



HAL
open science

Improving the performance of heterogeneous data centers through redundancy

Elene Anton, Urtzi Ayesta, M Jonckheere, Ina Maria Maaïke Verloop

► **To cite this version:**

Elene Anton, Urtzi Ayesta, M Jonckheere, Ina Maria Maaïke Verloop. Improving the performance of heterogeneous data centers through redundancy. 2020. hal-02491393v1

HAL Id: hal-02491393

<https://hal.science/hal-02491393v1>

Preprint submitted on 29 Feb 2020 (v1), last revised 9 Dec 2020 (v2)

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Improving the performance of heterogeneous data centers through redundancy

E. Anton^{1,3}, U. Ayesta^{1,2,3,4}, M. Jonckheere⁵ and I. M. Verloop^{1,3}

¹ CNRS, IRIT, 2 rue Charles Camichel, 31071 Toulouse, France

² IKERBASQUE - Basque Foundation for Science, 48011 Bilbao, Spain

³ Université de Toulouse, INP, 31071 Toulouse, France

⁴ UPV/EHU, University of the Basque Country, 20018 Donostia, Spain

⁵ Instituto de Cálculo - Conicet, Facultad de Ciencias Exactas y Naturales, Universidad de Buenos Aires (1428) Pabellón II, Ciudad Universitaria Buenos Aires, Argentina

ABSTRACT

We analyze the performance of redundancy in a multi-type job and multi-type server system. We assume the job dispatcher is unaware of the servers' capacities, and we set out to study under which circumstances redundancy improves the performance. With redundancy an arriving job dispatches redundant copies to all its compatible servers, and departs as soon as one of its copies completes service. As a benchmark comparison, we take the non-redundant system in which a job arrival is routed to only one randomly selected compatible server. Service times are generally distributed and all copies of a job are identical, i.e., have the same service requirement.

In our first main result, we characterize the sufficient and necessary stability conditions of the redundancy system. This condition coincides with that of a system where each job type only dispatches copies into its least-loaded servers, and those copies need to be fully served. In our second result, we compare the stability regions of the system under redundancy to that of no redundancy. We show that if the server's capacities are sufficiently heterogeneous, the stability region under redundancy can be much larger than that without redundancy. We apply the general solution to particular classes of systems, including redundancy- d and nested models, to derive simple conditions on the degree of heterogeneity required for redundancy to improve the stability. As such, our result is the first in showing that redundancy can improve the stability and hence performance of a system when copies are *non-i.i.d.*

1 INTRODUCTION

The main motivation of studying redundancy models comes from the fact that both empirical ([1, 2, 8, 26]) and theoretical ([10, 12, 17, 19, 20, 25]) evidence show that redundancy might improve the performance of real-world applications. Under redundancy, a job that arrives to the system dispatches multiple copies into the servers, and departs when a first copy completes service. By allowing for redundant copies, the aim is to minimize the latency of the system by exploiting the variability in the queue lengths and the capacity of the different servers.

Most of the theoretical results on redundancy system consider the performance analysis when either FCFS or Processor-Sharing (PS) service policies are implemented in the servers. Under the assumption that all the copies of a job are i.i.d (independent and identically distributed) and exponentially distributed, [3, 5, 12] show that the stability condition of the system is independent of the number of redundant copies and that performance (in terms of delay and number of jobs in the system) improves as the number of copies increases. However, [10] showed that the assumption that copies of a job are i.i.d. can be unrealistic, and that it might lead to theoretical results that do not reflect the results of replication schemes in real-life computer systems. The latter has triggered interest to consider other modeling assumptions for the correlation structure of the copies of a job. For example, for identical copies (all the copies of a job have the same size), [3] showed that under both FCFS and PS service policies, the stability region of the system with *homogeneous* servers decreases as the number of copies increases.

The above observation provides the motivation for our study: to understand when redundancy is beneficial. In order to do so, we analyze a general multi-type job and multi-type server system. A dispatcher needs to decide to which server(s) to route each incoming job. We assume that the dispatcher is oblivious to the capacities of the servers in the system. The latter can be motivated by (i) design constraints, (ii) (slowly) fluctuating capacity of a server due to external users, or (iii) the impossibility of exchanging information among dispatchers and servers. The only information that is available to the dispatcher is the type of job and its set of compatible servers.

We consider two different models: the redundancy model where the dispatcher sends a copy to all the compatible servers of the job type, and the Bernoulli model where a single copy is sent to a uniformly selected compatible server of the job type. The comparison between these two policies is fair under the assumption that the dispatcher only knows the type of jobs and the set of compatible servers. Hence, we do not compare the performance of redundancy with other routing policies – such as Join the Shortest Queue, Join the Idle Server, Power of d , etc. – that have more information on the state of the system. We hence aim to understand when redundancy is beneficial for the performance of the system in this context. Observe that the answer is not clear upfront as adding redundant copies has two opposite effects: on the one hand, redundancy helps

exploiting the variability across servers' capacities, but on the other hand, it induces a waste of resources as servers work on copies that do not end up being completely served.

To answer the above question, we analyze the stability of an arbitrary multi-type job and multi-type server system with redundancy. Job service requirements are generally distributed, and copies are identical. The scheduling discipline implemented by servers is PS, which is a common policy in server farms and web servers, see for example [14, Chapter 24]. In our main result, we derive sufficient and necessary stability conditions for the redundancy system. This general result allows us to characterize when redundancy can increase the stability region with respect to Bernoulli routing.

To the best of our knowledge, our analytical results are the first showing that, when copies are non-i.i.d., adding redundancy to the system can be beneficial from the stability point of view. We believe that our result can motivate further research in order to thoroughly understand when redundancy is beneficial in other settings. For example, for different scheduling disciplines, different correlation structures among copies, different redundancy schemes, etc. We discuss this in more detail in Section 9.

We briefly summarize the main findings of the paper:

- The characterization of sufficient and necessary stability condition of any general redundancy system with heterogeneous server capacities and arrivals.
- We prove that when servers are heterogeneous enough (see Conditions in Section 6), redundancy has a larger stability region than Bernoulli.
- By exploring numerically these conditions, we observe that the degree of heterogeneity needed in the servers for redundancy to be better, decreases in the number of servers, and increases in the number of redundant copies.

The rest of the paper is organized as follows. In Section 2 we discuss related work. Section 3 describes the model, and introduces the notion of *capacity-to-fraction-of-arrivals ratio* that plays a key role in the stability result. Section 4 gives an illustrative example in order to obtain intuition about the structure of the stability conditions. Section 5 states the stability condition for the redundancy model. Section 6 provides conditions on the heterogeneity of the system under which redundancy outperforms Bernoulli. The proof of the main result is given in Section 7. Simulations are given in Section 8, and concluding remarks are given in Section 9. For the sake of readability, proofs are deferred to the Appendix.

2 RELATED WORK

When copies of a job are i.i.d. and exponentially distributed, [5, 12] have shown that redundancy with FCFS employed in the servers does not reduce the stability region of the system. In this case, the stability condition is that for any subset of job types, the sum of the arrival rates must be smaller than the sum of service rates associated with these job types. In [23], the authors consider i.i.d. copies with highly variable service time distributions. They focus on redundancy- d systems where each job chooses a subset of d homogeneous servers uniformly at random. The authors show that with FCFS, the stability region increases (without bound) in both the number of copies, d , and in the parameter that describes the variability in service times.

In [18], the authors investigate when it is optimal to replicate a job. They show that for so-called New-Worse-Than-Used service time distributions, the best policy is to replicate as much as possible. In [11], the authors investigate the impact that scheduling policies have on the performance of so-called nested redundancy systems with i.i.d. copies. The authors show that when FCFS is implemented, the performance might not improve as the number of redundant copies increases, while under other policies proposed in the paper, such as Least-redundant-first or Primaries-first, the performance improves as the number of copies increases.

Anton et al. [3] study the stability conditions when the scheduling policies PS, Random Order of Service (ROS) or FCFS are implemented. For the redundancy- d model with homogeneous server capacities and i.i.d. copies, they show that the stability region is not reduced if either PS or Random Order of Service (ROS) is implemented. When instead copies belonging to one job are identical, [3] showed that (i) ROS does not reduce the stability region, (ii) FCFS reduces the stability region and (iii) PS dramatically reduces the stability region, and this coincides with the stability region of system where all copies need to be *fully* served, i.e., $\lambda < \frac{\mu K}{d}$. In [24], the authors show that the stability result for PS extends to generally distributed service times. In addition, they obtain the stability condition for general correlation structures among copies when studying the homogeneous redundancy- d model.

Hellems et al. [16] consider identical copies that are generally distributed. For a redundancy- d model with FCFS, they develop a numerical method to compute the workload and response time distribution when the number of servers tends to infinity, i.e., the mean-field regime. The authors can numerically infer whether the system is stable, but do not provide any characterization of the stability region. In a recent paper, Hellems et al. [15] extend this study to include many replication policies, and general correlation structure among the copies.

Gardner et al. [10] introduce a new dependency structure among the copies of a job, the S&X model. The service time of each copy of a job is decoupled into two components: one related to the inherent job size of the task, that is identical for all the copies of a job, and the other one related to the server's slowdown, which is independent among all copies. The paper proposes and analyzes the redundant-to-idle-queue scheme with homogeneous servers, and proves that it is stable, and performs well.

To the best of our knowledge, no analytical results were obtained so far for performance measures when servers are heterogeneous *and* copies are identical or of any other non i.i.d. structure.

3 MODEL DESCRIPTION

We consider a K parallel-server system with heterogeneous capacities μ_k , for $k = 1, \dots, K$. Each server has its own queue, where Processor Sharing (PS) service policy is implemented. We denote by $S = \{1, \dots, K\}$ the set of all servers.

Jobs arrive to the system according to a Poisson process of rate λ . Each job is labelled with a type c that represents the subset of compatible servers to which type- c jobs can be sent: i.e., $c = \{s_1, \dots, s_n\}$, where $n \leq K$, $s_1, \dots, s_n \in S$ and $s_i \neq s_l$, for all $i \neq l$. We denote by C the set of all types in the system, i.e., $C = \{c \in \mathcal{P}(S) : p_c > 0\}$, where $\mathcal{P}(S)$ contains all the possible subsets of

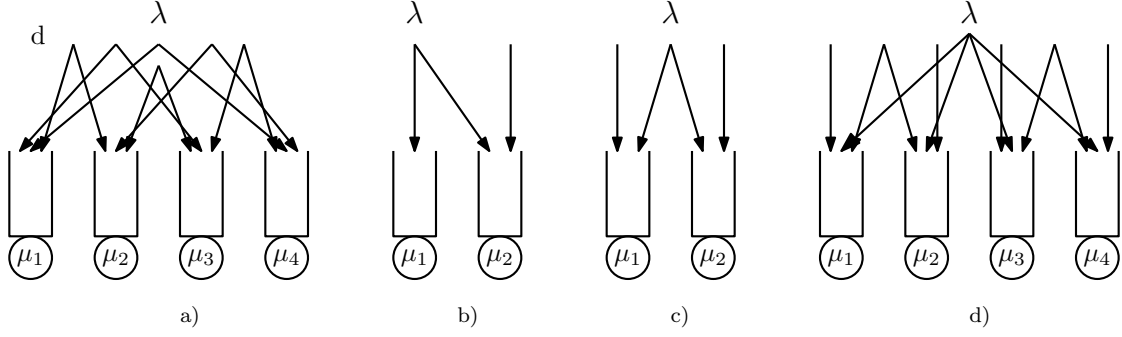


Figure 1: From left to right, the redundancy- d model (for $K = 4$ and $d = 2$), the N -model, the W -model and the WW -model.

S . Furthermore, we denote by $C(s)$ the subset of types that have server s as compatible server, that is, $C(s) = \{c \in C : s \in c\}$. A job is with probability p_c of type c , where $\sum_{c \in C} p_c = 1$.

In this paper, we consider two load balancing policies, which determine how the jobs are dispatched to the servers. Note that both load balancers are oblivious to the capacities of the servers.

- Bernoulli routing: a type- c job is sent with uniform probability to one of its compatible servers in c .
- Redundancy model: a type- c job sends identical copies to its $|c|$ compatible servers. The job (and corresponding copies) departs the system when one of its copies completes service.

We consider that jobs have identical copies, i.e., all the copies of a job have exactly the same size. Job sizes are distributed according to a general random variable X with cumulative distribution function F and unit mean. Additionally, we assume that

- (1) F has no atoms.
- (2) F is a light tailed distribution in the following sense,

$$\lim_{r \rightarrow \infty} \sup_{a \geq 0} \mathbf{E}[(X - a)1_{\{X - a > r\}} | X > a] = 0.$$

This holds, e.g., for exponential, hyper-exponential or Erlang distributions [21].

In this paper, we will characterize the stability condition under both load balancing policies. Stability will be understood as positive Harris recurrence for the underlying Markov processes. We define λ^R as the value of λ such that the redundancy model is stable if $\lambda < \lambda^R$ and unstable if $\lambda > \lambda^R$. Similarly, we define λ^B . We aim to characterize when $\lambda^R > \lambda^B$, that is, when does redundancy improve the stability condition compared to no redundancy (Bernoulli routing).

For Bernoulli, λ^B can be easily found. The Bernoulli system reduces to K independent servers, where server s receives arrivals at rate $\lambda(\sum_{c \in C(s)} \frac{p_c}{|c|})$ and has a departure rate μ_s , for all $s \in S$. The stability condition is hence,

$$\lambda < \lambda^B = \min_{s \in S} \left\{ \frac{\mu_s}{\sum_{c \in C(s)} \frac{p_c}{|c|}} \right\}. \quad (1)$$

In order to characterize λ^R , we need to study the system under redundancy in more detail. For that, we denote by $N_c(t)$ the number of type- c distinct jobs that are present in the redundancy system at time t and $\vec{N}(t) = (N_c(t), c \in C)$. Furthermore, we denote the

number of copies per server by $M_s(t) := \sum_{c \in C(s)} N_c(t)$, $s \in S$, and $\vec{M}(t) = (M_1(t), \dots, M_K(t))$. For the j -th type- c job, let b_{cj} denote the service requirement of this job, for $j = 1, \dots, N_c(t)$, $c \in C$. Let $a_{cjs}(t)$ denote the attained service in server s of the j -th type- c job at time t . We denote by $A_c(t) = (a_{cjs}(t))_{js}$ a matrix on \mathbb{R}_+ of dimension $N_c(t) \times |c|$. Note that the number of type- c jobs increases by one at rate λp_c , which implies that a row composed of zeros is added to $A_c(t)$. When one element $a_{cjs}(t)$ in matrix $A_c(t)$ reaches the required service b_{cj} , the corresponding job departs and all of its copies are removed from the system. Hence, row j in matrix $A_c(t)$ is removed. We further let $\phi_s(\vec{M}(t))$ be the capacity that each of the copies in server s obtains when in state $\vec{M}(t)$, which under PS is given by, $\phi_s(\vec{M}(t)) := \frac{\mu_s}{M_s(t)}$. The cumulative service that a copy in server s gets during the time interval (v, t) is

$$\eta_s(v, t) := \int_{x=v}^t \phi_s(\vec{M}(x)) dx.$$

In order to characterize the stability condition of this system, we define the *capacity-to-fraction-of-arrivals ratio* of a server in a subsystem:

Definition 3.1 (Capacity-to-fraction-of-arrival ratio). For any given set of servers $\tilde{S} \subseteq S$ and its associated set of job types $\tilde{C} = \{c \in C : c \subseteq \tilde{S}\}$, the *capacity-to-fraction-of-arrival ratio* of server $s \in \tilde{S}$ in this so-called \tilde{S} -subsystem is defined by $\frac{\mu_s}{\sum_{c \in \tilde{C}(s)} p_c}$, where $\tilde{C}(s) = \tilde{C} \cap C(s)$ is the subset of types in \tilde{C} that are served in server s .

Some common models

A well-known structure is *redundancy- d* , see Figure 1 a). Within this model, each job has d out of K compatible servers, where d is fixed. That is, $p_c > 0$ for all $c \in \mathcal{P}(S)$ with $|c| = d$, and $p_c = 0$ otherwise, so that there are $\binom{K}{d}$ types of jobs. If additionally, $p_c = 1/\binom{K}{d}$ for all $c \in C$, we say that the arrival process of jobs is homogeneously distributed over types. We will call this model the *redundancy- d model with homogeneous arrivals*. The particular case where server capacities are also homogeneous, i.e. $\mu_k = \mu$ for all $k = 1, \dots, K$ will be called the *redundancy- d model with homogeneous arrivals and server capacities*.

Another structure is the *nested model*. Then, for all $c, c' \in C$, either $i) c \subset c'$ or $ii) c' \subset c$ or $iii) c \cap c' = \emptyset$. First of all, note that

the redundancy- d model does not fit in the nested structure. The smallest nested system is the so called N -model (Figure 1 b)): this is a $K = 2$ server system with types $C = \{\{2\}, \{1, 2\}\}$. Another nested system is the W -model (Figure 1 c)), that is, $K = 2$ servers and types $C = \{\{1\}, \{2\}, \{1, 2\}\}$. In Figure 1 d), a nested model with $K = 4$ servers and 7 different jobs types, $C = \{\{1\}, \{2\}, \{3\}, \{4\}, \{1, 2\}, \{3, 4\}, \{1, 2, 3, 4\}\}$ is given. This model is referred to as the WW -model.

4 AN ILLUSTRATIVE EXAMPLE

Before formally stating the main results in Section 5.1, we first illustrate through a numerical example some of the key aspects of our proof, and in particular the essential role played by the capacity-to-fraction-of-arrival ratio defined in Definition 3.1. In Figure 2 we plot the trajectories of the number of copies per server with respect to time for a $K = 4$ redundancy-2 system (Figure 1), that is $C = \{\{1, 2\}, \{1, 3\}, \{1, 4\}, \{2, 3\}, \{2, 4\}, \{3, 4\}\}$. Our proof techniques will rely on fluid limits, and therefore we chose large initial points. Figures 2 a) and b) show the trajectories when servers and arrivals of types are homogeneous for $\lambda = 1.8$ and $\lambda = 2.1$, respectively. Figures 2 c) and d) consider a heterogeneous system (parameters see the legend) for $\lambda = 7.5$ and $\lambda = 9$, respectively.

The homogeneous example (Figure 2 a) and b)) falls within the scope of [3]. There it is shown that the stability condition is $\lambda < \frac{\mu K}{d}$. We note that this condition coincides with the stability condition of a system in which all the d copies need to be fully served. In Figure 2 a) and b), the value for λ is chosen such that they represent a stable and an unstable system, respectively. As formally proved in [3], at the fluid scale, when the system is stable the largest queue length decreases, whereas in the unstable case the minimum queue length increases. It thus follows, that in the homogeneous case, either *all classes* are stable, or unstable.

The behavior of the heterogeneous case is rather different. The parameters corresponding to Figures 2 c) and d) are such that the system is stable in c), but not in d). In Figure 2 c) we see that the trajectories of all queue lengths are not always decreasing, including the maximum queue length. In Figure 2 d), we observe that the number of copies in servers 3 and 4 are decreasing, whereas those of servers 1 and 2 are increasing.

When studying stability for the heterogeneous setting, one needs to reason recursively. First, assume that each server s needs to handle its full load, i.e., $\lambda \frac{\sum_{c \in C(s)} p_c}{\mu_s}$. Hence, one can simply compare the servers capacity-to-fraction-of-arrival ratios, $\mu_s / \sum_{c \in C(s)} p_c$, to see which server could potentially empty first. In this example, server 4 has the maximum capacity-to-fraction-of-arrival ratio, and, in fluid scale, will reach zero in finite time, and remain zero, since $\lambda < \mu_4 / \sum_{c \in C(4)} p_c = 5 / (p_{\{1,4\}} + p_{\{2,4\}} + p_{\{3,4\}}) = 11.11$ is larger than the arrival rate $\lambda = 7.5$.

Whenever, at fluid scale, server 4 is still positive, the other servers might either increase or decrease. However, the key insight is that once the queue length of server 4 reaches 0, the fluid behavior of the other classes no longer depend on the jobs that also have server 4 as compatible server. That is, we are sure that all jobs that have server 4 as compatible server, will be fully served in server 4, since server 4 is in fluid scale empty and all the other servers are overloaded. Therefore, jobs with server 4 as compatible server can

be ignored, and we are left with a subsystem formed by servers $\{1, 2, 3\}$ and without the job types served by server 4. Now again, we consider the maximum capacity-to-fraction-of-arrival ratio in order to determine the least-loaded server, but now for the subsystem $\{1, 2, 3\}$. This time, server 3 has the maximum capacity-to-fraction-of-arrival ratio, which is $4 / (p_{\{1,3\}} + p_{\{2,3\}}) = 10$. Since this value is larger than $\lambda = 7.5$, it is a sufficient condition for server 3 to empty.

Similarly, once server 3 is empty, we consider the subsystem with servers 1 and 2 only. Hence, there is only one type of jobs, $\{1, 2\}$. Now server 2 is the least-loaded server and its capacity-to-fraction-of-arrival ratio is $2 / p_{\{1,2\}} = 8$. This value being larger than the arrival rate, implies that server 2 (and hence server 1, because there is only one job type) will be stable too. Indeed, in the figure we also observe that once server 3 hits zero, both server 1 and server 2 are decreasing.

We can now explain the evolution observed in Figure 2 d) when $\lambda = 9$. The evolution for servers 4 and 3 can be argued as before: both their capacity-to-fraction-of-arrival ratios are larger than $\lambda = 9$, hence they empty in finite time. However, the capacity-to-fraction-of-arrival ratio of the subsystem with servers 1 and 2, which is 8, is strictly smaller than the arrival rate. We thus observe that, unlike in the homogeneous case, in the heterogeneous case some servers might be stable, while others (here server 1 and 2) are unstable.

Proposition 5.1 formalizes the above intuitive explanation, by showing that the stability of the system can be derived recursively.

The capacity-to-fraction-of-arrival ratio allows us now to reinterpret the homogeneous case depicted in Figure 2 a) and b). In this case, the capacity-to-fraction-of-arrival ratio of all the servers is the same, which implies (i) that either all servers will be stable, or all unstable, and (ii) from the stability viewpoint is as if all copies received service until completion.

5 STABILITY CONDITION

5.1 Multi-type job multi-type server system

In this section we discuss the stability condition of the general redundancy system with PS. In order to do so, we first define several sets of subsystems, similar to as what we did in the illustrative example of Section 4.

The first subsystem includes all servers, that is $S_1 = S$. We denote by \mathcal{L}_1 the set of servers with highest capacity-to-fraction-of-arrival ratio in the system $S_1 = S$. Thus,

$$\mathcal{L}_1 = \left\{ s \in S_1 : s = \arg \max_{\tilde{s} \in S_1} \left\{ \frac{\mu_{\tilde{s}}}{\sum_{c \in C} p_c} \right\} \right\}.$$

For $i = 2, \dots, K$, we define recursively

$$\begin{aligned} S_i &:= S \setminus \cup_{l=1}^{i-1} \mathcal{L}_l, \\ C_i &:= \{c \in C : c \subset S_i\}, \\ C_i(s) &:= C_i \cap C(s), \end{aligned}$$

$$\mathcal{L}_i := \left\{ s \in S_i : s = \arg \max_{\tilde{s} \in S_i} \left\{ \frac{\mu_{\tilde{s}}}{\sum_{c \in C_i(\tilde{s})} p_c} \right\} \right\}.$$

The S_i -subsystem will refer to the system consisting of the servers in S_i , with only jobs of types in the set C_i . The $C_i(s)$ is the subset of types that are served in server s in the S_i -subsystem. We let $C_i(s) = C$. The \mathcal{L}_i represents the set of servers s with highest capacity-to-fraction-of-arrival ratio in the S_i -subsystem, or in other

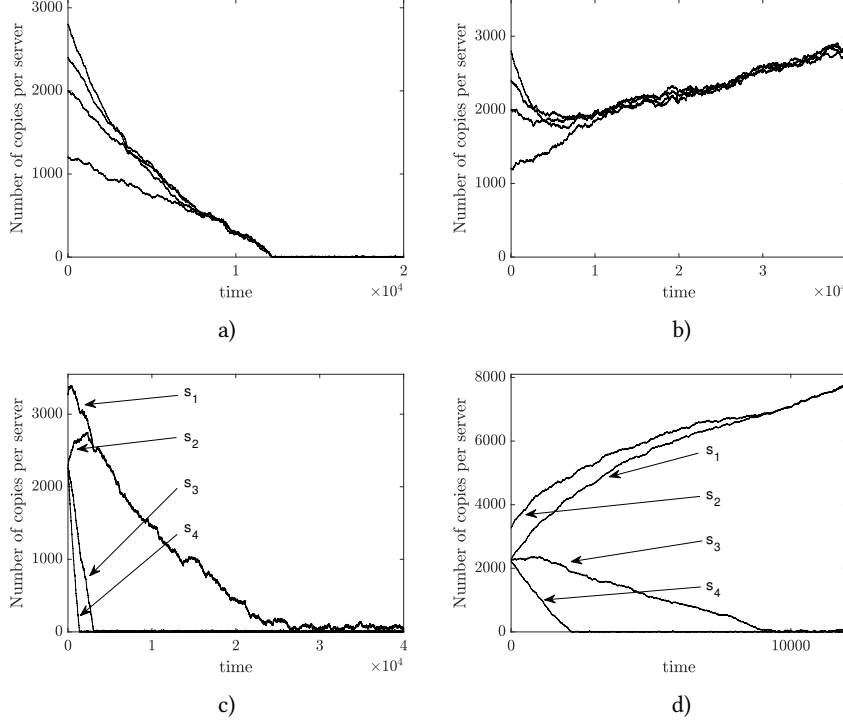


Figure 2: Trajectory of the number of copies per server with respect to time for a $K = 4$ redundancy-2 system with exponentially distributed job sizes. Figures a) and b) consider homogeneous capacities $\mu_k = 1$ for $k = 1, \dots, 4$ and homogeneous arrival rates per type, $p_c = 1/6$ for all $c \in C$, with a) $\lambda = 1.8$ and b) $\lambda = 2.1$. Figures c) and d) consider heterogeneous server capacities $\mu = (1, 2, 4, 5)$ and arrival rates per type $\vec{p} = \{0.25, 0.1, 0.1, 0.2, 0.2, 0.15\}$ respectively for types C , c) with $\lambda = 7.5$ and d) $\lambda = 9$.

words, the least-loaded servers in the S_i subsystem. Finally, we denote by $i^* := \arg \max_{i=1, \dots, K} \{C_i : C_i \neq \emptyset\}$ the last index i for which the subsystem S_i is not empty of job types.

Before continuing, we illustrate the above definitions using the illustrative example of Section 4. There, the first subsystem consists of servers $S_1 = \{1, 2, 3, 4\}$ and all job types, see Figure 3 a). The capacity-to-fraction-of-arrival ratios in the S_1 subsystem are: $\{2.2, 3.07, 8.8, 11.1\}$, for servers in the set S_4 . Hence $\mathcal{L}_1 = \{4\}$. The second subsystem is composed of servers $S_2 = \{1, 2, 3\}$ and all job types that have a copy in server 4 can be ignored, that is, $C_2 = \{\{1, 2\}, \{1, 3\}, \{2, 3\}\}$, see Figure 3 b). The capacity-to-fraction-of-arrival ratios for servers in the S_2 subsystem are given by $\{2.8, 4.4, 10\}$, and thus $\mathcal{L}_2 = \{3\}$. The third subsystem consists of servers $S_3 = \{1, 2\}$ and all job types that have a copy in servers 3 or 4 can be ignored, that is, $C_3 = \{\{1, 2\}\}$, see Figure 3 c). The capacity-to-fraction-of-arrival ratios for servers in the S_3 subsystem are given by $\{4, 8\}$. Hence, $\mathcal{L}_3 = \{2\}$. Then, $S_4 = \{1\}$, but $C_4 = \emptyset$, so that $i^* = 3$.

The value of the highest capacity-to-fraction-of-arrival ratio in the S_i -subsystem is denoted by

$$CAR_i := \max_{\bar{s} \in S_i} \left\{ \frac{\mu_{\bar{s}}}{\sum_{c \in C_i(\bar{s})} p_c} \right\}, \text{ for } i = 1, \dots, i^*.$$

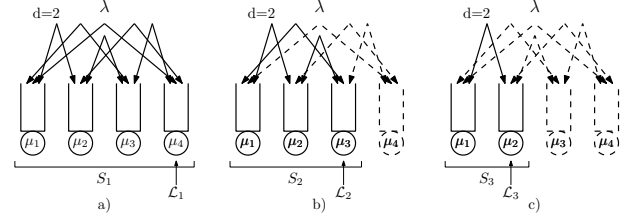


Figure 3: $K = 4$ server system under redundancy-2. In bold in a) subsystem S_1 , in b) subsystem S_2 and in c) subsystem S_3 .

Note that $CAR_i = \frac{\mu_s}{\sum_{c \in C_i(s)} p_c}$, for any $s \in \mathcal{L}_i$. Furthermore, we note that in a given S_i -subsystem, the server with highest capacity-to-fraction-of-arrivals ratio and the least-loaded one are actually the same.

In the following proposition we characterize the stability condition for servers in terms of the capacity-to-fraction-of-arrival ratio corresponding to each subsystem. It states that servers that have highest capacity-to-fraction-of-arrival ratio in subsystem S_i can only be stable if all servers in S_1, \dots, S_{i-1} are stable as well. The proof can be found in Section 7.

PROPOSITION 5.1. For a given $i \leq i^*$, servers $s \in \mathcal{L}_i$ are stable if $\lambda < CAR_l$, for all $l = 1, \dots, i$. Servers $s \in \mathcal{L}_i$ are unstable if there is an $l = 1, \dots, i$ such that $\lambda > CAR_l$.

COROLLARY 5.2. The redundancy system is stable if $\lambda < CAR_i$, for all $i = 1, \dots, i^*$. The redundancy system is unstable if there exists an $i \in \{1, \dots, i^*\}$ such that $\lambda > CAR_i$.

We note that CAR_l , $l = 1, \dots, i$, are not necessarily ordered with respect to l .

We now write an equivalent representation of the stability condition (proof see Appendix). Denote by $\mathcal{R}(c)$ the set of servers where job type c achieves maximum capacity-to-fraction-of-arrival ratio, or in other words, the set of least-loaded servers for type c :

$$\mathcal{R}(c) := \{s : \exists i, \text{ s.t. } c \in C_i(s) \text{ and } s \in \mathcal{L}_i\}.$$

Note that there is a unique subsystem S_i for which this happens, i.e., $\mathcal{R}(c) \subseteq \mathcal{L}_i$ for exactly one i . We note that for a type- c job, if c contains at least a server that was removed in the i th iteration, then $\mathcal{R}(c) \subseteq \mathcal{L}_i$. We further let $\mathcal{R} := \cup_{c \in \mathcal{C}} \mathcal{R}(c)$.

COROLLARY 5.3. The redundancy system is stable if $\lambda \sum_{c: s \in \mathcal{R}(c)} p_c < \mu_s$, for all $s \in \mathcal{R}$. The redundancy system is unstable if there exists an $s \in \mathcal{R}$ such that $\lambda \sum_{c: s \in \mathcal{R}(c)} p_c > \mu_s$.

From the above corollary, we directly observe that the stability condition for the redundancy system coincides with the stability condition corresponding to K individual servers where each job type c is only dispatched to its least-loaded servers.

5.2 Particular redundancy structures

In this subsection we discuss the stability condition for some particular cases of redundancy: redundancy- d and nested systems.

Redundancy- d

We focus here on the redundancy- d structure (defined in Section 3) with homogeneous arrivals, i.e. $p_c = \frac{1}{\binom{K}{d}}$ for all $c \in \mathcal{C}$.

We first assume that servers capacities are homogeneous, $\mu_k = \mu$ for all k . Since arrivals are homogeneous as well, the arrival rate to each server is $\lambda d/K$, thus the capacity-to-fraction-of-arrival ratio at every server is $\mu K/d$. This implies that $\mathcal{L}_1 = S$, $i^* = 1$ and $\mathcal{R}(c) = c$ for all $c \in \mathcal{C}$. From Corollary 5.2, we obtain that the system is stable if $\lambda d < \mu K$, which coincides with [3].

For heterogeneous servers capacities, we have the following:

COROLLARY 5.4. Under redundancy- d with homogeneous arrivals and $\mu_1 < \dots < \mu_K$, the system is stable if for all $i = d, \dots, K$, $\lambda \frac{\binom{i-1}{d-1}}{\binom{K}{d}} < \mu_i$. The system is unstable if there exists $i \in \{d, \dots, K\}$ such that $\lambda \frac{\binom{i-1}{d-1}}{\binom{K}{d}} > \mu_i$.

In the homogeneous case, it is easy to deduce that the stability condition, $\lambda d < \mu K$, decreases as d increases. However, in the heterogeneous case, both the numerator and denominator are non-monotone functions of d , and as a consequence it is not straightforward to deduce how the stability condition depends on d .

Nested systems

In this section we consider two nested redundancy systems.

5.2.1 N -model. The simplest nested model is the N -model. This is a $K = 2$ server system with capacities $\vec{\mu} = \{\mu_1, \mu_2\}$ and types $\mathcal{C} = \{\{2\}, \{1, 2\}\}$, see Figure 1 b). A job is of type $\{2\}$ with probability p and of type $\{1, 2\}$ with probability $1 - p$. The stability condition is $\lambda < \lambda^R$ where:

$$\lambda^R = \begin{cases} \mu_2, & 0 \leq p \leq \frac{\mu_2 - \mu_1}{\mu_2} \\ \mu_1/(1-p), & \left(\frac{\mu_2 - \mu_1}{\mu_2}\right)^+ \leq p \leq \frac{\mu_2}{\mu_1 + \mu_2} \\ \mu_2/p, & \frac{\mu_2}{\mu_1 + \mu_2} < p \leq 1 \end{cases}$$

The above is obtained as follows: The capacity-to-fraction-of-arrival ratio of the system is $\mu_1/(1-p)$ and μ_2 , respectively for servers 1 and 2. First assume $\mu_1/(1-p) > \mu_2$. Then $\mathcal{L}_1 = \{1\}$ and the second subsystem is composed of server $S_2 = \{2\}$ and $\mathcal{C}_2 = \{\{2\}\}$, with arrival rate λp to server 2. Hence the capacity-to-fraction-of-arrival ratio of server 2 in the S_2 -subsystem is μ_2/p . From Corollary 5.2, it follows that $\lambda^R = \min\{\mu_1/(1-p), \mu_2/p\}$. On the other hand, if $\mu_1/(1-p) < \mu_2$, then $\mathcal{L}_1 = \{2\}$, and $S_2 = \{1\}$, but $\mathcal{C}_2 = \emptyset$. Thus, $\lambda^R = \mu_2$. Lastly, if $\mu_1/(1-p) = \mu_2$, $\mathcal{L}_1 = \{1, 2\}$, thus $S_2 = \emptyset$ and $\mathcal{C}_2 = \emptyset$. Hence, $\lambda^R = \mu_2$.

We observe that the stability condition λ^R , is a continuous function with as maximum $\lambda^R = \mu_1 + \mu_2$ when $p = \mu_2/(\mu_1 + \mu_2)$. That is, for certain value of p it achieves the maximum stability condition. Note however that in this paper our focus is not on finding the best redundancy probabilities. Instead, in our model, we are given the probabilities p_c , which are determined by the characteristics of the job types and matchings, and we investigate whether or not the system can benefit from redundancy.

5.2.2 W -model. The W -model is a $K = 2$ server system with capacities $\vec{\mu} = \{\mu_1, \mu_2\}$ and types $\mathcal{C} = \{\{1\}, \{2\}, \{1, 2\}\}$, see Figure 1 c). A job is of type $\{1\}$ with probability $p_{\{1\}}$, type $\{2\}$ with probability $p_{\{2\}}$ and of type $\{1, 2\}$ with probability $p_{\{1,2\}}$. The stability condition is then given by:

$$\lambda^R = \begin{cases} \mu_2/(1-p_{\{1\}}), & p_{\{1\}} \leq \frac{\mu_1}{\mu_1 + \mu_2} \\ \mu_1/p_{\{1\}}, & p_{\{1\}} \geq \frac{\mu_1}{\mu_1 + \mu_2} \end{cases}$$

if $(1-p_{\{2\}})/\mu_1 > (1-p_{\{1\}})/\mu_2$ (i.e., the load on server 1 is larger than on server 2). And,

$$\lambda^R = \mu_2/(1-p_{\{1\}})$$

if $\mu_1(1-p_{\{1\}}) = \mu_2(1-p_{\{2\}})$. Similar to the N -model, the above can be obtained from Corollary 5.2. When $p = \mu_1/(\mu_1 + \mu_2)$, maximum stability $\lambda^R = \mu_1 + \mu_2$ is obtained.

6 WHEN DOES REDUNDANCY IMPROVE STABILITY

In this section, we compare the stability condition of the general redundancy system to that of the Bernoulli routing.

From Corollary 5.2, it follows that $\lambda^R = \min_{i=1, \dots, i^*} CAR_i$. Together with (1), we obtain the following sufficient and necessary conditions for redundancy to improve the stability condition.

COROLLARY 6.1. The stability condition under redundancy is larger than under Bernoulli routing if and only if

$$\min_{i=1, \dots, i^*, s \in \mathcal{L}_i} \left\{ \frac{\mu_s}{\sum_{c \in C_i(s)} p_c} \right\} \geq \min_{s \in S} \left\{ \frac{\mu_s}{\sum_{c \in C(s)} \frac{p_c}{|c|}} \right\}.$$

From inspecting the condition of Corollary 6.1, it is not clear upfront when redundancy would be better than Bernoulli. In the rest of the section, by applying Corollary 6.1 to redundancy- d and nested models, we will show that when the capacities of the servers are sufficiently heterogeneous, the stability of redundancy is larger than that of Bernoulli. In addition, numerical computations allow us to conclude that the degree of heterogeneity needed in the servers in order for redundancy to be beneficial, decreases in the number of servers, and increases in the number of redundant copies.

6.1 Redundancy- d

In this section, we compare the stability condition of the redundancy- d model with homogeneous arrivals to that of Bernoulli routing. From (1), we obtain that

$$\lambda^B = d \min_{i=1, \dots, K} \left\{ \frac{\mu_i}{\sum_{c \in C(s)} p_c} \right\} = K \min_{i=1, \dots, K} \mu_k. \quad (2)$$

From Corollary 5.4, we obtain that $\lambda^R = \min_{i=d, \dots, K} \left\{ \frac{\binom{K}{d}}{\binom{i-1}{d-1}} \mu_i \right\}$.

The following corollary is straightforward.

COROLLARY 6.2. *Assume $\mu_1 < \dots < \mu_K$. The system under redundancy- d and homogeneous arrivals has a strictly larger stability condition than the system under Bernoulli routing if and only if*

$$K\mu_1 < \min_{i=d, \dots, K} \left\{ \frac{\binom{K}{d}}{\binom{i-1}{d-1}} \mu_i \right\}.$$

The following is straightforward, since $\binom{i-1}{d-1}$ is increasing in i .

COROLLARY 6.3. *Assume $\mu_1 < \dots < \mu_K$ and homogeneous arrivals. The system under redundancy- d has a larger stability region than the Bernoulli routing if $\mu_1 d < \mu_d$.*

Hence, if there exists a redundancy parameter d such that $\mu_1 d < \mu_d$, then adding d redundant copies to the system improves its stability region. In that case, the stability condition of the system will improve by at least a factor $\frac{\mu_d}{d\mu_1}$.

In Table 1, we analyze how the heterogeneity of the server capacities impacts the stability of the system. We chose $\mu_k = \mu^{k-1}$, $k = 1, \dots, K$, so that the minimum capacity equals 1. Hence, for Bernoulli, $\lambda^B = K$. Under redundancy we have the following: For $\mu = 1$ the system is a redundancy- d system with homogeneous arrivals and server capacities, so that $\lambda^R = K/d$, [3]. Thus, $\lambda^R < \lambda^B$ in that case. We denote by μ^* the value of μ for which the stability region of the redundant system coincides with that of Bernoulli routing, i.e., the value of μ such that $\lambda^R = \lambda^B$. For $\mu < \mu^*$ (gray area in Table 1), Bernoulli has a larger stability region, while for $\mu > \mu^*$, redundancy outperforms Bernoulli.

First we observe that, for a fixed d , μ^* decreases as K increases, and is always less than $\mu = 2$. Therefore, as the number of servers increases, the level of heterogeneity that is needed in the servers in order to improve the stability under redundancy decreases. Second, for fixed K , we also observe that μ^* increases as d increases. This means that as the number of redundant copies d increases, the server capacities need to be more heterogeneous in order to improve the stability region under redundancy. Finally, focusing on the numbers in bold, we observe that when the number of servers K is

large enough and the servers are heterogeneous enough (large μ), the stability region increases in the number of redundant copies d .

Table 1: The maximum arrival rates λ^R and λ^B in a redundancy- d system with homogeneous arrivals and capacities $\mu_k = \mu^{k-1}$.

		$\mu = 1$	$\mu = 1.2$	$\mu = 1.4$	$\mu = 2$	$\mu = 3$	μ^*
$K = 3$	Red-2	1.5	2.16	2.94	6	9	1.41
	BR	3	3	3	3 3	3	
$K = 4$	Red-2	2	3.45	5.48	12	18	1.26
	BR	4	4	4	4	4	
$K = 5$	Red-2	2.5	5.18	9.14	20	30	1.19
	BR	5	5	5	5	5	
$K = 10$	Red-2	5	22.39	41.16	90	135	1.08
	BR	10	10	10	10	10	
$K = 4$	Red-3	1.33	2.30	3.65	10.66	36	1.44
	BR	4	4	4	4	4	
$K = 5$	Red-3	1.66	3.45	6.40	26.66	90	1.31
	BR	5	5	5	5	5	
$K = 10$	Red-3	3.33	17.19	60.23	320	1080	1.13
	BR	10	10	10	10	10	

In Table 2, we consider linearly increasing capacities on the interval $[1, M]$, that is $\mu_k = 1 + \frac{M-1}{K-1}(k-1)$, for $k = 1, \dots, K$. The grey area is where Bernoulli outperforms redundancy. For this specific system, the following corollary is straightforward.

COROLLARY 6.4. *Under a redundancy- d system with homogeneous arrivals and capacities $\mu_k = 1 + \frac{M-1}{K-1}(k-1)$, for $k = 1, \dots, K$, the redundancy system has stability condition: $\lambda^R = \frac{MK}{d}$, for $d > 1$, while $\lambda^B = K$. Hence, the redundancy system outperforms the stability condition of the Bernoulli routing if and only if $M \geq d$.*

Simple qualitative rules can be deduced. If $M \geq d$, redundancy is a factor M/d better than Bernoulli. Hence, increasing M , that is, the heterogeneity among the servers, is significantly beneficial for the redundancy system. However, the stability condition of the redundancy system degrades as the number of copies d increases.

Table 2: The maximum arrival rates λ^R and λ^B in a redundancy- d system with homogeneous arrivals and capacities $\mu_k = 1 + \frac{M-1}{K-1}(k-1)$.

		$M = 1$	$M = 2$	$M = 3$	$M = 4$	$M = 6$
$K = 3$	Red-2	1.5	3	4.5	6	9
	BR	3	3	3	3	3
$K = 4$	Red-2	2	4	6	8	12
	BR	4	4	4	4	4
$K = 5$	Red-2	2.5	5	7.5	10	15
	BR	5	5	5	5	5
$K = 10$	Red-2	5	10	15	20	30
	BR	10	10	10	10	10
$K = 4$	Red-3	1.33	2.66	4	5.33	8
	BR	4	4	4	4	4
$K = 5$	Red-3	1.66	3.33	5	6.66	10
	BR	5	5	5	5	5
$K = 10$	Red-3	3.33	6.66	10	13.33	20
	BR	10	10	10	10	10

6.2 Nested systems

6.2.1 *N-model*. The stability condition of the N -model with Bernoulli routing is given by the following expression:

$$\lambda^B = \begin{cases} 2 \min\{\mu_1, \mu_2\}, & p = 0 \\ 2\mu_1/(1-p), & 0 \leq p \leq \left(\frac{\mu_2 - \mu_1}{\mu_1 + \mu_2}\right)^+ \\ 2\mu_2/(1+p), & \left(\frac{\mu_2 - \mu_1}{\mu_1 + \mu_2}\right)^+ < p \leq 1 \end{cases}$$

The above set of conditions is obtained from the fact that under Bernoulli routing, $\lambda^B = \min\{\mu_1/(1-p), \mu_2/(p + \frac{1}{2}(1-p))\}$. Note that λ^B is a continuous function with a maximum $\mu_1 + \mu_2$ at the point $p = \frac{\mu_2 - \mu_1}{\mu_1 + \mu_2}$. Now, comparing λ^B to λ^R as obtained in Section 5.2.1 leads to the following:

COROLLARY 6.5. *Given an N -model. The stability condition under redundancy is larger than under Bernoulli routing under the following conditions: If $\mu_2 \leq \mu_1$, then $p \in \left(\left(\frac{2\mu_2 - \mu_1}{2\mu_2 + \mu_1}\right)^+, 1\right)$. If $\mu_2 > \mu_1$, then $p \in \left(0, \left(\frac{\mu_2 - 2\mu_1}{\mu_2}\right)^+\right) \cup \left(\frac{2\mu_2 - \mu_1}{2\mu_2 + \mu_1}, 1\right)$.*

From the above we conclude that if μ_1 is larger than $2\mu_2$, then redundancy is always better than Bernoulli, independent of the arrival rates of job types. For the case $\mu_2 > \mu_1$, we observe that for μ_2 large enough, redundancy will outperform Bernoulli.

6.2.2 *W-based nested systems*. We consider the following structure of nested systems: W (see Figure 1 c), WW (Figure 1 d) and $WWWW$. The latter is a $K = 8$ server system that is composed of 2 WW models and an additional job type $c = \{1, \dots, 8\}$ for which all servers are compatible. For all three models, we assume that a job is with prob. $p_c = 1/|C|$ of type c .

In Table 3, we analyze how heterogeneity in the server capacities impacts the stability. First of all, note that $\lambda^B = K$. For redundancy, the value of λ^R depends on the server capacities. In the upper part of the table, we let $\mu_k = \mu^{k-1}$ for $k = 1, \dots, K$. We denote by μ^* the value of μ for which $\lambda^R = \lambda^B$. We observe that as the number of servers duplicate, the μ^* decreases, and is always smaller than 1.5. So that, as the number of servers increases, the level of heterogeneity that is needed in order for redundancy to outperform Bernoulli decreases too.

In the second part of the table we assume $\mu_k = 1 + \frac{M-1}{K-1}(k-1)$ for $k = 1, \dots, K$. We observe that when $M \geq K$ the stability condition under redundancy equals $\lambda^R = |C|$, which is always larger than $\lambda^B = K$. However, as the number of servers increases, the maximum capacity of the servers, M , needs to increase M in order for redundancy to outperform Bernoulli.

7 PROOF OF THE STABILITY CONDITION

In the following, we outline the proof that the condition in Proposition 5.1 is sufficient and necessary for the respective subsystems to be stable. As we observe in Section 4, there are two main issues concerning the evolution of redundancy systems with heterogeneous capacities. First of all, the trajectory of a particular server decreases, only if a certain subset of servers is already in steady state. Secondly, for a particular server $s \in S$, the instantaneous departure of that server might be larger than μ_s due to copies leaving in servers other than s . This makes the dynamics of the system complex. In order to determine the stability conditions, we construct upper and

Table 3: The maximum arrival rates λ^R and λ^B in nested systems.

$\mu_k = \mu^{k-1}$		$\mu = 1$	$\mu = 1.2$	$\mu = 1.4$	$\mu = 2$	μ^*
$K = 2$	W-model BR	1.5 2	1.8 2	2.10 2	3 2	1.33
$K = 4$	WW-model BR	2.33 4	4.03 4	4.90 4	7 4	1.19
$K = 8$	WWWW-model BR	3.75 8	8.64 8	10.5 8	15 8	1.17
$\mu_k = 1 + \frac{M-1}{K-1}(k-1)$		$M = 1$	$M = 2$	$M = 4$	$M = 6$	$M = 8$
$K = 2$	W-model BR	1.5 2	3 2	3 2	3 2	3 2
$K = 4$	WW-model BR	2.33 4	4.66 4	7 4	7 4	7 4
$K = 8$	WWWW-model BR	3.75 8	7.14 8	10.71 8	12.85 8	15 8

lower bounds of our system for which the dynamics are easier to characterize. Proving that the upper bound (lower bound) is stable (unstable) directly implies that the original system is also stable (unstable). This will be done in Corollary 7.7 and Corollary 7.12.

Before doing so, we introduce some notation: We denote by $E_c(t) = \max\{j : U_{cj} < t\}$ the number of type- c jobs that arrived during the time interval $(0, t)$ and by U_{cj} the instant of time at which the j th type- c job arrived to the system. We recall that b_{cj} denotes its service realization. We denote by $b'_{c_m s}$ the residual job size of the m th eldest type- c job in server s that is already in service at time 0. We define the random variable $\tau_s = \inf\{t > 0 : M_s(t) = 0\}$ as the first time at which server s is empty.

Sufficient stability condition

We define the Upper Bound (UB) system as follows. Upon arrival, each job is with probability p_c of type c and sends identical copies to all servers $s \in C$. Recall the set $\mathcal{R}(c)$, which denotes the set of servers where type c receives maximum capacity-to-fraction-of-arrivals ratio. In the UB system, a type- c job departs the system **only when all copies in the set of servers $\mathcal{R}(c)$ are fully served**. When this happens, the remaining copies that are still in service (necessarily not in a server in $\mathcal{R}(c)$) are immediately removed from the system. We denote by $N_c^{UB}(t)$ the number of type- c jobs present in the UB system at time t .

In the remainder of this section, we will show that the condition of Proposition 5.1 is a sufficient stability condition for the upper bound system. Intuitively, this can be explained as follows. Given a server $s \in \mathcal{L}_1$ and any type $c \in C(s)$, it holds that $\mathcal{R}(c) \subseteq \mathcal{L}_1(c)$. Hence, a server in \mathcal{L}_1 will need to fully serve all arriving copies. Therefore each server s , with $s \in \mathcal{L}_1$, behaves as an M/G/1 PS queue, which is stable if and only if its arrival rate of copies, $\lambda \sum_{c \in C_1(s)} p_c$, is strictly smaller than its departure rate, μ_s . Assume now that for all $l = 1, \dots, i-1$ the subsystems S_l are stable and we want to show that servers in \mathcal{L}_i are stable as well. First of all, note that in the fluid limit, all types c that do not exist in the S_i -subsystem, i.e., $c \notin C_i(s)$, will after a finite amount of time equal (and remain) zero, since they are stable. For the remaining types c that have copies in server $s \in \mathcal{L}_i$, i.e., $s \in c$ with $s \in \mathcal{L}_i$, it will hold that their servers with maximum capacity-to-fraction-of-arrivals ratio are $\mathcal{R}(c) \subseteq \mathcal{L}_i$. Due to the characteristics of the upper-bound system, all copies sent to these servers will need to be served. Hence, a server $s \in \mathcal{L}_i$ behaves in the fluid limit as an M/G/1 PS queue with arrival rate

$\lambda \sum_{c \in C_i(s)} p_c$ and departure rate μ_s . In particular, it is stable if and only if $\lambda \sum_{c \in C_i(s)} p_c < \mu_s$.

Remark 7.1. In [3], the authors show that for the redundancy- d system with homogeneous arrivals and server capacities, the system where all the copies need to be served is an upper bound. We note that this upper bound coincides with our upper bound (in that case $\mathcal{L}_1 = S$). There, the proof followed directly, as each server behaved as an M/G/1 PS queue. In the heterogeneous server setting studied here, the difficulty lies in the fact that similar arguments can be used only in a recursive manner. In order to see a server as a PS queue in the fluid regime, one first needs to argue that the types that have copies in higher capacity-to-fraction-of-arrivals servers are 0 at a fluid scale.

In the following, we prove that UB provides an upper bound on the original system. To do so, we show that every job departs earlier in the original system than in the UB system. In the statement, we assume that in case a job has already departed in the original system, but not in the UB system, then its attained service in all its servers in the original system is set equal to its service requirement b_{cj} .

PROPOSITION 7.2. *Assume $N_c(0) = N_c^{UB}(0)$ and $a_{cjs}(0) = a_{cjs}^{UB}(0)$, for all c, j, s . Then, $N_c(t) \leq N_c^{UB}(t)$ and $a_{cjs}(t) \geq a_{cjs}^{UB}(t)$, for all c, j, s and $t \geq 0$.*

We first describe the dynamics of the number of type- c jobs in the UB system. We recall that a type- c job departs only when all the copies in the set of servers $\mathcal{R}(c)$ are completely served. We define by $\eta_{\mathcal{R}(c)}^{min}(v, t) = \min_{\tilde{s} \in \mathcal{R}(c)} \{\eta_{\tilde{s}}(v, t)\}$ the minimum cumulative amount of capacity received by the copy in servers $\mathcal{R}(c)$. Therefore,

$$N_c^{UB}(t) = \sum_{m=1}^{N_c^{UB}(0)} \mathbf{1}(\{\exists \tilde{s} \in \mathcal{R}(c) : b'_{cm\tilde{s}} > \eta_{\tilde{s}}(0, t)\}) + \sum_{j=1}^{E_c(t)} \mathbf{1}(b_{cj} > \eta_{\mathcal{R}(c)}^{min}(U_{cj}, t)).$$

We denote the number of type- c copies in server s by $M_{s,c}^{UB}(t)$. We note that for a type- c job in server s there are two possibilities:

- if $s \in \mathcal{R}(c)$, the copy of the type- c job leaves the server as soon as it is completely served. The cumulative amount of capacity that the copy receives during (v, t) is $\eta_s(v, t)$.
- If $s \notin \mathcal{R}(c)$, the copy of the type- c job in server s leaves the system either if it is completely served or if all copies of this type- c job in the servers $\mathcal{R}(c)$ are served. We note that for any $\tilde{s} \in \mathcal{R}(c)$, $\tilde{s} \in \mathcal{L}_l$, with $l < i$.

Hence, the number of type- c jobs in server $s \in \mathcal{L}_i$ is given by the following expression. If $s \in \mathcal{R}(c)$,

$$M_{s,c}^{UB}(t) = \sum_{m=1}^{M_{s,c}^{UB}(0)} \mathbf{1}(b'_{cms} > \eta_s(0, t)) + \sum_{j=1}^{E_c(t)} \mathbf{1}(b_{cj} > \eta_s(U_{cj}, t))$$

and if $s \notin \mathcal{R}(c)$,

$$M_{s,c}^{UB}(t) = \sum_{m=1}^{M_{s,c}^{UB}(0)} \mathbf{1}(\{\exists \tilde{s} \in \mathcal{R}(c) : b'_{cm\tilde{s}} > \eta_{\tilde{s}}(0, t)\} \cap b'_{cms} > \eta_s(0, t)) + \sum_{j=1}^{E_c(t)} \mathbf{1}(b_{cj} > \eta_{\mathcal{R}(c),s}(U_{cj}, t)),$$

where $\eta_{\mathcal{R}(c),s}(v, t) = \max\{\eta_{\mathcal{R}(c)}^{min}(v, t), \eta_s(v, t)\}$. The first terms in both equations correspond to the type- c jobs that were already in the system by time $t = 0$, the second terms correspond to the type- c jobs that arrived during the time interval $(0, t)$.

In the following we obtain the number of copies per server. Before doing so, we need to introduce some additional notation. Let $\mathcal{D}^l(s) = \{c \in C(s) : \mathcal{R}(c) \subseteq \mathcal{L}_l(c)\}$ be the set of types in server s for which the set of servers where these types receive maximum capacity-to-fraction-of-arrivals ratio is $\mathcal{R}(c) \subseteq \mathcal{L}_l(c)$. If $s \in \mathcal{L}_i$, then, by definition, $\mathcal{D}^l(s) \neq \emptyset$ if $l \leq i$ and $\{\mathcal{D}^l(s)\}_{l=1}^i$ forms a partition of $C(s)$. Furthermore, $\mathcal{D}^i(s) = C_i(s)$, for all $s \in \mathcal{L}_i$. Therefore, for a server $s \in \mathcal{L}_i$, the number of copies in the server is given by the following expression:

$$M_s^{UB}(s) = \sum_{c \in C(s)} M_{s,c}^{UB}(t) = \sum_{l=1}^{i-1} \sum_{c \in \mathcal{D}^l(s)} M_{s,c}^{UB}(t) + \sum_{c \in C_i(s)} M_{s,c}^{UB}(t) \quad (3)$$

The first term of the RHS of the equation corresponds to the type- c jobs in server s that have $\mathcal{R}(c) \subseteq \mathcal{L}_l(c)$. The second term of the RHS corresponds to type- c jobs in server s that have $\mathcal{R}(c) \subseteq \mathcal{L}_i(c)$. Particularly, we note that in the UB system, $M_s^{UB}(t) \leq \sum_{c \in C(s)} N_c^{UB}(t)$, since copies might have left, while the job is still present.

In order to prove the stability condition, we investigate the fluid-scaled system. The fluid-scaling consists in studying the rescaled sequence of systems indexed by parameter r . For $r > 0$, denote by $N_{c,s}^{UB,r}(t)$ the system where the initial state satisfies $M_{s,c}^{UB}(0) = r m_{s,c}^{UB}(0)$, for all $c \in C$ and $s \in S$. We define, $\tau_s^r = \tau_s/r$,

$$\begin{aligned} M_{s,c}^{UB,r}(t) &= \frac{M_{s,c}^{UB}(rt)}{r}, \\ M_s^{UB,r}(t) &= \frac{M_s^{UB}(rt)}{r}, \\ \eta_s^r(u, t) &= \eta_s(ru, rt), \\ \eta_{\mathcal{R}(c)}^{min,r}(u, t) &= \eta_{\mathcal{R}(c)}^{min}(ru, rt), \\ \eta_{\mathcal{R}(c),s}^r(u, t) &= \eta_{\mathcal{R}(c),s}(ru, rt). \end{aligned}$$

In the following, we give the characterization of the fluid model.

Definition 7.3. Non-negative continuous functions $m_s^{UB}(\cdot)$ are a fluid model solution if they satisfy the functional equations

$$\begin{aligned} m_s^{UB}(t) &= \sum_{l=1}^{i-1} \sum_{c \in \mathcal{D}^l(s)} \left[m_{s,c}^{UB}(0) \left(1 - G \left(\bar{\eta}_{\mathcal{R}(c),s}(0, t) \right) \right) \right. \\ &\quad \left. + \lambda p_c \left(\int_{x=0}^t 1 - F \left(\bar{\eta}_{\mathcal{R}(c),s}(x, t) \right) dx \right) \right] \\ &+ \sum_{c \in \mathcal{C}_i(s)} \left[m_{s,c}^{UB}(0) (1 - G(\bar{\eta}_s(0, t))) \right. \\ &\quad \left. + \lambda p_c \int_{x=0}^t (1 - F(\bar{\eta}_s(x, t))) dx \right], \quad (4) \end{aligned}$$

for $s \in \mathcal{L}_i$ and $i = 1, \dots, i^*$, where $G(\cdot)$ is the distribution of the remaining service requirements, $F(\cdot)$ the service time distribution of arriving jobs,

$$\begin{aligned} \bar{\eta}_s(v, t) &= \int_{x=v}^t \phi_s(\bar{m}^{UB}(x)) dx, \\ \bar{\eta}_{\mathcal{R}(c)}^{min}(v, t) &= \min_{\tilde{s} \in \mathcal{R}(c)} \{\bar{\eta}_{\tilde{s}}(v, t)\}, \text{ and} \\ \bar{\eta}_{\mathcal{R}(c),s}(v, t) &= \max\{\bar{\eta}_{\mathcal{R}(c)}^{min}(v, t), \bar{\eta}_s(v, t)\}. \end{aligned}$$

The existence and convergence of the fluid limit to the fluid model can now be proved.

PROPOSITION 7.4. *The limit point of any convergent subsequence of $(\bar{M}^{UB,r}(t); t \geq 0)$ is almost surely a solution of the fluid model (4).*

We can now prove that the UB system is Harris recurrent. Note that the concept of Harris recurrence is needed here since the state space is obviously not countable, (as we need to keep track of residual service times). Hence, we need to prove that there exists a petite-set C for which $P(\tau_C < \infty) = 1$ where τ_C is the stopping time of C , see e.g., [4, 6, 22] for the corresponding definitions. To do so, we first establish the fluid stability, that is, its associate fluid limit is 0 in finite time. The latter is useful, as we can use the results of [21] that establish that under some suitable conditions, fluid stability implies Harris recurrency, see the lemma below.

LEMMA 7.5. *If the fluid limit is fluid stable, then the stochastic system is Harris recurrent.*

We now give a further characterization of the fluid limit.

PROPOSITION 7.6. *Let $i \leq i^*$ and assume $\lambda \sum_{c \in \mathcal{C}_i(s)} p_c < \mu_s$ for all $l \leq i-1$ and $s \in \mathcal{L}_l$. Then, there is a time $T \geq 0$, such that for $t \geq T$ and for $s \in \cup_{l=1}^{i-1} \mathcal{L}_l$,*

$$m_s^{UB}(t) = 0$$

and for $s \in \mathcal{L}_i$

$$\begin{aligned} m_s^{UB}(t) &= \sum_{c \in \mathcal{C}_i(s)} \left[m_{s,c}^{UB}(0) (1 - G(\bar{\eta}_s(0, t))) \right. \\ &\quad \left. + \lambda p_c \int_{x=0}^t (1 - F(\bar{\eta}_{s,i}(x, t))) dx \right], \quad (5) \end{aligned}$$

with

$$\bar{\eta}_{s,i}(v, t) := \int_{x=v}^t \phi_{s,i}(\bar{m}(x)) dx,$$

$$\text{and } \phi_{s,i}(\bar{m}(x)) := \frac{\mu_s}{\sum_{c \in \mathcal{C}_i(s)} m_{s,c}(x)}.$$

If $\lambda < CRL_l$, for all $l = 1, \dots, i$, then from Lemma 7.5 and Proposition 7.6 we conclude that for servers $s \in \mathcal{L}_i$, the associated stochastic number of copies in server s is Harris recurrent. Simply by noting that (5) coincides with the fluid limit of an $M/G/1$ PS system with arrival rate $\lambda \sum_{c \in \mathcal{C}_i(s)} p_c$ and server speed μ_s . Since $\lambda \sum_{c \in \mathcal{C}_i(s)} p_c < \mu_s$ ($\lambda < CRL_i$), (5) equals zero in finite time.

COROLLARY 7.7. *For $i \leq i^*$, the set of servers $s \in \mathcal{L}_i$ in the UB system is stable if $\lambda < CRL_l$, for all $l = 1, \dots, i$.*

Necessary stability condition

In this section we prove the necessary stability condition of Proposition 5.1. Let us first define

$$i := \min \{l = 1, \dots, i^* : \lambda > CRL_l\}.$$

We note that for any $i < i$, $\lambda < CRL_i$. So that the servers in \mathcal{L}_i , with $i < i$ are stable, see Corollary 7.7. We are left to prove that the servers in \mathcal{S}_i cannot be stable. In order to do so, we construct a lower-bound system.

In the S_i subsystem, the capacity-to-fraction-of-arrivals ratios are such that for all $s \in S_i$, $\mu_s / (\sum_{c \in \mathcal{C}_i(s)} p_c) \leq CRL_i$. We construct a lower bound (LB) system in which the resulting capacity-to-fraction-of-arrivals ratio is CRL_i for all servers $s \in S_i$. We use the superscript LB in the notation to refer to this system, which is defined as follows. First of all, we only want to focus on the S_i system, hence, we set the arrival rate $p_c^{LB} = 0$ for types $c \in C \setminus C_i$, whereas the arrival rate for types $c \in C_i$ remain unchanged, i.e., $p_c^{LB} = p_c$. The capacity of servers $s \in S_i$ in the LB-system is set to

$$\mu_s^{LB} := \mu_s \frac{\sum_{c \in \mathcal{C}_i(s)} p_c}{\sum_{c \in \mathcal{C}_i(\tilde{s})} p_c} = \gamma \cdot \left(\sum_{c \in \mathcal{C}_i(s)} p_c \right),$$

where $\tilde{s} \in \mathcal{L}_i$ and $\gamma = CAR_i$. Additionally, in the LB-system, we assume that each copy of a type- c job receives the same amount of capacity, which is equal to the highest value of $\mu_s^{LB} / M_s^{LB}(t)$, $s \in c$. We therefore define the service rate for a job of type c by $\phi_c^{LB}(\bar{N}^{LB}(t))$, where $c \in C_i$ (instead of $\phi_s(\cdot)$ for a copy in server s in the original system). This is given by

$$\phi_c^{LB}(\bar{N}^{LB}(t)) := \max_{s \in c} \left\{ \frac{\mu_s^{LB}}{M_s^{LB}(t)} \right\},$$

The cumulative amount of capacity that a type- c job receives is

$$\eta_c^{LB}(v, t) := \int_{x=v}^t \phi_c^{LB}(\bar{N}^{LB}(x)) dx, \text{ for } c \in C_i.$$

Remark 7.8. The lower bound in [3] was simpler as it only needed to modify the service rates, similar to $\phi_c^{LB}(\cdot)$. In the heterogeneous case, we also need to modify the capacities of the servers. This change increases the rate of every server in such a way that the capacity-to-fraction-of-arrivals ratio is equal to γ , the maximum ratio in the subsystem allowing to find the fluid limits.

The number of type- c jobs in the system is given by

$$N_c^{LB}(t) = 0, \quad c \in C \setminus C_i$$

$$N_c^{LB}(t) = \left[\sum_{m=1}^{N_c^{LB}(0)} 1 \left(b'_{cms} > \eta_c^{LB}(0, t) \right) + \sum_{j=1}^{E_c(t)} 1 \left(b_{cj} > \eta_c^{LB}(U_{cj}, t) \right) \right], \quad c \in C_i.$$

In the following, we prove that LB provides a lower bound for the original system.

LEMMA 7.9. *Assume $N_c(0) = N_c^{LB}(0)$, for all c . Then, $N_c(t) \geq_{st} N_c^{LB}(t)$, for all $c \in C$ and $t \geq 0$.*

In order to show that the LB system is unstable, we investigate the fluid-scaled system. For $r > 0$, denote by $N_c^{LB,r}(t)$ the system where the initial state satisfies $N_c^{LB}(0) = rn_c^{LB}(0)$, for all $c \in C$. We write for the fluid-scaled number of jobs per type

$$N_c^{LB,r}(t) = \frac{N_c^{LB}(rt)}{r},$$

$$\eta_c^{LB,r}(u, t) = \eta_c^{LB}(ru, rt).$$

In the following we give the characterization of the fluid model.

Definition 7.10. Non-negative continuous functions $n_c^{LB}(\cdot)$ are a fluid model solution if they satisfy the functional equations

$$n_c^{LB}(t) = 0, \quad c \in C \setminus C_i$$

$$n_c^{LB}(t) = n_c^{LB}(0) \left(1 - G \left(\bar{\eta}_c^{LB}(0, t) \right) \right) + \lambda p_c \left(\int_{x=0}^t 1 - F \left(\bar{\eta}_c^{LB}(x, t) \right) dx \right), \quad c \in C_i$$

where $G(\cdot)$ is the distribution of the remaining service requirements of initial jobs, $F(\cdot)$ the service time distribution of arriving jobs and

$$\bar{\eta}_c^{LB}(v, t) = \int_{x=v}^t \phi_c^{LB}(\bar{\eta}_c^{LB}(x)) dx, \quad \text{with } c \in C_i.$$

The existence and convergence of the fluid limit to the fluid model can be proved as before. The statement of Proposition 7.4, indeed directly translates to the process $\bar{N}^{LB,r}(t)$, since $\eta_c^{LB}(v, t)$ is both decreasing and continuous in v . Therefore, it is left out.

Next, we characterize the fluid limit of the \bar{N}^{LB} system.

LEMMA 7.11. *Let us assume that the initial condition is such that $n_c^{LB}(0) = 0$ for all $c \in C \setminus C_i$ and for $c \in C_i$, $n_c^{LB}(0)$ are such that $m_s^{LB}(0)/\mu_s^{LB} = \alpha(0)$ for all $s \in S_i$.*

Let

$$\alpha(t) = \alpha(0)(1 - G(\bar{\eta}_\alpha^{LB}(0, t))) + \frac{\lambda}{\gamma} \int_{x=0}^t (1 - F(\bar{\eta}_\alpha^{LB}(x, t))) dx, \quad (6)$$

where $\bar{\eta}_\alpha^{LB}(v, t) = \int_{x=v}^t \phi_\alpha^{LB}(\alpha(x)) dx$, with $\phi_\alpha^{LB}(\alpha(t)) = \frac{1}{\alpha(t)}$.

Then, $n_c^{LB}(t) = 0$ for all $t \geq 0$ and $c \in C \setminus C_i$, and $m_s^{LB}(t)/\mu_s^{LB} = \alpha(t)$ for all $t \geq 0$ and $s \in S_i$.

We note that Equation (6) corresponds to the fluid limit of an $M/G/1$ system with PS, arrival rate λ/γ and server speed 1. Assuming $\lambda > \gamma$ (or equivalently $\lambda > CAR_i$), it follows that the fluid limit $\alpha(t)$, and hence $m_s^{LB}(t)$, $s \in S_i$, diverges. Now, by using similar

arguments as in Dai [7], the fact that the limit diverges implies that the corresponding stochastic process can not be tight, and hence cannot be stable.

COROLLARY 7.12. *In the LB-system, the set of servers $s \in \mathcal{L}_i$ is unstable if $\lambda > CRL_i$.*

8 NUMERICAL ANALYSIS

We have implemented a simulator in order to assess the impact of redundancy, and to compare numerically the performance with Bernoulli routing. We ran these simulations for a large number of busy periods (10^6), so that the variance and confidence intervals of the mean number of jobs in the system are sufficiently small.

Exponential service time distributions: In Figure 4 we consider the W -model with exponential service time distributions. We set $p_{\{1\}} = 0.35$ and $p_{\{2\}} + p_{\{1,2\}} = 0.65$, and vary the value of $p_{\{1,2\}}$. We consider two different sets of server capacities, namely $\bar{\mu} = \{1, 2\}$ and $\bar{\mu} = \{2, 1\}$. The only redundant job type is $\{1, 2\}$, thus as its value increases, we can observe how increasing the fraction of redundant jobs affects the performance. We also note that when $p_{\{1,2\}}$ increases, the load in server 1 increases as well, whereas the load in server 2 stays constant. In a) and b) we depict the mean number of jobs with redundancy and Bernoulli routing for two different arrival rates, and in c) we plot λ^R and $\lambda^{*,B}$ using the analysis of Section 5.2.2.

Qualitatively, plots a) and b) show a different behavior as $p_{1,2}$ increases. In the $\bar{\mu} = \{2, 1\}$ case, redundancy performs better than Bernoulli for most values of $p_{\{1,2\}}$, but the difference is not large, particularly for relatively large values of $p_{\{1,2\}}$. On the other hand, in the $\bar{\mu} = \{1, 2\}$ case, redundancy does much better than Bernoulli routing for large values of $p_{\{1,2\}}$. This is due to the fact that for large values of $p_{\{1,2\}}$, redundancy does better in exploiting the larger capacity of server 1. In c) we observe that redundancy consistently has a larger stability region than Bernoulli.

In Figure 5 we simulate the performance of the W model for different values of μ_2 , while keeping fixed \bar{p} and $\mu_1 = 1$. In a) we plot the mean number of jobs and we see that for both configurations of \bar{p} , the performance of both redundancy and Bernoulli improve as μ_2 increases, and that redundancy consistently performs better or similar to Bernoulli routing. The gap between redundancy and Bernoulli is significant in the case of $\bar{p} = \{1/2, 1/4, 1/4\}$. The reason can be deduced from b), where we plot λ^R and λ^B with respect to μ_2 . We observe in b) that the stability region of redundancy is, for every value of μ_2 , larger than that of Bernoulli, and that the gap is larger for the case $\bar{p} = \{1/2, 1/4, 1/4\}$.

General service time distributions: In Figure 6 a) we investigate the performance for several non-exponential distributions. In particular, we consider the following distributions for the service times: deterministic, hyperexponential, and Bounded Pareto. With the hyperexponential distribution, job sizes are exponentially distributed with parameter μ_1 (μ_2) with probability p ($1 - p$). For Pareto the density function is $\frac{1-(k/x)^\alpha}{(1-(k/q)^\alpha)}$, for $k \leq x \leq q$. We choose the parameters so that the mean service time equals 1. In a), we plot the mean number of jobs as a function of λ for the N , W , WW , and redundancy-2 ($K = 5$), and redundancy-4 ($K = 5$) models. The respective parameters \bar{p} are chosen such that the system is stable for the simulated arrival rates. We observe that for the five systems,

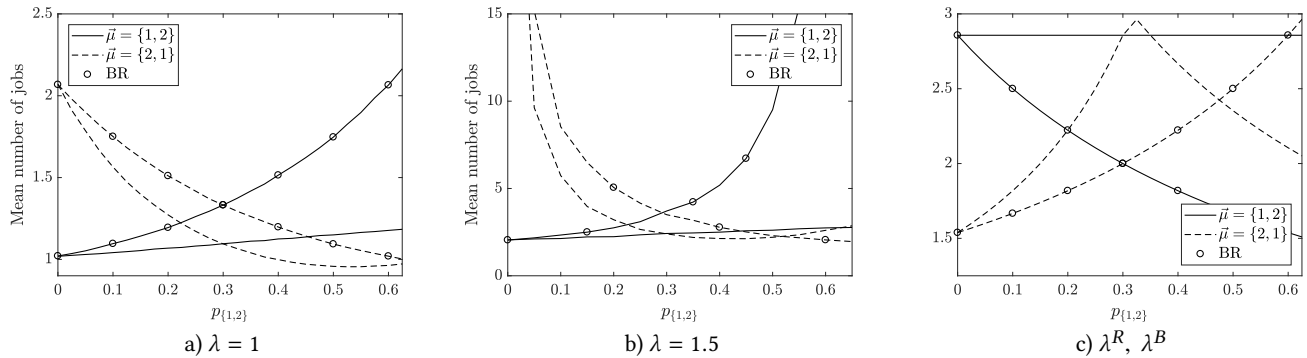


Figure 4: W -model with $p_{\{1\}} = 0.35$, $p_{\{2\}} = 1 - p_{\{1\}} - p_{\{1,2\}}$. Figures a) and b) depict the mean number of jobs under redundancy and Bernoulli routing (BR) for $\lambda = 1$ and $\lambda = 1.5$. Figure c) depicts the stability regions λ^R and λ^B .

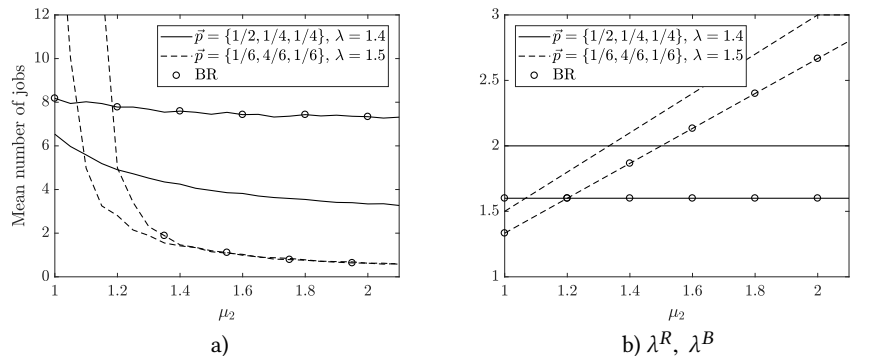


Figure 5: W -model with fixed parameters \vec{p} and $\mu_1 = 1$: a) depicts the mean number of jobs under redundancy and Bernoulli routing, and b) depicts the stability regions λ^R and λ^B .

Table 4: The Dolly(1,12) empirical distribution for the slow-down [1]. The capacity is set to $1/S$.

S	1	2	3	4	5	6	7	8	9	10	11	12
Prob	0.23	0.14	0.09	0.03	0.08	0.10	0.04	0.14	0.12	0.021	0.007	0.002

performance seems to be nearly insensitive to the service time distribution, beyond its mean value.

Markov-modulated capacities: In Figure 6 b) we consider a variation of our model where servers' capacities fluctuate over time. More precisely, we assume that each server has an exponential clock, with mean ϵ . Every time the clock rings, the server samples a new value for S from Dolly(1,12), see Table 4 and sets its capacity equal to $1/S$. The Dolly(1,12) distribution is a 12-valued discrete distribution that was empirically obtained by analyzing traces in Facebook and Microsoft clusters, see [1, 10].

In Figure 6 b) we plot the mean number of jobs for a $K = 5$ server system with redundancy-2 and redundancy-4, and for the W -model under redundancy, and we compare it with Bernoulli routing. Arrival rates are equal for all classes. It can be seen that

with Bernoulli routing, both redundancy-2 and redundancy-4 become equivalent systems, and hence their respective curves overlap. The general observation is that in this setting with identical servers, Bernoulli routing performs better than redundancy. Further research is needed to understand whether with heterogeneous Markov-modulated servers, redundancy can be beneficial.

9 CONCLUSION

With exponentially distributed jobs, and i.i.d. copies, it has been shown that redundancy does not reduce the stability region of a system, and that it improves the performance. This happens in spite of the fact that redundancy necessarily provokes a waste of computation resources in servers that work on copies that are canceled before being fully served. The modeling assumptions play thus a crucial role, and as argued in several papers, e.g. [10], the i.i.d. assumption might lead to insights that are qualitatively wrong.

In the present work, we consider the more realistic situation in which copies are identical, and the service times are generally distributed. In our main result we have shown that redundancy can help improve the performance in case the server's capacities are sufficiently heterogeneous. To the best of our knowledge, this is the first positive result on redundancy with identical copies, and it

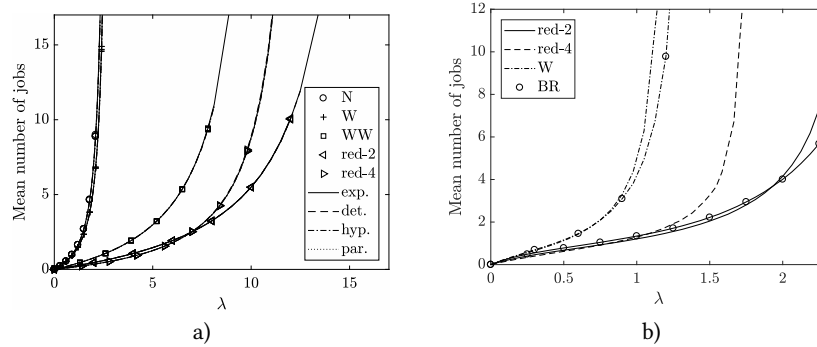


Figure 6: Mean number of jobs in the system with respect to λ : a) Non-exponential service times and models N , W , WW , and redundancy-2 ($K = 5$), and redundancy-4 ($K = 5$) models. We chose $\vec{\mu} = (1, 2)$ for the N and W model, $\vec{\mu} = (1, 2, 4, 6)$ for the WW model, and $\vec{\mu} = (1, 2, 4, 6, 8)$ for redundancy- d . b) Markov modulated server capacities in the W -model.

illustrates that the negative result proven in [3] critically depends on the fact that the capacities were homogeneous.

We thus believe that our work opens the avenue for further research to understand when redundancy is beneficial in other settings. For instance, it would be interesting to investigate what happens in case servers implement other scheduling policies. It is also important to consider other cross-correlation structures for the copies, in particular the S&X model recently proposed in the literature. Another interesting situation is when the capacities of the servers fluctuate over time. Other possible extension is to consider the cancel-on-start variant of redundancy, in which as soon as one copy enters service, all the others are removed. For conciseness purposes, in this paper we have restricted ourselves to what we considered one of the most basic, yet interesting and relevant setting.

REFERENCES

- [1] Ganesh Ananthanarayanan, Ali Ghodsi, Scott Shenker, and Ion Stoica. 2013. Effective Straggler Mitigation: Attack of the Clones. In *NSDI*, Vol. 13. 185–198.
- [2] Ganesh Ananthanarayanan, Srikanth Kandula, Albert G Greenberg, Ion Stoica, Yi Lu, Bikas Saha, and Edward Harris. 2010. Reining in the Outliers in Map-Reduce Clusters using Mantri. In *OSDI’10 Proceedings of the 9th USENIX conference on Operating systems design and implementation*. 265–278.
- [3] Elene Anton, Urtzi Ayesta, Matthieu Jonckheere, and Ina Maria Verloop. March, 2019. On the stability of redundancy models. *arXiv:1903.04414* (March, 2019).
- [4] Soeren Asmussen. 2002. *Applied Probability and Queues*. Springer.
- [5] Thomas Bonald and Céline Comte. 2017. Balanced fair resource sharing in computer clusters. *Performance Evaluation* 116 (2017), 70–83.
- [6] Maury Bramson. 2008. *Stability of Queueing Networks*. Springer.
- [7] Jim G. Dai. 1996. A fluid limit model criterion for instability of multiclass queueing networks. *The Annals of Applied Probability* 6 (1996), 751–757. <https://doi.org/10.1214/aoap/1034968225>
- [8] Jeffrey Dean and Luiz André Barroso. 2013. The tail at scale. *Commun. ACM* 56, 2 (2013), 74–80.
- [9] Regina Egorova. 2009. *Sojourn time tails in processor-sharing systems*, Technische Universiteit Eindhoven. Ph.D. Dissertation. <https://doi.org/10.6100/IR639849>
- [10] Kristen Gardner, Mor Harchol-Balter, Alan Scheller-Wolf, and Benny van Houdt. 2017. A Better Model for Job Redundancy: Decoupling Server Slowdown and Job Size. *IEEE/ACM Transactions on Networking* 25, 6 (2017), 3353–3367.
- [11] Kristen Gardner, Esa Hyttiä, and Rhonda Righter. 2019. A Little Redundancy Goes a Long Way: Convexity in Redundancy Systems. *Performance Evaluation (2019)* (2019). <https://doi.org/10.1016/j.peva.2019.02.001>
- [12] Kristen Gardner, Samuel Zbarsky, Sherwin Doroudi, Mor Harchol-Balter, Esa Hyttiä, and Alan Scheller-Wolf. 2016. Queueing with redundant requests: exact analysis. *Queueing Systems* 83, 3-4 (2016), 227–259.
- [13] H. Christian Gromoll, Philippe Robert, and Bert Zwart. 2008. Fluid Limits for Processor Sharing Queues with Impatience. *Math. Oper. Res.* 33 (05 2008), 375–402. <https://doi.org/10.1287/moor.1070.0298>
- [14] Mor Harchol-Balter. 2013. *Performance Modeling and Design of Computer Systems: Queueing Theory in Action*. Cambridge University Press.
- [15] Tim Hellemans, Tejas Bodas, and Benny van Houdt. 2019. Performance Analysis of Workload Dependent Load Balancing Policies. *POMACS* 3, 2 (2019), 35:1–35:35. <https://doi.org/10.1145/3341617.3326150>
- [16] Tim Hellemans and Benny van Houdt. 2018. Analysis of redundancy(d) with identical Replicas. *Performance Evaluation Review* 46, 3 (2018), 1–6.
- [17] Gauri Joshi, Emina Soljanin, and Gregory Wornell. 2015. Queues with redundancy: Latency-cost analysis. *ACM SIGMETRICS Performance Evaluation Review* 43, 2 (2015), 54–56.
- [18] Ger Koole and Rhonda Righter. 2007. Resource allocation in grid computing. *Journal of Scheduling* (2007).
- [19] Kangwook Lee, Ramtin Pedarsani, and Kannan Ramchandran. 2017. On scheduling redundant requests with cancellation overheads. *IEEE/ACM Transactions on Networking (TON)* 25, 2 (2017), 1279–1290.
- [20] Kangwook Lee, Nihar B. Shah, Longbo Huang, and Kannan Ramchandran. 2017. The mds queue: Analysing the latency performance of erasure codes. *IEEE Transactions on Information Theory* 63, 5 (2017), 2822–2842.
- [21] Nam H. Lee. 2008. *A sufficient condition for stochastic stability of an Internet congestion control model in terms of fluid model stability*, UC San Diego. Ph.D. Dissertation.
- [22] Sean Meyn and Richard Tweedie. 1993. Generalized resolvents and Harris recurrence of Markov processes. *Contemp. Math.* 149 (1993), 227–250. <https://doi.org/10.1090/conm/149/01273>
- [23] Youri Raaijmakers, Sem Borst, and Onno Boxma. 2019. Redundancy scheduling with scaled Bernoulli service requirements. *Queueing Systems Volume* 93 (2019), Issue 1-2.
- [24] Youri Raaijmakers, Sem Borst, and Onno Boxma. December, 2019. Stability of Redundancy Systems with Processor Sharing. *arXiv:1912.00681* (December, 2019).
- [25] Nihar B. Shah, Kangwook Lee, and Kannan Ramchandran. 2016. When do redundant requests reduce latency? *IEEE Transactions on Communications* 64, 2 (2016), 715–722.
- [26] Ashish Vulimiri, Philip Brighten Godfrey, Radhika Mittal, Justine Sherry, Sylvia Ratnasamy, and Scott Shenker. 2013. Low latency via redundancy. In *Proceedings of the ACM conference on Emerging networking experiments and technologies*. ACM, 283–294.

APPENDIX A: PROOFS OF SECTION 5

Proof of Corollary 5.3

Let us consider $s \in \mathcal{R}$. Let i be such that $s \in \mathcal{L}_i$, which is unique since $\{\mathcal{L}_i\}_{i=1}^I$ is a partition of \mathcal{R} . We will show that for this s and i , it holds that $CAR_i = \frac{\mu_s}{\sum_{c:s \in \mathcal{R}(c)} P_c}$. Hence, together with Corollary 5.2 this concludes the result.

First, note that $CAR_i = \frac{\mu_s}{\sum_{c \in C_i(s)} P_c}$. Hence, we need to prove that $\sum_{c:s \in \mathcal{R}(c)} P_c = \sum_{c \in C_i(s)} P_c$, or equivalently, $\{c : s \in \mathcal{R}(c)\} = C_i(s)$.

For any $c \in C(s)$, $\mathcal{R}(c) = \mathcal{L}_l(c)$ with $l \leq i$. We note that $C_i(s) = C(s) \setminus \{c \in C(s) : \mathcal{R}(c) = \mathcal{L}_l(c) \text{ with } l < i\}$. Therefore, for $s \in \mathcal{L}_i$, $C_i(s) = \{c \in C : s \in c, c \in C_i, s \in \mathcal{L}_i(c)\} = \{c \in C : s \in \mathcal{R}(c)\}$. The last equality holds by definition of $\mathcal{R}(c)$. \square

Proof of Corollary 5.4.

The stability condition of such a system is given by Corollary 5.2. We note that each server $s \in S$ receives $C(s) = \binom{K-1}{d-1}$ different job types, that is, by fixing a copy in server s , all possible combinations of $d-1$ servers out of $K-1$. Thus, $\mathcal{L}_1 = \arg \max_{s \in S_1} \left\{ \binom{K}{d-1} \mu_s \right\} = K$, $S_2 = S - \{K\}$ and condition $\lambda \binom{K-1}{d} < \mu_K$.

We note each server $s \in S_i$ receives $\binom{|S_i|-1}{d-1}$ different job types, for $i = 1, \dots, i^*$ and thus, the maximum capacity-to-fraction-of-arrivals ratio in the subsystem with servers S_i , only depends on the capacities of servers in S_i , that is $\mathcal{L}_i = \arg \max_{s \in S_i} \{\mu_s\}$. Additionally since, $\mu_1 < \dots < \mu_K$, one obtains that $\mathcal{L}_i = K - i + 1$, for $i = 1, \dots, K-d+1$. The associated conditions are $\lambda \binom{K-i+1}{d} < \mu_{K-i+1}$ for $i = 1, \dots, K-d+1$. This set of conditions is equivalent to that in Corollary 5.4. \square

APPENDIX B: PROOFS OF SECTION 7

Proof of Proposition 7.2

We assume that both systems are coupled as follows: at time $t = 0$, both systems start at the same initial state $N_c(0) = N_c^{UB}(0)$ and $a_{cjs}(0) = a_{cjs}^{UB}(0)$ for all c, j, s . Arrivals and service times are also coupled. For simplicity in notation, we assume that when in the original system a type- c copy reaches its service requirement b , the attained service of its $d-1$ additional copies is fixed to b and the job remains in the system until the copy of that same job in the UB system is fully served at all servers in $\mathcal{R}(c)$.

We prove this result by induction on t . It holds at time $t = 0$. We assume that for $u \leq t$ it holds that $N_c(t) \leq N_c^{UB}(t)$ and $a_{cjs}(t) \geq a_{cjs}^{UB}(t)$ for all c, j, s . We show that this inequality holds for t^+ .

We first assume that at time t , it holds that $N_c(t) = N_c^{UB}(t)$ for some $c \in C$. The inequality is violated only if there is a job for which the copy in the UB system is fully served at all servers $\mathcal{R}(c)$, but none of the copies in the original system is completed. That means, there exist a j such that $a_{cjs}(t) < a_{cjs}^{UB}(t) = b_j$ for all $s \in c$. However, this can not happen, since by hypothesis $a_{cjs}(t) \geq a_{cjs}^{UB}(t)$ for all $s \in c$.

We now assume that at time t , $a_{cjs}(t) = a_{cjs}^{UB}(t)$ for some c, j, s . There are now two cases. If this copy (and job) has already left in the original system, then $a_{cjs}(t) = a_{cjs}(t^+) = b_{cj}$ and hence $a_{cjs}(t^+) \geq a_{cjs}^{UB}(t^+)$. If instead the copy has not left in the original system, then by hypothesis it holds that $N_c(t) \leq N_c^{UB}(t)$ and thus, $M_s(t) \leq M_s^{UB}(t)$ and $\frac{\mu_s}{M_s(t)} \geq \frac{\mu_s}{M_s^{UB}(t)}$. That means that the copy in the original system has a higher service rate at time t than the same copy in the UB system. Hence, $a_{cjs}(t^+) \leq a_{cjs}^{UB}(t^+)$. \square

Proof of Proposition 7.4

The proof is identical to the the proof of Theorem 5.2.1 in [9] (which is itself based on Lemma 5 in [13]). We only need to ensure that $\bar{\eta}_{\mathcal{R}(c)}^{min}(v, t)$ and $\bar{\eta}_{\mathcal{R}(c),s}(v, t)$ are decreasing in v and continuous on $v \in [\psi_s(t) + \epsilon, t]$, where $\psi_s(t) = \sup\{v \in [0, t] : m_s(u) = 0\}$.

Let us verify that $\bar{\eta}_{\mathcal{R}(c)}^{min}(v, t)$ and $\bar{\eta}_{\mathcal{R}(c),s}(v, t)$ are decreasing and continuous on v . We note that the function $\eta_s(\cdot, t)$ that gives the cumulative service that a copy in server s received during time interval (\cdot, t) , is a Lipschitz continuous function, increasing for $t < \tau_s$ and non decreasing for $t > \tau_s$.

If $\bar{\eta}_{\mathcal{R}(c)}^{min}(v, t) = \bar{\eta}_{s_1}(v, t)$ and $\bar{\eta}_{\mathcal{R}(c),s}(v, t) = \bar{\eta}_{s_2}(v, t)$ for all $v \in [0, t)$ and some $s_1, s_2 \in S$, then both $\bar{\eta}_{\mathcal{R}(c)}^{min}(v, t)$ and $\bar{\eta}_{\mathcal{R}(c),s}(v, t)$ are decreasing and continuous on v , since by definition $\bar{\eta}_s(v, t)$ is decreasing and continuous on v for all $s \in S$.

Let us assume that for $v_0 \in [0, t)$ is such that $\bar{\eta}_{\mathcal{R}(c)}^{min}(v, t) = \bar{\eta}_{s^1}(v, t)$ for $v \leq v_0$ and $\bar{\eta}_{\mathcal{R}(c)}^{min}(v_0^+, t) = \bar{\eta}_{s^2}(v_0, t)$, for some $s^1, s^2 \in \mathcal{R}(c)$. We first verify that $\bar{\eta}_{\mathcal{R}(c)}^{min}(v, t)$ is continuous on $v = v_0$. Since, $\bar{\eta}_{s^1}(v, t)$ and $\bar{\eta}_{s^2}(v, t)$ are continuous on $v = v_0$, then

$$\lim_{x^+ \rightarrow v_0} \bar{\eta}_{\mathcal{R}(c)}^{min}(x, t) = \bar{\eta}_{s^1}(v_0, t) = \bar{\eta}_{s^2}(v_0, t) = \lim_{x^+ \rightarrow v_0} \bar{\eta}_{\mathcal{R}(c)}^{min}(x, t).$$

Therefore, we conclude that $\bar{\eta}_{\mathcal{R}(c)}^{min}(x, t)$ is continuous on $v \in [0, t)$. Analogously, one can verify that $\bar{\eta}_{\mathcal{R}(c),s}^{min}(x, t)$ is continuous on $v \in [0, t)$.

We now verify that $\bar{\eta}_{\mathcal{R}(c)}^{min}(x, t)$ is decreasing on $v \in [0, t)$. Let us consider $0 < t_1 < v_0 < t_2 < t$. Then for $\bar{\eta}_{\mathcal{R}(c)}^{min}(v, t)$,

$$\bar{\eta}_{\mathcal{R}(c)}^{min}(t_1, t) = \bar{\eta}_{s^1}(t_1, t) \leq \bar{\eta}_{s^1}(t_2, t) \leq \bar{\eta}_{s^2}(t_2, t) = \bar{\eta}_{\mathcal{R}(c)}^{min}(t_2, t),$$

where the first inequality holds since $\bar{\eta}_{s^1}(v, t)$ is decreasing on v . We conclude that $\bar{\eta}_{\mathcal{R}(c)}^{min}(v, t)$ is decreasing v .

Let us verify that $\bar{\eta}_{\mathcal{R}(c),s}(v, t)$ is decreasing on v . W.l.o.g. we assume that there exists $v_0 \in [0, t)$, such that $\bar{\eta}_{\mathcal{R}(c),s}(v, t) = \bar{\eta}_{\mathcal{R}(c)}^{min}(v, t)$ for $v < v_0$ and $\bar{\eta}_{\mathcal{R}(c),s}(v, t) = \bar{\eta}_s(v, t)$ for $t > v > v_0$. Then,

$$\bar{\eta}_{\mathcal{R}(c),s}(t_1, t) = \bar{\eta}_{\mathcal{R}(c)}^{min}(t_1, t) \leq \bar{\eta}_s(t_1, t) \leq \bar{\eta}_s(t_2, t) = \bar{\eta}_{\mathcal{R}(c),s}(t_2, t)$$

where the first inequality holds since $\bar{\eta}_{s^1}(v, t)$ is decreasing on v . We conclude that $\bar{\eta}_{\mathcal{R}(c),s}(x, t)$ is decreasing v . \square

Proof of Lemma 7.5

In [21], the authors consider bandwidth sharing networks (with processor sharing policies), and show that under mild conditions, the stability of the fluid model (describing the Markov process of the number of per-class customers with their residual job sizes) is sufficient for stability (positive Harris recurrence).

Our system, though slightly different from theirs satisfies the same assumptions, and as a consequence their results are directly applicable to our model.

More precisely, given the assumptions on the service time distribution, our model satisfies the assumptions given in [21, Section 2.2] for inter-arrival times and job-sizes. (In particular exponential inter-arrival times satisfy the conditions given in [21, Assumption 2.2.2].) \square

Proof of Proposition 7.6

For simplicity in notation, we remove the superscript UB throughout the proof.

First assume $s \in \mathcal{L}_1$. Since $\mathcal{D}^0 = \emptyset$, from Equation (4), we directly obtain

$$m_s(t) = \sum_{c \in \mathcal{C}_1(s)} [m_{s,c}(0)(1 - G(\bar{\eta}_s(0, t))) + \lambda p_c \int_{x=0}^t (1 - F(\bar{\eta}_s(x, t))) dx], \quad \forall t > 0.$$

This expression coincides with the fluid limit of a $M/G/1$ PS queue with arrival rate $\lambda \sum_{c \in \mathcal{C}_1(s)} p_c$ and server speed μ_s . Since $\lambda \sum_{c \in \mathcal{C}_1(s)} p_c < \mu_s$, we know that there exists a $\bar{\tau}_s$ such that $m_s(t) = 0$, for all $t \geq \bar{\tau}_s$.

The remainder of the proof is by induction. Consider now a server $s \in \mathcal{L}_l$ and assume there exists a time \tilde{T} such that $m_s(t) = 0$, for all $t \geq \tilde{T}$ and $s \in \cup_{j=1}^{l-1} \mathcal{L}_j$. Thus, for $t \geq \tilde{T}$, also $m_{s,c}(t) = 0$ for all $s \in \cup_{j=1}^{l-1} \mathcal{L}_j$, $c \in \mathcal{D}^j(s)$, $j = 1, \dots, l-1$. We consider server $s \in \mathcal{L}_l$. From (4) its drift is then given by:

$$\begin{aligned} m_s(t) &= \sum_{j=1}^{l-1} \sum_{c \in \mathcal{D}^j(s)} m_{s,c}(t) + \sum_{c \in \mathcal{C}_l(s)} m_{s,c}(t) \\ &= \sum_{c \in \mathcal{C}_l(s)} [m_{s,c}(0)(1 - G(\bar{\eta}_s(0, t))) \\ &\quad + \lambda p_c \int_{x=0}^t (1 - F(\bar{\eta}_s(x, t))) dx], \end{aligned}$$

for all $t \geq \tilde{T}$. Now note that $\phi_s(\vec{m}(t)) = \frac{\mu_s}{m_s(t)} = \frac{\mu_s}{\sum_{c \in \mathcal{C}_l(s)} m_{s,c}(t)} = \phi_{s,l}(\vec{m}(t))$, where the second equality follows from the fact that $m_{s,c}(t) = 0$ for all $s \in \cup_{j=1}^{l-1} \mathcal{L}_j$, $c \in \mathcal{D}^j(s)$, $j = 1, \dots, l-1$.

To finish the proof, (5) coincides with the fluid limit of an $M/G/1$ system with PS, arrival rate $\lambda \sum_{c \in \mathcal{C}_l(s)} p_c$ and server speed μ_s . Hence, if $l < i$, the standard PS queue is stable, and we are sure that it equals and remains zero in finite time. \square

Proof of Lemma 7.9

We note that for all $c \in C \setminus \mathcal{C}_i$ the result is direct since $p_c^{LB} = 0$ for all $c \in C \setminus \mathcal{C}_i$. Then, let us consider $c \in \mathcal{C}_i$. For any \vec{N} and \vec{N}^{LB} such that $\vec{N} \geq \vec{N}^{LB}$, the following inequalities hold:

$$\begin{aligned} \phi_s(\vec{N}) &= \frac{\mu_s}{M_s} \\ &= \frac{\mu_s (\sum_{c \in \mathcal{C}_i(s)} p_c)}{(\sum_{c \in \mathcal{C}_i(s)} p_c) M_s} = \frac{(\sum_{c \in \mathcal{C}_i(s)} p_c) \mu_s / (\sum_{c \in \mathcal{C}_i(s)} p_c)}{\sum_{c \in C(s) \setminus \mathcal{C}_i(s)} N_c + \sum_{c \in \mathcal{C}_i(s)} N_c} \\ &\leq \frac{(\sum_{c \in \mathcal{C}_i(s)} p_c) \mu_s / (\sum_{c \in \mathcal{C}_i(s)} p_c)}{\sum_{c \in \mathcal{C}_i(s)} N_c} \\ &\leq \frac{\gamma \times (\sum_{c \in \mathcal{C}_i(s)} p_c)}{\sum_{c \in \mathcal{C}_i(s)} N_c^{LB}} \leq \max_{s \in c} \left\{ \frac{\mu_s^{LB}}{M_s^{LB}} \right\} = \phi_c^{LB}(\vec{N}^{LB}). \end{aligned}$$

The second last inequality holds since $\gamma \geq \mu_s / (\sum_{c \in \mathcal{C}_i(s)} p_c)$ for all $s \in S_i$ and $N_c^{LB} \leq N_c$ for all $c \in \mathcal{C}_i$. We note that $\sum_{c \in \mathcal{C}_i(s)} N_c^{LB} = M_s^{LB}(t)$. It follows from straight forward sample-path arguments that $N_c^{LB}(t) \leq N_c(t)$ for all $t \geq 0$ and $c \in \mathcal{C}_i$. \square

Proof of Lemma 7.11

From Definition 7.10, we obtain that for each server $s \in S_i$,

$$\begin{aligned} \frac{m_s^{LB}(t)}{\mu_s^{LB}} &= \frac{1}{\mu_s^{LB}} \sum_{c \in \mathcal{C}_i(s)} n_c^{LB}(t) \\ &= \sum_{c \in \mathcal{C}_i(s)} \left[\frac{n_c^{LB}(0)}{\mu_s^{LB}} \left(1 - G(\bar{\eta}_c^{LB}(0, t)) \right) \right. \\ &\quad \left. + \frac{\lambda p_c}{\mu_s^{LB}} \left(\int_{x=0}^t 1 - F(\bar{\eta}_c^{LB}(x, t)) dx \right) \right]. \end{aligned}$$

The proof follows from observing that the given initial conditions imply that $\bar{\eta}_c^{LB}(v, t) = \bar{\eta}_c^{LB}(v, t)$ for all $0 \leq v \leq t$ and $t \geq 0$, which in turn imply the same fluid dynamics. \square