned}$$

Now note that the same computation also gives

$$E\left[\sum_{\mathbf{j} \in \tilde{A}_r} e^{\beta \tilde{S}_{\mathbf{j}}}\right] = \sum_{\mathbf{j} \in \mathbb{N}_+^r} E\left[e^{\beta \tilde{S}_{\mathbf{j}}} 1(\mathbf{j} \in \tilde{A}_r)\right] = E[\tilde{N}] \left(E\left[\tilde{N} e^{\beta \tilde{X}}\right]\right)^{r-1} E\left[e^{\beta \tilde{X}}\right].$$

The proof for $E\left[\sum_{\mathbf{j} \in A_r} e^{\beta S_{\mathbf{j}}}\right]$ is obtained by simply dropping the ~ from all the random variables. □

We are now ready to prove the stability part of Theorem 1.

Proof of Theorem 1 **(Stability)**. Note that since $(\hat{N}, \hat{\chi}, \hat{\tau}) \stackrel{\mathcal{D}}{=} (\tilde{N}, \tilde{\chi}, \tilde{\tau})$, we have $E\left[\hat{N} e^{\beta \hat{X}}\right] = E\left[\tilde{N} e^{\beta \tilde{X}}\right]$. Hence, we need to show that provided $E[\tilde{N} e^{\beta \tilde{X}}] < 1$ the limit $\lim_{t_0 \to -\infty} W_\emptyset^{(n,t_0)}$ exists and is finite a.s. To this end, recall that

$$W_\emptyset^{(n,t_0)} = \bigvee_{r=0}^{\kappa} \bigvee_{\mathbf{j} \in \mathcal{B}_{\emptyset,r}} \left(\hat{X}_{\mathbf{j}|1} + \hat{X}_{\mathbf{j}|2} + \cdots + \hat{X}_{\mathbf{j}}\right), \tag{17}$$

and note that as $t_0 \to -\infty$ we have that $\kappa \to \infty$ a.s. and $\mathcal{B}_{\emptyset,r} \uparrow \hat{A}_r$. It follows by monotone convergence that

$$\lim_{t_0 \to -\infty} W_\emptyset^{(n,t_0)} = \bigvee_{r=0}^{\infty} \bigvee_{\mathbf{j} \in \hat{A}_r} \hat{S}_{\mathbf{j}} = \bigvee_{r=0}^{\infty} \hat{U}_r = W^{(n)} \qquad \text{a.s.}$$

Therefore, it only remains to verify that $W^{(n)} < \infty$ a.s.

To this end, note that it suffices to show that

$$P(\hat{U}_r > 0 \text{ i.o.}) = 0.$$

This in turn will follow from the Borel-Cantelli Lemma once we show that for $r \geq 1$,

$$P(\hat{U}_r > 0) \leq H c^{r-1} \tag{18}$$

for some constants $0 < H < \infty$ and $0 < c < 1$. To see that this is the case, use Markov's inequality to obtain that

$$P\left(\hat{U}_r > 0\right) = P\left(e^{\beta \hat{U}_r} > 1\right) = P\left(\max_{\mathbf{j} \in \hat{A}_r} e^{\beta \hat{S}_{\mathbf{j}}} > 1\right) \leq E\left[\bigvee_{\mathbf{j} \in \hat{A}_r} e^{\beta \hat{S}_{\mathbf{j}}}\right] \leq E\left[\sum_{\mathbf{j} \in \hat{A}_r} e^{\beta \hat{S}_{\mathbf{j}}}\right].$$

Now use Lemma 2 to obtain that

$$E\left[\sum_{\mathbf{j} \in \hat{A}_r} e^{\beta \hat{S}_{\mathbf{j}}}\right] = E[\tilde{N}] E[e^{\beta \tilde{X}}] \left(E\left[\tilde{N} e^{\beta \tilde{X}}\right]\right)^{r-1}.$$

Setting $H = E[\tilde{N}] E[e^{\beta \tilde{X}}]$ and $c = E\left[\tilde{N} e^{\beta \tilde{X}}\right] < 1$ completes the proof. □

For the second part of the main theorem we will prove that

$$W^{(n)} \xrightarrow{d_p} W \qquad \text{as } n \to \infty \tag{19}$$

for any $p \geq 1$, where d_p denotes the Wasserstein distance of order p (see, e.g., [33], Chapter 6). This is equivalent to convergence in distribution plus convergence of all the moments of order up to p (see Theorem 6.8 in [33]).

We will do this in three main steps. First, we will show that if μ_n is the probability measure of $W^{(n)}$ and $\hat{\nu}_k$ is the probability measure of $\bigvee_{r=0}^k \hat{U}_r$, then

$$d_p(\mu_n, \hat{\nu}_{r_n}) \to 0, \qquad n \to \infty, \tag{20}$$

for any $r_n \to \infty$. Next, will show that if $\tilde{\nu}_k$ is the probability measure of $\bigvee_{r=0}^k \tilde{U}_r$ and μ is the probability measure of W then

$$d_p(\tilde{\nu}_{r_n}, \mu) \to 0, \qquad n \to \infty, \tag{21}$$

for any $r_n \to \infty$. Finally, we will use a coupling argument between $\hat{\nu}_k$ and $\tilde{\nu}_k$ to prove that for a well-chosen $r_n \to \infty$,

$$d_p(\hat{\nu}_{r_n}, \tilde{\nu}_{r_n}) \to 0, \qquad n \to \infty. \tag{22}$$

The triangle inequality combined with (20), (21), and (22) yields (19).

We start with a preliminary lemma regarding the moments of $\hat{U}_r$, $\tilde{U}_r$ and U_r, as well as those of $\tilde{V}_1^{(k)}$ and $V_1^{(k)}$.

Lemma 3. *For any $p \geq 1$, $\beta > 0$ and $r \geq 1$,*

$$E\left[(\hat{U}_r^+)^p\right] \leq C_{\beta,p} E[\tilde{N}] \left(E[\tilde{N} e^{\beta \tilde{X}}]\right)^r,$$
$$E\left[(\tilde{U}_r^+)^p\right] \leq C_{\beta,p} E[\tilde{N}] \left(E[\tilde{N} e^{\beta \tilde{X}}]\right)^r,$$
$$\text{and} \qquad E\left[(U_r^+)^p\right] \leq C_{\beta,p} E[N] \left(E[N e^{\beta X}]\right)^r,$$

where $C_{\beta,p} = pE[\mathcal{E}^{p-1}]/\beta^p$ and $\mathcal{E}$ is an exponential random variable with rate one. In addition, for any $k \geq 1$ we have

$$E\left[\left((\tilde{V}_1^{(k)})^+\right)^p\right] \leq C_{\beta,p} \sum_{r=1}^k \left(E[\tilde{N} e^{\beta \tilde{X}}]\right)^r$$
$$\text{and} \qquad E\left[\left((V_1^{(k)})^+\right)^p\right] \leq C_{\beta,p} \sum_{r=1}^k \left(E[N e^{\beta X}]\right)^r.$$

Proof. We only give the proof for $\hat{U}_r$ and $V_1^{(k)}$ since the other expectations are essentially the same. Start by using Markov's inequality to obtain that

$$E\left[\left(\hat{U}_r^+\right)^p\right] = \int_0^\infty P\left(\left(\hat{U}_r^+\right)^p > x\right) dx = \int_0^\infty P\left(e^{\beta \hat{U}_r} > e^{\beta x^{1/p}}\right) dx$$
$$\leq E\left[e^{\beta \hat{U}_r}\right] \int_0^\infty e^{-\beta x^{1/p}} dx = E\left[\bigvee_{\mathbf{j} \in \hat{A}_r} e^{\beta \hat{S}_{\mathbf{j}}}\right] \frac{p}{\beta^p} \int_0^\infty u^{p-1} e^{-u} du,$$

where $\int_0^\infty u^{p-1} e^{-u} du = E[\mathcal{E}^{p-1}] < \infty$ with $\mathcal{E}$ exponentially distributed with rate one. Letting $C_{\beta,p} = pE[\mathcal{E}^{p-1}]/\beta^p$ gives

$$E\left[\left(\hat{U}_r^+\right)^p\right] \leq C_{\beta,p}\, E\left[\bigvee_{\mathbf{j} \in \hat{A}_r} e^{\beta \hat{S}_{\mathbf{j}}}\right] \leq C_{\beta,p}\, E\left[\sum_{\mathbf{j} \in \hat{A}_r} e^{\beta \hat{S}_{\mathbf{j}}}\right].$$

Now use Lemma 2 and the observation that $E[e^{\beta \tilde{X}}] \leq E[\tilde{N} e^{\beta \tilde{X}}]$ to complete the first part of the proof.

For $V_1^{(k)}$ note that for $k \geq 1$, the same arguments used above give

$$E\left[\left((V_1^{(k)})^+\right)^p\right] \leq C_{\beta,p} E\left[\bigvee_{r=1}^k \bigvee_{(1,\mathbf{j}) \in A_r} e^{\beta S_{(i,\mathbf{j})}}\right] \leq C_{\beta,p} E\left[\sum_{r=1}^k \sum_{(1,\mathbf{j}) \in A_r} e^{\beta S_{(i,\mathbf{j})}}\right] = C_{\beta,p} \sum_{r=1}^k E\left[\sum_{(1,\mathbf{j}) \in A_r} e^{\beta S_{(i,\mathbf{j})}}\right].$$

Now use Lemma 2 to obtain that

$$E\left[\sum_{(1,\mathbf{j}) \in A_r} e^{\beta S_{(i,\mathbf{j})}}\right] = E[e^{\beta X}] \left(E[N e^{\beta X}]\right)^{r-1} \leq \left(E[N e^{\beta X}]\right)^r,$$

which yields

$$E\left[\left(V_1^{(k)})^+\right)^p\right] \leq C_{\beta,p} \sum_{r=1}^k \left(E[N e^{\beta X}]\right)^r.$$

This completes the proof. □

The following result gives (20), which we state alongside its symmetric version for the truncation of the branches of the limit W.

Lemma 4. *Suppose Assumption 1 is satisfied and let ν_k be the probability measure of $\bigvee_{r=0}^k U_r$. Then, for any $r_n \to \infty$ as $n \to \infty$, and any $p \geq 1$, we have that*

$$\lim_{n\to\infty} E\left[\left|W^{(n)} - \bigvee_{r=0}^{r_n} \hat{U}_r\right|^p\right] = 0 \qquad \text{and} \qquad \lim_{n\to\infty} E\left[\left|W - \bigvee_{r=0}^{r_n} U_r\right|^p\right] = 0.$$

In particular, this implies that, as $n \to \infty$,

$$d_p(\mu_n, \hat{\nu}_{r_n}) \to 0 \qquad \text{and} \qquad d_p(\nu_{r_n}, \mu) \to 0.$$

Proof. Let $\beta > 0$ be the one from Assumption 1 (ii) and let $\rho_\beta = E[Ne^{\beta\chi}]E\left[e^{-\beta\tau}\right] < 1$. Fix $0 < \epsilon < 1 - \rho_\beta$ and note that

$$E[\tilde{N}e^{\beta\tilde{\chi}}]E\left[e^{-\beta\tilde{\tau}}\right] = E[\tilde{N}e^{\beta\tilde{\chi}}]\left(\frac{\lambda_n^*}{\lambda_n^* + \beta}\right).$$

By Assumption 1 (i) and (iii) we have that $E[\tilde{N}] \to E[N]$ and $E[\tilde{N}e^{\beta\tilde{\chi}}] \to E[Ne^{\beta\chi}]$, and therefore $\lambda_n^* \to \lambda^*$ as $n \to \infty$ and

$$\lim_{n\to\infty} E[\tilde{N}e^{\beta\tilde{X}}] = \lim_{n\to\infty} E[\tilde{N}e^{\beta\tilde{\chi}}]E\left[e^{-\beta\tilde{\tau}}\right] = \rho_\beta. \tag{23}$$

It follows that for large enough n,

$$E[\tilde{N}e^{\beta\tilde{X}}] \leq \rho_\beta + \epsilon < 1.$$

Next, note that

$$\begin{aligned} E\left[\left|W^{(n)} - \bigvee_{r=0}^{r_n} \hat{U}_r\right|^p\right] &= E\left[\left(\left(\bigvee_{r=r_n+1}^{\infty} \hat{U}_r - \bigvee_{r=0}^{r_n} \hat{U}_r\right)^+\right)^p\right] \leq E\left[\left(\left(\bigvee_{r=r_n+1}^{\infty} \hat{U}_r\right)^+\right)^p\right] \\ &= E\left[\bigvee_{r=r_n+1}^{\infty} \left(\hat{U}_r^+\right)^p\right] \leq \sum_{r=r_n+1}^{\infty} E\left[\left(\hat{U}_r^+\right)^p\right]. \end{aligned}$$

Now use Lemma 3 to obtain that

$$E\left[\left(\hat{U}_r^+\right)^p\right] \leq C_{\beta,p}E[\tilde{N}]\left(E[\tilde{N}e^{\beta\tilde{X}}]\right)^r \leq C_{\beta,p}E[\tilde{N}]\left(\rho_\beta + \epsilon\right)^r.$$

It follows that, for sufficiently large n,

$$\sum_{r=r_n+1}^{\infty} E\left[\left(\hat{U}_r^+\right)^p\right] \leq C_{\beta,p}E[\tilde{N}] \sum_{r=r_n+1}^{\infty} (\rho_\beta + \epsilon)^r = O\left((\rho_\beta + \epsilon)^{r_n}\right)$$

as $n \to \infty$, since $\rho_\beta + \epsilon < 1$.

The proof involving W and the $\{U_r\}_{r\geq 0}$ is essentially the same and is therefore omitted. □

We now move on to the proof of (21), which involves the coupling of weighted branching processes with different mark distributions. The details of the coupling we will use are contained in the following lemma. We use $\|X\|_p = (E[|X|^p])^{1/p}$ to denote the L_p norm.

LEMMA 5. *Let $\zeta^{(1)}, \zeta^{(2)}, \zeta^{(3)}$ be i.i.d. Uniform[0,1] random variables. Construct $(\tilde{N}, \tilde{\chi}, \tilde{\tau}, N, \chi, \tau)$ according to:*

$$\begin{aligned} \tilde{\tau} &= -\frac{1}{\lambda_n^*} \log \zeta^{(1)} && \text{and} && \tau = -\frac{1}{\lambda^*} \log \zeta^{(1)}, \\ \tilde{N} &= F_n^{-1}(\zeta^{(2)}) && \text{and} && N = F^{-1}(\zeta^{(2)}), \\ \tilde{\chi} &= G_{\tilde{N}}^{-1}(\zeta^{(3)}) && \text{and} && \chi = G_N^{-1}(\zeta^{(3)}), \end{aligned}$$

where $g^{-1}(t) = \inf\{x \in \mathbb{R} : g(x) \geq t\}$ (this is the standard inverse transform construction). Let $(\tilde{N}, \tilde{X}, N, X) = (\tilde{N}, \tilde{\chi} - \tilde{\tau}, N, \chi - \tau)$. Then, for any $p \geq 1$,

$$\|\tilde{X} - X\|_p \leq \frac{E[|N - \tilde{N}|] \|\mathcal{E}\|_p}{\lambda} + \frac{2(E[|N - \tilde{N}|])^{1/p} |\log E[|N - \tilde{N}|]|}{\beta} \left(1 + \frac{pE[e^{\beta\chi} + e^{\beta\tilde{\chi}}]}{|\log E[|N - \tilde{N}|]|}\right)^{1/p},$$

where $\mathcal{E}$ is an exponential random variable with rate one.

Proof. Start by noting that

$$\begin{aligned} \|\tilde{X} - X\|_p &\leq \|\tilde{\tau} - \tau\|_p + \|\tilde{\chi} - \chi\|_p = \left| \frac{1}{\lambda_n^*} - \frac{1}{\lambda^*} \right| \|\mathcal{E}\|_p + \| |\tilde{\chi} - \chi| 1(\tilde{N} \neq N)\|_p \\ &\leq \frac{|E[N] - E[\tilde{N}]|}{\lambda E[N] E[\tilde{N}]} \|\mathcal{E}\|_p + \|\tilde{\chi} 1(\tilde{N} \neq N)\|_p + \|\chi 1(\tilde{N} \neq N)\|_p. \end{aligned}$$

Next, let $\beta > 0$ be the one from Assumption 1 (ii), set $a_n = \beta^{-p} |\log E[|N - \tilde{N}|]|^p$ and note that

$$\begin{aligned} E\left[\tilde{\chi}^p 1(\tilde{N} \neq N)\right] &\leq a_n P(\tilde{\chi}^p \leq a_n, \tilde{N} \neq N) + E\left[\tilde{\chi}^p 1(\tilde{\chi}^p > a_n, \tilde{N} \neq N)\right] \\ &\leq a_n P(\tilde{N} \neq N) + \int_{a_n}^{\infty} P(\tilde{\chi}^p > t) dt \\ &\leq a_n E[|N - \tilde{N}|] + E[e^{\beta\tilde{\chi}}] \int_{a_n}^{\infty} e^{-\beta t^{1/p}} dt \\ &\leq a_n E[|N - \tilde{N}|] + (p/\beta) a_n^{1-1/p} \int_{a_n}^{\infty} (\beta/p) t^{1/p-1} e^{-\beta t^{1/p}} dt \\ &= a_n E[|N - \tilde{N}|] + (p/\beta) E[e^{\beta\tilde{\chi}}] a_n^{1-1/p} e^{-\beta a_n^{1/p}} \\ &= a_n E[|N - \tilde{N}|] \left(1 + \frac{pE[e^{\beta\tilde{\chi}}]}{|\log E[|N - \tilde{N}|]|}\right). \end{aligned}$$

Similarly,

$$E\left[\chi^p 1(\tilde{N} \neq N)\right] \leq a_n E[|N - \tilde{N}|] \left(1 + \frac{pE[e^{\beta\chi}]}{|\log E[|N - \tilde{N}|]|}\right).$$

Combining the two we obtain:

$$\|\tilde{\chi} 1(\tilde{N} \neq N)\|_p + \|\chi 1(\tilde{N} \neq N)\|_p \leq 2a_n^{1/p} (E[|N - \tilde{N}|])^{1/p} \left(1 + \frac{pE[e^{\beta\chi} + e^{\beta\tilde{\chi}}]}{|\log E[|N - \tilde{N}|]|}\right)^{1/p}.$$

The result now follows. □

The proof of (21) is given in the following theorem. We point out that if $f_n \equiv f$, then $(\tilde{N}, \tilde{\chi}, \tilde{\tau}) \stackrel{\mathcal{D}}{=} (N, \chi, \tau)$, and (21) would immediately follow from Lemma 4. The result below establishes (21) provided only that $f_n \xrightarrow{d_1} f$.

Theorem 3. *Suppose that Assumption 1 is satisfied. Then, for $\tilde{\nu}_k$, the probability measure of $\bigvee_{r=0}^k \tilde{U}_r$, μ, the probability measure of $\bigvee_{r=0}^\infty U_r$, any $p \geq 1$ and any $r_n \to \infty$, we have that*

$$d_p(\tilde{\nu}_{r_n}, \mu) \to 0 \qquad n \to \infty.$$

Proof. We start by constructing two marked Galton-Watson processes using an i.i.d. sequence $\{(\tilde{N}_{\mathbf{i}}, \tilde{X}_{\mathbf{i}}, N_{\mathbf{i}}, X_{\mathbf{i}})\}_{\mathbf{i} \in U}$ of copies of the vector $(\tilde{N}, \tilde{X}, N, X)$ from Lemma 5. Set $\tilde{A}_0 = \{\emptyset\} = A_0$ and $\tilde{A}_r = \{(\mathbf{i}, i_r) : \mathbf{i} \in \tilde{A}_{r-1}, 1 \leq i_r \leq \tilde{N}_{\mathbf{i}}\}$, $A_r = \{(\mathbf{i}, i_r) : \mathbf{i} \in A_{r-1}, 1 \leq i_r \leq N_{\mathbf{i}}\}$ for $r \geq 1$.

Fix $0 < \epsilon < \rho_\beta = E[N e^{\beta X}] < 1$ and assume from now on that n is sufficiently large to ensure that $E[|N - \tilde{N}|] \leq 1$ and $\tilde{\rho}_\beta = E[\tilde{N} e^{\beta \tilde{X}}] \leq \rho_\beta + \epsilon$. Define

$$k_n = \frac{(1-\epsilon)|\log E[|N - \tilde{N}|]|}{\log E[N]} 1(E[N] > 1) + (E[|N - \tilde{N}|])^{(\epsilon - 1)/p} 1(E[N] = 1).$$

Define the processes U_r and $\tilde{U}_r$, $r \geq 1$, on their corresponding trees and note that

$$\begin{aligned} d_p(\tilde{\nu}_{r_n}, \mu) &\leq \left\| \bigvee_{r=0}^{r_n} \tilde{U}_r - \bigvee_{r=0}^{\infty} U_r \right\|_p \\ &\leq \left\| \bigvee_{r=0}^{r_n \wedge k_n} \tilde{U}_r - \bigvee_{r=0}^{r_n \wedge k_n} U_r \right\|_p + \left\| \bigvee_{r=r_n \wedge k_n + 1}^{\infty} \tilde{U}_r^+ \right\|_p + \left\| \bigvee_{r=r_n \wedge k_n + 1}^{\infty} U_r^+ \right\|_p. \end{aligned}$$

Now use Lemma 3 to obtain that

$$E\left[\left(\bigvee_{r=r_n \wedge k_n + 1}^{\infty} U_r^+ \right)^p \right] \leq \sum_{r=r_n \wedge k_n + 1}^{\infty} E[(U_r^+)^p] \leq C_{\beta,p} E[N] \sum_{r=r_n \wedge k_n + 1}^{\infty} \rho_\beta^r,$$

and symmetrically,

$$E\left[\left(\bigvee_{r=r_n \wedge k_n + 1}^{\infty} \tilde{U}_r^+ \right)^p \right] \leq C_{\beta,p} E[\tilde{N}] \sum_{r=r_n \wedge k_n + 1}^{\infty} (\tilde{\rho}_\beta)^r.$$

It follows that as $n \to \infty$,

$$d_p(\tilde{\nu}_{r_n}, \mu) \leq \left\| \bigvee_{r=0}^{r_n \wedge k_n} \tilde{U}_r - \bigvee_{r=0}^{r_n \wedge k_n} U_r \right\|_p + O\left((\rho_\beta + \epsilon)^{r_n \wedge k_n} \right).$$

It remains to show that the norm in the expression above converges to zero. To this end, note that for any $k \geq 1$ we can write

$$\bigvee_{r=0}^{k} \tilde{U}_r = 0 \vee \bigvee_{i=1}^{\tilde{N}_\emptyset} \tilde{V}_i^{(k)} = \bigvee_{i=1}^{\tilde{N}_\emptyset} (\tilde{V}_i^{(k)})^+ \qquad \text{and} \qquad \bigvee_{r=0}^{k} U_r = 0 \vee \bigvee_{i=1}^{N_\emptyset} V_i^{(k)} = \bigvee_{i=1}^{N_\emptyset} (V_i^{(k)})^+.$$

Now use the inequality

$$\left| \max_{1 \leq i \leq k} x_i - \max_{1 \leq i \leq k} y_i \right| \leq \max_{1 \leq i \leq k} |x_i - y_i| \tag{24}$$

for any sequences of real numbers $\{x_i\}_{i \geq 1}$ and $\{y_i\}_{i \geq 1}$ and any $k \geq 1$, to obtain

$$\left| \bigvee_{r=0}^{k} \tilde{U}_r - \bigvee_{r=0}^{k} U_r \right| \leq \bigvee_{i=1}^{\tilde{N}_\emptyset \wedge N_\emptyset} \left| (\tilde{V}_i^{(k)})^+ - (V_i^{(k)})^+ \right| \vee \bigvee_{i=\tilde{N}_\emptyset \wedge N_\emptyset + 1}^{\tilde{N}_\emptyset} (\tilde{V}_i^{(k)})^+ \vee \bigvee_{i=\tilde{N}_\emptyset \wedge N_\emptyset + 1}^{N_\emptyset} (V_i^{(k)})^+,$$

with the convention that $\bigvee_{i=a}^b x_i \equiv 1$ and $\sum_{i=a}^b x_i \equiv 0$ if $a > b$. It follows that

$$\begin{aligned}
&E\left[\left|\bigvee_{r=0}^k \tilde{U}_r - \bigvee_{r=0}^k U_r\right|^p\right] \\
&\le E\left[\bigvee_{i=1}^{\tilde{N}_\emptyset \wedge N_\emptyset} \left|(\tilde{V}_i^{(k)})^+ - (V_i^{(k)})^+\right|^p \vee \bigvee_{i=\tilde{N}_\emptyset \wedge N_\emptyset + 1}^{\tilde{N}_\emptyset} ((\tilde{V}_i^{(k)})^+)^p \vee \bigvee_{i=\tilde{N}_\emptyset \wedge N_\emptyset + 1}^{N_\emptyset} ((V_i^{(k)})^+)^p\right] \\
&\le E\left[\sum_{i=1}^{\tilde{N}_\emptyset \wedge N_\emptyset} \left|(\tilde{V}_i^{(k)})^+ - (V_i^{(k)})^+\right|^p + \sum_{i=\tilde{N}_\emptyset \wedge N_\emptyset + 1}^{\tilde{N}_\emptyset} ((\tilde{V}_i^{(k)})^+)^p + \sum_{i=\tilde{N}_\emptyset \wedge N_\emptyset + 1}^{N_\emptyset} ((V_i^{(k)})^+)^p\right] \\
&\le E[N] E\left[\left|(\tilde{V}_1^{(k)})^+ - (V_1^{(k)})^+\right|^p\right] + E[(\tilde{N} - N)^+] E\left[((V_1^{(k)})^+)^p\right] \\
&\quad + E[(N - \tilde{N})^+] E\left[((\tilde{V}_1^{(k)})^+)^p\right].
\end{aligned}$$

Now use Lemma 3 to obtain that

$$E\left[((V_1^{(k)})^+)^p\right] \le C_{\beta,p} \sum_{r=1}^k \rho_\beta^r \qquad \text{and} \qquad E\left[((\tilde{V}_1^{(k)})^+)^p\right] \le C_{\beta,p} \sum_{r=1}^k (\tilde{\rho}_\beta)^r \le C_{\beta,p} \sum_{r=1}^k (\rho_\beta + \epsilon)^r,$$

from where it follows that

$$\left\|\bigvee_{r=0}^k \tilde{U}_r - \bigvee_{r=0}^k U_r\right\|_p \le (E[N])^{1/p} \left\|(\tilde{V}_1^{(k)})^+ - (V_1^{(k)})^+\right\|_p + (E[|\tilde{N} - N|] K)^{1/p},$$

with $K = K(\beta, p, \epsilon) = C_{\beta,p}(1 - \rho_\beta - \epsilon)^{-1}$.

We now repeat the same type of arguments to obtain a recursive inequality for $u_k \triangleq \|(\tilde{V}_1^{(k)})^+ - (V_1^{(k)})^+\|_p$. To do this recall that

$$V_1^{(k)} = \max\left\{X_1, X_1 + \bigvee_{l=1}^{N_1} V_{(1,l)}^{(k-1)}\right\} \stackrel{\mathcal{D}}{=} \max\left\{X,\, X + \bigvee_{i=1}^N V_i^{(k-1)}\right\} = X + \bigvee_{i=1}^N (V_i^{(k-1)})^+,$$

with the $\{V_i^{(k-1)}\}_{i\ge1}$ i.i.d. copies of $V_1^{(k-1)}$ independent of (N, X). The same is true for the process with the $\sim$ notation. It follows from using (24) and Minkowski's inequality that

$$\begin{aligned}
u_k &= \left\|\max\left\{0,\, \tilde{X} + \bigvee_{i=1}^{\tilde{N}} (\tilde{V}_i^{(k-1)})^+\right\} - \max\left\{0,\, X + \bigvee_{i=1}^N (V_i^{(k-1)})^+\right\}\right\|_p \\
&\le \left\|\tilde{X} + \bigvee_{i=1}^{\tilde{N}} (\tilde{V}_i^{(k-1)})^+ - X - \bigvee_{i=1}^N (V_i^{(k-1)})^+\right\|_p \\
&\le \|\tilde{X} - X\|_p + \left\|\bigvee_{i=1}^{\tilde{N}} (\tilde{V}_i^{(k-1)})^+ - \bigvee_{i=1}^N (V_i^{(k-1)})^+\right\|_p.
\end{aligned}$$

Now use (24) again to obtain that

$$E\left[\left|\bigvee_{i=1}^{\tilde{N}} (\tilde{V}_i^{(k-1)})^+ - \bigvee_{i=1}^N (V_i^{(k-1)})^+\right|^p\right]$$

$$
\begin{aligned}
&\leq E\left[\bigvee_{i=1}^{\tilde{N}\wedge N} \left| (\tilde{V}_i^{(k-1)})^+ - (V_i^{(k-1)})^+ \right|^p \vee \bigvee_{i=\tilde{N}\wedge N+1}^{\tilde{N}} ((\tilde{V}_i^{(k-1)})^+)^p \vee \bigvee_{i=\tilde{N}\wedge N+1}^{N} ((V_i^{(k-1)})^+)^p \right] \\
&\leq E\left[\sum_{i=1}^{\tilde{N}\wedge N} \left| (\tilde{V}_i^{(k-1)})^+ - (V_i^{(k-1)})^+ \right|^p + \sum_{i=\tilde{N}\wedge N+1}^{\tilde{N}} ((\tilde{V}_i^{(k-1)})^+)^p + \sum_{i=\tilde{N}\wedge N+1}^{N} ((V_i^{(k-1)})^+)^p \right] \\
&\leq E[N] E\left[\left| (\tilde{V}_1^{(k-1)})^+ - (V_1^{(k-1)})^+ \right|^p \right] + E[(\tilde{N}-N)^+] E[((\tilde{V}_1^{(k-1)})^+)^p] \\
&\quad + E[(N-\tilde{N})^+] E[((V_1^{(k-1)})^+)^p] \\
&\leq E[N] u_{k-1}^p + E[|\tilde{N}-N|] K.
\end{aligned}
$$

Therefore, for $k \geq 2$,

$$
\begin{aligned}
u_k &\leq \|\tilde{X} - X\|_p + \left(E[N] u_{k-1}^p + E[|\tilde{N}-N|] K \right)^{1/p} \\
&\leq \|\tilde{X} - X\|_p + (E[N])^{1/p} u_{k-1} + (E[|\tilde{N}-N|] K)^{1/p}.
\end{aligned}
$$

In general, iterating the above and using the boundary condition $u_1 = \|X - \tilde{X}\|_p$ gives

$$
\begin{aligned}
u_k &\leq \left(\|\tilde{X} - X\|_p + (E[|\tilde{N}-N|] K)^{1/p} \right) \sum_{i=0}^{k-2} (E[N])^{i/p} + (E[N])^{(k-1)/p} u_1 \\
&\leq \left(\|\tilde{X} - X\|_p + (E[|\tilde{N}-N|] K)^{1/p} \right) \sum_{i=0}^{k-1} (E[N])^{i/p}.
\end{aligned}
$$

Now use Lemma 5 and the earlier assumption $E[|\tilde{N}-N|] \leq 1$ to obtain that

$$
\|\tilde{X} - X\|_p + (E[|\tilde{N}-N|] K)^{1/p} \leq H (E[|\tilde{N}-N|])^{1/p} |\log E[|\tilde{N}-N|]|,
$$

for some other constant $H = H(\beta, p, \epsilon, \lambda) \geq K^{1/p}$.

We have thus shown that for any $k \geq 1$,

$$
\begin{aligned}
\left\| \bigvee_{r=0}^{k} \tilde{U}_r - \bigvee_{r=0}^{k} U_r \right\|_p &\leq (E[N])^{1/p} H (E[|\tilde{N}-N|])^{1/p} |\log E[|\tilde{N}-N|]| \sum_{i=0}^{k-1} (E[N])^{i/p} + (E[|\tilde{N}-N|] K)^{1/p} \\
&\leq H (E[|\tilde{N}-N|])^{1/p} |\log E[|\tilde{N}-N|]| \sum_{i=0}^{k} (E[N])^{i/p}.
\end{aligned}
$$

Our choice of k_n guarantees that

$$
\begin{aligned}
\left\| \bigvee_{r=0}^{r_n \wedge k_n} \tilde{U}_r - \bigvee_{r=0}^{r_n \wedge k_n} U_r \right\|_p &\leq H (E[|\tilde{N}-N|])^{1/p} |\log E[|\tilde{N}-N|]| \sum_{i=0}^{k_n} (E[N])^{i/p} \\
&= O\left((E[|\tilde{N}-N|])^{\epsilon/p} |\log E[|\tilde{N}-N|]| \right)
\end{aligned}
$$

as $n \to \infty$. Since $E[|\tilde{N}-N|]$ is the Wasserstein distance of order one (d_1) between distributions f_n and f (see, e.g., [19, 31]), Assumption 1 (i) and Theorem 6.8 in [33] imply that

$$
\lim_{n\to\infty} E[|\tilde{N}-N|] = 0.
$$

This completes the proof. □

The last result needed to complete the proof of Theorem 1 is to establish (22), which we will do by showing that we can choose $r_n \to \infty$ in such a way that with high probability no paths merge in the predecessor graph at distance up to r_n from the tagged job.

To aide with the identification of the merging paths, we need some additional index notation. We use $|\mathbf{i}| = |(i_1, i_2, \dots, i_k)| = k$ to denote the length of $\mathbf{i}$, and write $\mathbf{j} \prec \mathbf{i}$ if either (i) $|\mathbf{j}| < |\mathbf{i}|$ or (ii) $|\mathbf{j}| = |\mathbf{i}|$ and $\min\{s \geq 1 : j_s < i_s\} \leq |\mathbf{i}|$. We use this ordering of the labels to identify the "redundant" labels, i.e., the labels representing jobs in the predecessor graph that are also represented by a smaller label. More precisely, if we write $\mathbf{i} \sim \mathbf{j}$ to denote that $\mathbf{i}$ and $\mathbf{j}$ belong to the same job in the predecessor graph, then the "redundant" labels are those in the sets

$$M_r = \{\mathbf{i} \in \hat{A}_r : \mathbf{i} \sim \mathbf{j} \text{ for some } \mathbf{j} \prec \mathbf{i}\}, \qquad r \geq 1.$$

Note that "redundant" labels are created when paths merge in the predecessor graph.

The following result establishes (22), which completes the proof of Theorem 1.

THEOREM 4. *Suppose that Assumption 1 is satisfied. Then, for $\hat{\nu}_k$, the probability measure of $\bigvee_{r=0}^k \hat{U}_r$, and $\tilde{\nu}_k$, the probability measure of $\bigvee_{r=0}^k \tilde{U}_r$, we have that for any $p \geq 1$ and any $r_n \to \infty$ such that*

$$\lim_{n\to\infty} \frac{1}{n} \left(\sum_{s=0}^{r_n} E[\tilde{N}]^s \right)^2 = 0, \tag{25}$$

we have that

$$d_p(\hat{\nu}_{r_n}, \tilde{\nu}_{r_n}) \to 0 \qquad n \to \infty.$$

Proof. From the definition of the Wasserstein metric, we need to construct a coupling of $\hat{\nu}_{r_n}$ and $\tilde{\nu}_{r_n}$ for which we can show that their L_p distance converges to zero. To do this, let time run in reverse in the queueing system and construct the predecessor graph of the tagged job. Recall that the inter-arrival times between a job with label $\mathbf{i}$ and each of its $\hat{N}_{\mathbf{i}}$ immediate predecessors are given by $\{\hat{\tau}_{(\mathbf{i},1)}, \dots, \hat{\tau}_{(\mathbf{i},\hat{N}_{\mathbf{i}})}\}$, and the marginal distribution of each of the $\hat{\tau}_{(\mathbf{i},j)}$ is exponential with rate λ_n^*. We now define the coupled marked Galton-Watson process as follows.

To describe the coupling we will use the set of redundant labels defined earlier, M_r. In essence, to construct the sets $\tilde{A}_r$ we will copy all the labels, along with their marks, from the set $\hat{A}_r \cap M_r^c$, and every time we encounter a label $\mathbf{i} \in M_r$ such that $(\mathbf{i}|r-1) \notin M_{r-1}$, we will sample an independent subtree from that point onwards. To keep track of the nodes in $\tilde{A}_r$ that were not copied from $\hat{A}_r$ we will recursively construct a discrepancy set $\tilde{M}_r$. To start, set $\tilde{N}_\emptyset = \hat{N}_\emptyset$ and $\tilde{A}_0 = \{\emptyset\}$ and $\tilde{M}_\emptyset = \varnothing$. The construction of $\tilde{A}_r$ for $r \geq 1$ is done as follows:

i) Initialize $\tilde{M}_r = \varnothing$.
ii) For each $\mathbf{i} \in \tilde{A}_{r-1} \cap \tilde{M}_{r-1}^c$ and $1 \leq j \leq \tilde{N}_{\mathbf{i}}$:
   a. If $(\mathbf{i}, j) \in \hat{A}_r \cap M_r^c$, set $(\tilde{N}_{(\mathbf{i},j)}, \tilde{\chi}_{(\mathbf{i},j)}, \tilde{\tau}_{(\mathbf{i},j)}) = (\hat{N}_{(\mathbf{i},j)}, \hat{\chi}_{(\mathbf{i},j)}, \hat{\tau}_{(\mathbf{i},j)})$;
   b. Else, generate $(\tilde{N}_{(\mathbf{i},j)}, \tilde{\chi}_{(\mathbf{i},j)}, \tilde{\tau}_{(\mathbf{i},j)})$ independently of all the random variables in the predecessor graph and any other random variables in the coupled tree. Add $(\mathbf{i}, j)$ to $\tilde{M}_r$.
iii) For each $\mathbf{i} \in \tilde{A}_{r-1} \cap \tilde{M}_{r-1}$ and $1 \leq j \leq \tilde{N}_{\mathbf{i}}$: generate $(\tilde{N}_{(\mathbf{i},j)}, \tilde{\chi}_{(\mathbf{i},j)}, \tilde{\tau}_{(\mathbf{i},j)})$ independently of all the random variables in the predecessor graph and any other random variables in the coupled tree. Add $(\mathbf{i}, j)$ to $\tilde{M}_r$.

Note that this construction ensures that $\hat{A}_k \cap M_k^c \subseteq \tilde{A}_k$, however, paths in M_k may not even be in $\tilde{A}_k$. Next, define the event

$$B_k = \bigcap_{r=1}^k \{M_r = \varnothing\},$$

and note that if B_k happens, then

$$\bigvee_{r=0}^{k} \hat{U}_r = \bigvee_{r=0}^{k} \tilde{U}_r.$$

Now that we have a coupling, note that by the Cauchy-Schwartz inequality,

$$E\left[\left| \bigvee_{r=0}^{k} \hat{U}_r - \bigvee_{r=0}^{k} \tilde{U}_r \right|^p \right] = E\left[\left| \bigvee_{r=0}^{k} \hat{U}_r - \bigvee_{r=0}^{k} \tilde{U}_r \right|^p 1(B_k^c) \right] \leq \left\| \bigvee_{r=0}^{k} \hat{U}_r - \bigvee_{r=0}^{k} \tilde{U}_r \right\|_{2p}^p P(B_k^c)^{1/2}$$
$$\leq \left(\left\| \bigvee_{r=0}^{k} \hat{U}_r \right\|_{2p} + \left\| \bigvee_{r=0}^{k} \tilde{U}_r \right\|_{2p} \right)^p P(B_k^c)^{1/2}.$$

Moreover, by Lemma 3 we have

$$\left\| \bigvee_{r=0}^{k} \hat{U}_r \right\|_{2p} \leq \left(E\left[\sum_{r=1}^{k} (\hat{U}_r^+)^{2p} \right] \right)^{1/(2p)} \leq \left(\sum_{r=1}^{k} C_{\beta,2p} E[\tilde{N}] (\tilde{\rho}_\beta)^r \right)^{1/(2p)},$$

where $\tilde{\rho}_\beta = E[\tilde{N} e^{\beta(\tilde{\chi} - \tilde{\tau})}]$. The same inequality holds for the $\{\tilde{U}_r\}$. Since under Assumption 1 we have that $E[\tilde{N}] \to E[N]$ and $\tilde{\rho}_\beta \to \rho_\beta < 1$ as $n \to \infty$, we have that for any $0 < \epsilon < 1 - \rho_\beta$ and sufficiently large n,

$$\left\| \bigvee_{r=0}^{k} \hat{U}_r \right\|_{2p} + \left\| \bigvee_{r=0}^{k} \tilde{U}_r \right\|_{2p} \leq \left(2C_{\beta,2p} (E[N] + \epsilon) \sum_{r=1}^{\infty} (\rho_\beta + \epsilon)^r \right)^{1/(2p)} < \infty.$$

Hence, it only remains to show that $P(B_{r_n}^c) \to 0$ as $n \to \infty$ for r_n satisfying (25).

Define $x_n^{-2} = n^{-1} \left(\sum_{s=0}^{r_n+1} E[\tilde{N}]^s \right)^2$ and note that $x_n \to \infty$ as $n \to \infty$. Let $D_k = \bigcap_{s=1}^{k} \{ |\tilde{A}_s| \leq E[\tilde{N}]^s x_n \}$ and note that for any $k \geq 1$ we have

$$P(B_k^c) \leq P(B_k^c \cap D_k) + P(D_k^c) = \sum_{r=1}^{k} P(D_k \cap B_{r-1} \cap B_r^c) + P(D_k^c)$$
$$= \sum_{r=1}^{k} P\left(D_r \cap B_{r-1} \cap \bigcup_{\mathbf{i} \in \hat{A}_r} \{ \mathbf{i} \sim \mathbf{j} \text{ for some } \mathbf{j} \prec \mathbf{i} \} \right) + P(D_k^c).$$

To compute each of these probabilities, let $L_{\mathbf{i}} = \{ \mathbf{j} : \mathbf{j} \prec \mathbf{i} \text{ and } \mathbf{j} \sim \mathbf{i} \}$ denote the set of smaller labels that the job represented by node $\mathbf{i}$ may have. Define $\mathcal{F}_k = \sigma\left((\hat{N}_{\mathbf{i}}, L_{\mathbf{i}}), \tilde{N}_{\mathbf{j}} : \mathbf{i} \in \bigcup_{r=0}^{k} \hat{A}_r, \mathbf{j} \in \bigcup_{r=0}^{k} \tilde{A}_r \right)$ and note that B_k, D_{k+1} and $\hat{A}_{k+1}$ are all measurable with respect to $\mathcal{F}_k$. We then have

$$P\left(D_r \cap B_{r-1} \cap \bigcup_{\mathbf{i} \in \hat{A}_r} \{ \mathbf{i} \sim \mathbf{j} \text{ for some } \mathbf{j} \prec \mathbf{i} \} \right)$$
$$= E\left[1(D_r \cap B_{r-1}) P\left(\left. \bigcup_{\mathbf{i} \in \hat{A}_r} \{ \mathbf{i} \sim \mathbf{j} \text{ for some } \mathbf{j} \prec \mathbf{i} \} \right| \mathcal{F}_{r-1} \right) \right].$$

To analyze each of the conditional probabilities, note that the job with label $(\mathbf{i}|r-1)$ has up to $\hat{N}_{(\mathbf{i}|r-1)}$ immediate predecessors, and we can associate each of them with a job fragment. Suppose that $\mathbf{i}$ is the immediate predecessor of fragment l ($l \in \{1, 2, \ldots, \hat{N}_{(\mathbf{i}|r-1)}\}$). For there to exist a $\mathbf{j} \prec \mathbf{i}$

such that $\mathbf{i} \sim \mathbf{j}$, it must be that the job labeled $\mathbf{i}$ must also be an immediate predecessor to at least one other job in the predecessor graph with smaller label than $\mathbf{i}$. Let $K_{\mathbf{i}} = \left| \left\{ \mathbf{j} \in \bigcup_{s=0}^r \hat{A}_s : \mathbf{j} \prec \mathbf{i} \right\} \right|$. It follows that since all tags of length $\hat{N}_{\mathbf{i}}$ are equally likely, and provided $K_{\mathbf{i}} + \hat{N}_{\mathbf{i}} \leq n$,

$$P(\mathbf{i} \sim \mathbf{j} \text{ for some } \mathbf{j} \prec \mathbf{i} | \mathcal{F}_{r-1}, \hat{N}_{\mathbf{i}}, K_{\mathbf{i}}) \leq 1 - \frac{\binom{n-K_{\mathbf{i}}-1}{\hat{N}_{\mathbf{i}}-1}}{\binom{n-1}{\hat{N}_{\mathbf{i}}-1}},$$

with the inequality, rather than equality, due to the possibility that some of the jobs counted in $K_{\mathbf{i}}$ may not have arrived by the time the job labeled $\mathbf{i}$ arrives. Let Y denote a hypergeometric random variable with parameters $(n-1, K_{\mathbf{i}}, \hat{N}_{\mathbf{i}} - 1)$ and note that,

$$1 - \frac{\binom{n-K_{\mathbf{i}}-1}{\hat{N}_{\mathbf{i}}-1}}{\binom{n-1}{\hat{N}_{\mathbf{i}}-1}} = 1 - P(Y=0) = P(Y \geq 1) \leq E[Y] = \frac{K_{\mathbf{i}}(\hat{N}_{\mathbf{i}} - 1)}{n-1}.$$

Now let $\hat{Z}_r = \sum_{s=0}^r |\hat{A}_s|$ and note that $K_{\mathbf{i}} \leq \hat{Z}_r$ for any $\mathbf{i} \in \hat{A}_r$, and therefore,

$$P(\mathbf{i} \sim \mathbf{j} \text{ for some } \mathbf{j} \prec \mathbf{i} | \mathcal{F}_{r-1}, \hat{N}_{\mathbf{i}}, K_{\mathbf{i}}) \leq 1(K_{\mathbf{i}} + \hat{N}_{\mathbf{i}} > n) + \frac{K_{\mathbf{i}}(\hat{N}_{\mathbf{i}} - 1)}{n-1} \leq 1(\hat{Z}_r + \hat{N}_{\mathbf{i}} > n) + \frac{\hat{Z}_r(\hat{N}_{\mathbf{i}} - 1)}{n-1}.$$

It follows that for any $1 \leq r \leq r_n$,

$$\begin{aligned}
&P\left(\left. \bigcup_{\mathbf{i} \in \hat{A}_r} \{\mathbf{i} \sim \mathbf{j} \text{ for some } \mathbf{j} \prec \mathbf{i}\} \right| \mathcal{F}_{r-1} \right) \\
&\leq \sum_{\mathbf{i} \in \hat{A}_r} P(\mathbf{i} \sim \mathbf{j} \text{ for some } \mathbf{j} \prec \mathbf{i} | \mathcal{F}_{r-1}) \\
&\leq \sum_{\mathbf{i} \in \hat{A}_r} E\left[\left. 1(\hat{Z}_r + \hat{N}_{\mathbf{i}} > n) + \frac{\hat{Z}_r(\hat{N}_{\mathbf{i}} - 1)}{n-1} \right| \mathcal{F}_{r-1} \right] \\
&\leq \sum_{\mathbf{i} \in \hat{A}_r} \left(1(\hat{Z}_r > n/2) + \frac{2E[\tilde{N}]}{n} + \frac{\hat{Z}_r E[\tilde{N} - 1]}{n-1} \right) \qquad \text{(by Markov's inequality)} \\
&\leq |\hat{A}_r| \left(1(\hat{Z}_r > n/2) + \frac{2(\hat{Z}_r + 1)E[\tilde{N}]}{n} \right).
\end{aligned}$$

Next, set $\tilde{Z}_r = \sum_{s=0}^r |\tilde{A}_s|$ and note that on the event B_{r-1} we can replace $(\hat{A}_r, \hat{Z}_r)$ with $(\tilde{A}_r, \tilde{Z}_r)$, and therefore, for any $1 \leq r \leq r_n$,

$$\begin{aligned}
&P\left(D_r \cap B_{r-1} \cap \bigcup_{\mathbf{i} \in \hat{A}_r} \{\mathbf{i} \sim \mathbf{j} \text{ for some } \mathbf{j} \prec \mathbf{i}\} \right) \\
&\leq E\left[1(D_r \cap B_{r-1}) |\hat{A}_r| \left(1(\hat{Z}_r > n/2) + \frac{2(\hat{Z}_r + 1)E[\tilde{N}]}{n} \right) \right] \\
&\leq E\left[1(D_r) |\tilde{A}_r| \left(1(\tilde{Z}_r > n/2) + \frac{2(\tilde{Z}_r + 1)E[\tilde{N}]}{n} \right) \right] \\
&\leq E\left[E[\tilde{N}]^r x_n \left(1(\tilde{Z}_r > n/2) + \frac{2(\tilde{Z}_r + 1)E[\tilde{N}]}{n} \right) \right] \\
&\leq E[\tilde{N}]^r x_n \left(\frac{2E[\tilde{Z}_r]}{n} + \frac{2E[\tilde{Z}_r + 1]E[\tilde{N}]}{n} \right) \qquad \text{(by Markov's inequality)}
\end{aligned}$$

$$= \frac{2E[\tilde{N}]^r x_n}{n} \left(\sum_{s=0}^{r} E[\tilde{N}]^s + \sum_{s=1}^{r+1} E[\tilde{N}]^s + E[\tilde{N}] \right) \leq \frac{4E[\tilde{N}]^r x_n}{n} \sum_{s=0}^{r+1} E[\tilde{N}]^s.$$

It follows that

$$\begin{aligned} P(B_{r_n}^c) &\leq \sum_{r=1}^{r_n} \frac{4E[\tilde{N}]^r x_n}{n} \sum_{s=0}^{r+1} E[\tilde{N}]^s + P(D_{r_n}^c) \\ &\leq \frac{4x_n}{n} \left(\sum_{r=0}^{r_n+1} E[\tilde{N}]^r \right)^2 + P(D_{r_n}^c) = 4x_n^{-1} + P(D_{r_n}^c). \end{aligned}$$

By Doob's martingale inequality we have $P(D_{r_n}^c) \leq x_n^{-1}$, so we obtain that $P(B_{r_n}^c) \leq 5x_n^{-1} \to 0$ as $n \to \infty$. This completes the proof. □

The last proof in the paper is that of Theorem 2.

Proof of Theorem 2. We need to verify the conditions of Theorem 3.4 in [18] for both the case when N is independent of χ and when they are not. We start with the independent case.

The non-arithmetic condition is immediate from the observation that the $\{\tau_i\}$ are exponentially distributed and independent of $(N, \chi_1, \ldots, \chi_N)$. The derivative and root conditions follow from the assumptions by noting that

$$E\left[\sum_{i=1}^{N} e^{\theta(\chi_i - \tau_i)} (\chi_i - \tau_i) \right] = E[N] E\left[e^{\theta(\chi-\tau)} (\chi - \tau) \right] = E\left[N e^{\theta(\chi-\tau)} (\chi - \tau) \right] > 0,$$

and

$$E\left[\sum_{i=1}^{N} e^{\theta(\chi_i - \tau_i)} \right] = E[N] E\left[e^{\theta(\chi-\tau)} \right] = E\left[N e^{\theta(\chi-\tau)} \right] = 1.$$

To verify condition 1 note that for $\theta > 1$ and $p = \lceil \theta \rceil$, Lemma 4.1 in [16] gives (using $C_i \equiv 1$ for all i),

$$\begin{aligned} E\left[\left(\sum_{i=1}^{N} e^{\chi_i - \tau_i} \right)^{\theta} \right] &\leq \left(E[e^{(p-1)(\chi-\tau)}] \right)^{\theta/(p-1)} E\left[N^{\theta} \right] + E\left[\sum_{i=1}^{N} e^{\theta(\chi_i - \tau_i)} \right] \\ &\leq E\left[e^{\theta(\chi-\tau)} \right] E[N^{\theta}] + 1 \qquad \text{(by Jensen's inequality)}, \\ &= E\left[N^{\theta} e^{\theta(\chi-\tau)} \right] + 1. \end{aligned}$$

For $0 < \theta \leq 1$, the same arguments give for any $0 < \epsilon < 1$,

$$\begin{aligned} E\left[\left(\sum_{i=1}^{N} e^{\theta(\chi_i - \tau_i)/(1+\epsilon)} \right)^{1+\epsilon} \right] &\leq \left(E[e^{\theta(\chi-\tau)/(1+\epsilon)}] \right)^{1+\epsilon} E\left[N^{1+\epsilon} \right] + E\left[\sum_{i=1}^{N} e^{\theta(\chi_i - \tau_i)} \right] \\ &\leq E\left[e^{\theta(\chi-\tau)} \right] E[N^{1+\epsilon}] + 1 \\ &= E\left[N^{1+\epsilon} e^{\theta(\chi-\tau)} \right] + 1. \end{aligned}$$

Since $E[N^{\theta \vee (1+\epsilon)} e^{\theta(\chi-\tau)}] < \infty$ by assumption, all the conditions of the theorem are satisfied and the asymptotic behavior of W follows.

For the general case we have that the fact that the $\tilde{\tau}$ is exponentially distributed and independent of (N, χ) gives the non-arithmetic condition again, and all other conditions are specifically stated as assumptions in the theorem (this case corresponds to taking $(Q, N, C_1, C_2, \ldots) = (e^{\chi-\tau}, N, e^{\chi-\tau}, e^{\chi-\tau}, \ldots)$). It now follows from Theorem 3.4 in [18] that

$$P(V > x) \sim H e^{-\theta x}$$

with

$$H = \frac{E\left[e^{\theta(\chi-\tau)} \vee \bigvee_{i=1}^{N} e^{\theta(\chi-\tau)} e^{\theta V_i} - \sum_{i=1}^{N} e^{\theta(\chi-\tau)} e^{\theta V_i}\right]}{\theta E\left[N e^{\theta(\chi-\tau)}(\chi-\tau)\right]}$$
$$= \frac{E\left[e^{\theta(\chi-\tau)} \left(\bigvee_{i=1}^{N} e^{\theta V_i^+} - \sum_{i=1}^{N} e^{\theta V_i}\right)\right]}{\theta E\left[N e^{\theta(\chi-\tau)}(\chi-\tau)\right]}.$$

To obtain the asymptotic behavior of W note that

$$P(W > x) = P\left(\bigvee_{i=1}^{N} V_i^+ > x\right) = E\left[1 - P(V^+ \le x)^N\right],$$

and use Jensen's inequality to obtain that

$$P(W > x) = E\left[1 - P(V^+ \le x)^N\right] \le 1 - P(V^+ \le x)^{E[N]} \sim E[N]P(V^+ > x)$$

as $x \to \infty$. To obtain a lower bound note that

$$P(W > x) = E\left[1 - e^{N \log(1 - P(V^+ > x))}\right] \ge E\left[1 - e^{-NP(V^+ > x)}\right]$$
$$= P(V^+ > x) E\left[N\left(\frac{1 - e^{-NP(V^+ > x)}}{NP(V^+ > x)}\right)\right],$$

and use Fatou's lemma to obtain that

$$\liminf_{x\to\infty} \frac{P(W > x)}{P(V^+ > x)} \ge E\left[N \liminf_{x\to\infty}\left(\frac{1 - e^{-NP(V^+ > x)}}{NP(V^+ > x)}\right)\right] = E[N].$$

We conclude that

$$P(W > x) \sim E[N]P(V^+ > x) = E[N]P(V > x) \qquad x \to \infty.$$

This completes the proof. □

Acknowledgments. We would like to thank three anonymous referees who provided many thoughtful comments that helped improve the presentation of the paper.

Mariana Olvera-Cravioto's work was funded by grant NSF CMMI-1537638.

References

[1] Aldous D, Bandyopadhyay A (2005) A survey of max-type recursive distributional equation. *Ann. Appl. Probab.* 15(2):1047–1110.

[2] Alsmeyer G, Biggins J, Meiners M (2012) The functional equation of the smoothing transform. *Ann. Probab.* 40(5):2069–2105.

[3] Alsmeyer G, Damek E, Mentemeier S (2011) Precise tail index of fixed points of the two-sided smoothing transform. Löwe GAM, ed., *Random Matrices and Iterated Random Functions* (Münster, Germany: Springer Proceedings in Mathematics & Statistics).

[4] Alsmeyer G, Meiners M (2012) Fixed points of inhomogeneous smoothing transforms. *J. Diff. Equations Appl.* 18(8):1287–1304.

[5] Alsmeyer G, Meiners M (2013) Fixed points of the smoothing transform: Two-sided solutions. *Probab. Theory Relat. Fields* 155(1-2):165–199.

[6] Alsmeyer G, Rösler U (2008) A stochastic fixed point equation for weighted minima and maxima. *Annales de l'Institiut Henri Poincaré* 44(1):89–103.

[7] Asmussen S (2003) *Applied Probability and Queues*. Stochastic Modelling and Applied Probability (New York: Springer), 2nd edition.

[8] Baccelli F, Brémaud P (2003) *Elements of Queueing Theory*. Stochastic Modelling and Applied Probability (Berlin: Springer), 2nd edition.

[9] Baccelli F, Foss S (2011) Poisson hail on a hot ground. *J. Appl. Probab.* 48A:343–366.

[10] Biggins J (1998) Lindley-type equations in the branching random walk. *Stochastic Process. Appl.* 75:105–133.

[11] Chen N, Olvera-Cravioto M (2015) Efficient simulation for branching linear recursions. *Proceedings of the Winter Simulation Conference 2015*, 2716–2727.

[12] Courcoubetis C, Reiman M, Simon B (1987) Stability of a queueing system with concurrent service and locking. *SIAM Journal on Computing* 16(1):169–178.

[13] Goldie C (1991) Implicit renewal theory and tails of solutions of random equations. *Ann. Appl. Probab.* 1(1):126–166.

[14] Green L (1980) A queueing system in which customers require a random number of servers. *Operations Research* 28(6):1335 – 1346.

[15] Harrison P, Zertal S (2003) Queueing models with maxima of service times. *Computer Performance Evaluation. Modelling Techniques and Tools. 13th International Conference, TOOLS 2003*, volume 2794 of *Lecture Notes in Computer Science*, 152–168 (Springer).

[16] Jelenković P, Olvera-Cravioto M (2012) Implicit renewal theory and power tails on trees. *Adv. Appl. Prob.* 44(2):528–561.

[17] Jelenković P, Olvera-Cravioto M (2012) Implicit renewal theory for trees with general weights. *Stochastic Process. Appl.* 122(9):3209–3238.

[18] Jelenković P, Olvera-Cravioto M (2015) Maximums on trees. *Stochastic Process. Appl.* 125:217–232.

[19] Kantorovich L, Rubinstein G (1958) On a space of completely additive functions. *Vestn. Leningr. Univ. Math.* 13(7):52–59.

[20] Karpelevich F, Kelbert M, Suhov Y (1994) Higher-order Lindley equations. *Stochastic Process. Appl.* 53:65–96.

[21] Kelly F (1987) One-dimensional circuit-switched networks. *Ann. Probab.* 15(3):1166–1179.

[22] Kulkarni V, Puryear L (1994) A reader-writer queue with reader preference. *Queueing Systems Theory Appl.* 15:81–97.

[23] Kulkarni V, Puryear L (1995) A reader-writer queue with alternating exhaustive priorities. *Queueing Systems Theory Appl.* 19:81–103.

[24] Lebrecht A, Dingle N, Knottenbelt W (2008) A response time distribution model for zoned raid. *15th International Conference, ASMTA 2008*, 144–157, Lecture Notes in Computer Science (Springer).

[25] Lebrecht A, Knottenbelt W (2007) Response time approximations in fork-join queues. *23rd Annual UK Performance Engineering Workshop (UKPEW)* .

[26] Mitra D (1985) Probabilistic models and asymptotic results for concurrent processing with exclusive and non-exclusive locks. *SIAM J. Comput.* 14(4):1030–1051.

[27] Mitra D, Weinberger P (1984) Probabilistic models of database locking: Solutions, computational algorithms, and asymptotics. *J. ACM* 31(4):855–878.

[28] Olvera-Cravioto M (2019) Convergence of the Population Dynamics algorithm in the Wasserstein metric. *Electronic Journal of Probability* 24(61):1–27.

[29] Reiman M, Wright P (1991) Performance analysis of concurrent-read exclusive-write. *ACM SIGMETRICS Performance Evaluation Review* 19(1):168–177.

[30] Rösler U (1993) The weighted branching process. *Dynamics of complex and irregular systems (Bielefeld, 1991)* 154–165, bielefeld Encounters in Mathematics and Physics VIII, World Science Publishing, River Edge, NJ.

[31] Rüshendorf L (1985) The Wasserstein distance and approximation theorems. *Z. Wahrscheinlichkeitstheorie verw. Gebiete* 70:117–129.

[32] Tsimashenka I, Knottenbelt W, Harrison P (2014) Controlling variability in split-merge systems and its impact on performance. *Ann. Oper. Res.* .

[33] Villani C (2009) *Optimal transport, old and new* (New York: Springer).

[34] Whitt W (1985) Blocking when service is required from several facilities simultaneously. *AT&T Technical Journal* 64(8):1807–1985.

[35] Whitt W (2002) *Stochastic-Process Limits* (New York: Springer).