



HAL
open science

A Survey of Stability Results for Redundancy Systems

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

► **To cite this version:**

Elene Anton, Urtzi Ayesta, Matthieu Jonckheere, Ina Maria Maaïke Verloop. A Survey of Stability Results for Redundancy Systems. Alexey Piunovskiy; Yi Zhang. Modern Trends in Controlled Stochastic Processes: Theory and Applications, Volume III, 41, Springer, pp.266-283, 2021, Emergence, Complexity and Computation (ECC), 978-3-030-76927-7. 10.1007/978-3-030-76928-4_13 . hal-03168397

HAL Id: hal-03168397

<https://hal.science/hal-03168397>

Submitted on 13 Mar 2021

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.

A Survey of Stability Results for Redundancy Systems

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

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

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

³Université de Toulouse, INP, 31071 Toulouse, France

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

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

Abstract

Redundancy mechanisms consist in sending several copies of a same job to a subset of servers. It constitutes one of the most promising ways to exploit diversity in multi-servers applications. However, its pros and cons are still not sufficiently understood in the context of realistic models with generic statistical properties of service-times distributions and correlation structures of copies. We aim at giving a survey of recent results concerning the stability - arguably the first benchmark of performance - of systems with cancel-on-completion redundancy. We also point out open questions and conjectures.

Keywords: redundancy, load balancing, stability

1 Introduction

While there are several variants of redundancy-based systems, the general notion of redundancy is to dispatch multiple copies of each job to a subset of servers and to consider the result of whichever copy completes service first. By allowing for redundant copies, the aim is to minimize the system latency by exploiting the variability in the queue lengths of the different queues. The potential of redundancy mechanisms lies in finding the right trade-off between exploiting variability and the waste of resources induced by having redundant copies.

Several empirical ([2, 3, 11, 12, 38, 41]) and numerical studies ([15, 16, 26, 29, 30]) suggest that redundancy might potentially improve the performance of real-world computer system applications. In particular, Vulimiri et al. [41] consider a 10 DNS servers system and compare the system where each arriving query dispatches 10 copies to all the 10 DNS servers, to an alternative system where queries are assigned to a single server chosen uniformly at random. The authors observe that the fraction of queries with a service time exceeding 500 ms is reduced by a factor 6.5, and the fraction exceeding 1.5 sec is reduced by a factor 50. Another interesting study is provided by Dean and Barroso [12] who underline that several redundancy techniques are applied in Google's BigTable in order to improve the latency of incoming queries. They show that a redundancy system with two copies reduces the median response time by 16% and the 99.9th-percentile of the tail of the response time distribution by nearly 40% compared to the non-redundant system.

Broadly speaking, depending on when replicas are deleted, we can consider two classes of redundancy systems: cancel-on-start (*c.o.s.*) and cancel-on-completion (*c.o.c.*). In redundancy systems with

c.o.c., once one of the copies has completed service, the other copies are deleted and the job is said to have received service. In redundancy systems with *c.o.s.*, copies are deleted as soon as one copy starts being served, and as a consequence, *c.o.s.* does not waste any computation resources.

In this survey, we will provide an overview on stability results in redundancy systems. From the point of view of stability, *c.o.s.* does not have any negative impact, and for this reason we focus on stability results when *c.o.c.* is implemented.

Let us illustrate through a simple example how redundancy affects the stability region. Consider a system with K homogeneous servers in which copies of each arriving job are dispatched to $d \leq K$ servers chosen uniformly at random. We assume that jobs arrive according to a Poisson process of rate λ and jobs have general service times with unit mean. Without redundancy, i.e. $d = 1$, the stability condition under any work-conserving policy is given by $\lambda < \mu K$, where μ is the capacity of the servers. Now, let us assume that the service times of copies are i.i.d. and that $d = K$. In this case, the system behaves as a single server system with arrival rate λ and server capacity μK , and the stability condition is again $\lambda < \mu K$. However, if all the copies had the same service time as the original job (identical copies), servers are synchronized and the instantaneous departure rate is just μ . Therefore, the system behaves as a single server system with arrival rate λ and server capacity μ , for which the stability condition is $\lambda < \mu$. This simple example illustrates how the modeling assumptions and the degree of redundancy can dramatically impact the stability condition of the system.

One of the main lessons we draw from the results available in the literature, is that the stability region depends strongly on the scheduling policy employed at the servers and the correlation structure of copies. Somewhat surprisingly, we also identify situations for which it was shown that adding redundant copies does not reduce the stability region. Overall, we believe more research is needed in order to design efficient redundancy algorithms.

The rest of the survey is organized as follows. Section 2 describes the main model assumptions and notation, Section 3 deals with the case in which the service times of the copies are *i.i.d.*, and Section 4 with identical and correlated copies. In Section 5, we present a brief account of results on redundancy that, even though not directly related to stability, are relevant from the performance point of view. We conclude with Section 6 where we discuss several open problems and state various conjectures.

2 Model Description and Preliminaries

We consider a K parallel heterogeneous server system. That is, we have a set of servers $S = \{1, \dots, K\}$ and server s has capacity μ_s , for $s \in S$. Jobs arrive to the system according to a Poisson process of rate λ . Arriving jobs have service times that are independent across jobs and are identically distributed with mean 1.

Jobs are labeled by types $c = \{s_1, \dots, s_i\} \subset S$, where i is the number of copies and c is the set of servers to which this job will dispatch copies. We let \mathcal{C} be the set of all possible types. A job is of type c with probability p_c , where $\sum_{c \in \mathcal{C}} p_c = 1$.

We consider redundancy models that are *c.o.c.*, that is, as soon as a copy is fully served, the additional copies of that job are removed from the system. This cancellation process induces a correlation in the departure process at the servers. Thus, within a server s there is a departure of a copy due to the following two events: *i*) a local copy departs due to completion in server s , or *ii*) a copy in another server completes that induces a departure in server s .

Model Topology. A well-known symmetric topology is the one in which each job sends a copy to d out of K servers. In case the server are chosen uniformly at random, that is, $p_c = 1/\binom{K}{d}$, and servers have the same capacity μ , we refer to this model as the redundancy- d model, see Figure 1 (a). The number of copies, d , is referred to as the redundancy degree.

Two other examples of redundancy topologies are the so-called N -model and W -model, see Figure 1 (b) and (c). Both models are non-symmetric, with two servers. The set of possible job types is $\mathcal{C} =$

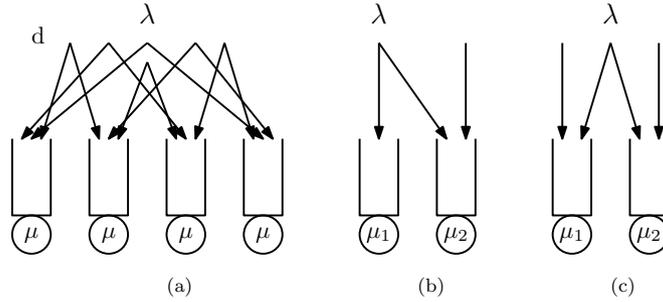


Figure 1: (a) The redundancy- d model for $K = 4$ and $d = 2$. (b) The N -model. (c) The W -model.

$\{\{2\}, \{1, 2\}\}$ in the N -model, and $\mathcal{C} = \{\{1\}, \{2\}, \{1, 2\}\}$ in the W -model.

When no specific structure is assumed, we refer to it in the sequel as a general topology.

Scheduling Policy. A scheduling policy determines how copies are served within each server. As we will see, the choice of the scheduling policy can have a dramatic impact on the stability region. First-Come-First-Served (FCFS) and Processor Sharing (PS) are widely implemented in real-world computer systems ([21]), and are thus common policies considered in the literature on redundancy. Random-Order-of-Service (ROS) is not a common discipline in systems, but as we will see in the ensuing, it yields very good performance in terms of stability for a redundancy system. These three policies represent the main focus of our survey. To the best of our knowledge, other policies such as Last-Come-First-Served (LCFS), Shortest-Remaining-Processing-Time (SRPT), and Least-Attained-Service (LAS) have not been considered so far.

Correlation Structure Among Copies. This describes how the service times of the copies of a given job are related. Formally, the service times X_1, \dots, X_k of the copies of one job can be sampled from a joint distribution $F(x_1, \dots, x_k)$. Two extreme cases are *i.i.d. copies* and *identical copies*. Under *i.i.d. copies*, all copies have independent service times sampled from the same distribution, whereas with *identical copies*, all the copies of a job have the same service time. Another interesting framework is the so-called $S\&X$ model introduced in [16]. Here, the service time of each copy is decomposed into two components; the inherent job size, which is identical for all the copies of a job, and the experienced slowdown on the server it is being served.

Existing Stability Results. Table 1 summarizes the main stability results for *c.o.c.* redundancy models available in the literature and discussed in this survey. The table is organized by scheduling policy, service time distribution, redundancy topology and correlation structure. In brackets we specify the additional assumptions that the authors consider in their respective paper. The term “red- d ” refers to the redundancy- d system and the term “gen.” refers to a general redundancy topology.

The stability condition when jobs have *i.i.d.* copies is the main topic of Section 3, first for exponential service times (Section 3.1) and then for scaled Bernoulli distributions (Section 3.2). These are the results in the first two columns of Table 1. Correlated copies are discussed in Section 4, first for identical copies (Section 4.1, middle two columns in Table 1) and then for general correlation structures (Section 4.2, last two columns of Table 1). In Section 6, we discuss open problems and state various conjectures regarding stability conditions. In Table 1, these conjectures are indicated with a **X**.

3 Independent and Identically Distributed Copies

In this section we assume that jobs have *i.i.d.* copies.

Table 1. Stability results for *c.o.c.* redundancy models.

	Service time dist.	i.i.d. copies		identical copies		General correlation	
		red- <i>d</i>	gen.	red- <i>d</i>	gen.	red- <i>d</i>	gen.
FCFS	Exp.	[17]	[8, 20]	[24] (Mean field)			
	General	[35] (Scaled Bernoulli)		[4]		[32] (Suf. cond.)	[34] (Comparison result)
PS	Exp.	[4]	X	[4]			
	General			[36]	[5]	[36] (Nec. cond.)	
ROS	Exp.	[4]	X	[4]			X
	General						

3.1 Exponential Service Times

We first discuss results on FCFS and exponentially distributed service times, a setting studied by Gardner et al. [17, 20] and Bonald and Comte [8]. It was shown in [8] that this model fits the framework of Order Independent queues (see [28, Chapter 2]), which is a large class of systems that have a product-form steady-state distribution. This can be seen as follows. Since copies are i.i.d., we can describe the system through the Markovian state descriptor $(c_n, c_{n-1}, \dots, c_2, c_1)$. Here, n is the number of jobs in the system, c_1 is the type of the eldest job in the system and c_i is the type of the i th eldest job. Because of FCFS, the eldest job is served in all of its compatible servers c_1 . The i -th eldest job is in service at servers $s \in c_i \setminus \bigcup_{j=1}^{i-1} c_j$, for $i = 1, \dots, n$. Due to the exponentially distributed service times and i.i.d. copies, the instantaneous departure rate of the i th job is given by the sum of the rates in the servers where the job is in service, that is, $\sum_{s \in c_i \setminus c_1, \dots, c_{i-1}} \mu_s$. Hence, the total instantaneous departure rate out of state $(c_n, c_{n-1}, \dots, c_2, c_1)$ is $\sum_{s \in \bigcup_{j=1}^n c_j} \mu_s$, which depends on the set of classes present in the system, but not on their ordering in the state descriptor, i.e., the so-called order independent property.

The characterization of the steady-state distribution facilitates the derivation of performance measures such as the stability condition and mean response times. The proposition below states the stability result for this model.

Proposition 1 ([8, 20]). *For a redundancy system with general topology under FCFS with exponentially distributed service times and i.i.d. copies, the system is stable if for all $C \subseteq \mathcal{C}$,*

$$\lambda \sum_{c \in C} p_c < \sum_{s \in S(C)} \mu_s, \quad (1)$$

where $S(C) = \bigcup_{c \in C} \{s \in c\}$. *The system is unstable if there exists $\tilde{C} \subseteq \mathcal{C}$ such that*

$$\lambda \sum_{c \in \tilde{C}} p_c > \sum_{s \in S(\tilde{C})} \mu_s.$$

Informally, Equation (1) states that the arrival rate to any subset of job types must be less than the total capacity of the associated compatible servers. For exponential service times, this is the *maximum stability* condition, i.e., the system cannot be stable if one of these inequalities were not satisfied. Thus, we conclude that the stability region is not reduced due to adding redundant copies. The latter might seem counter-intuitive at first, since even if servers waste resources serving copies that are not fully served, the stability condition is as large as if there was no redundancy (see also the simple example in the introduction).

Extending Proposition 1 to other scheduling policies is an important open problem (see Section 6 for more details). To the best of our knowledge, this has only been achieved for the redundancy- d model. In this case, it is easy to see that Equation (1) reduces to $\lambda < \mu K$, and it has been shown that this stability condition remains valid when either PS or ROS is implemented.

Proposition 2 ([4]). *For the redundancy- d model under either PS or ROS with exponentially distributed service times and i.i.d. copies, the system is stable when $\lambda < K\mu$ and unstable when $\lambda > K\mu$.*

Hence, under PS, ROS and FCFS, the redundancy- d model is maximum stable. This however does not hold true in general. In the example below (originally in [4]), we describe priority policy that is not maximum stable, i.e., the system can become unstable even though $\lambda < K\mu$.

Example: Priority Policy. Consider the redundancy- d system with $K = 3$, $d = 2$ and $\mu = 1$. There are three different types of jobs: $\mathcal{C} = \{\{1, 2\}, \{1, 3\}, \{2, 3\}\}$. In server 1, FCFS is implemented. In server 2 and server 3, jobs of types $\{1, 2\}$ and $\{1, 3\}$ have preemptive priority over jobs of type $\{2, 3\}$, respectively. Additionally, within a type, jobs are served in order of arrival.

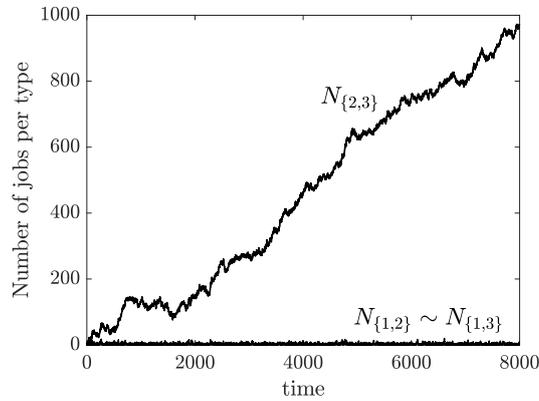


Figure 2: The trajectory of the number of jobs per type when $\lambda = 2.9$.

In Figure 2 we plot the sample-path of the number of jobs when $\lambda = 2.9 < 3 = \mu K$. We observe that the number of type- $\{2, 3\}$ jobs in the system grows large, while the number of type- $\{1, 2\}$ and type- $\{1, 3\}$ jobs stay close to 0. Hence, the system is clearly unstable, even though $\lambda < \mu K$. This can intuitively be explained by the inefficiency induced by the priority mechanism as the type- $\{2, 3\}$ jobs are preempted by type- $\{1, 2\}$ and type- $\{1, 3\}$ jobs in servers 2 and 3, respectively. We refer to [4] for more details.

3.2 General Service Times

To the best of our knowledge, no stability results exist for general service times with i.i.d. copies. In this section, we present the stability result obtained for scaled Bernoulli service times, defined as

$$\begin{cases} X \cdot M, & \text{with probability } 1/M \\ 0, & \text{with probability } 1 - 1/M, \end{cases}$$

where $M > 0$ and X is a strictly positive random variable with $E[X] = 1$. In this setting, Raaijmakers et al. [35] characterize the stability condition for the redundancy- d model where FCFS is implemented and the number of servers grows large.

Proposition 3 ([35]). *Consider the redundancy- d model under FCFS with scaled Bernoulli service times and i.i.d. copies. Then, $\lambda < \frac{M^{d-1}}{E[\min(X_1, \dots, X_d)]}$ is a sufficient stability condition for any M . In addition, for any ϵ , it holds that $(1 - \epsilon)\lambda < \frac{M^{d-1}}{E[\min(X_1, \dots, X_d)]}$ is a necessary condition, for M sufficiently large.*

We observe that the stability condition is independent of the number of servers, but strongly depends on the number of copies d . The latter is in contrast to the exponentially distributed service times, where the stability condition does depend on the number of servers but is independent of d (see Proposition 2). Thus, we observe that when copies are i.i.d., the stability condition strongly depends on the service time distribution. In addition, we observe that as M grows large (and hence the variance of the service times grows large), the stability region increases by a factor M^{d-1} , by taking advantage of a greater diversity in service times.

4 Correlated Copies

Several studies (e.g., [42]) have shown that the i.i.d. copies assumption can be unrealistic, since large jobs remain large when replicated. Hence, having additional copies could lead to high response times and even instability. Motivated by the latter, stability results with correlated copies have been the focus of recent literature.

4.1 Identical Copies

In this section, we assume that jobs have identical copies, i.e., all copies belonging to one job have the same size. This correlation makes that a job can only depart due to its copy that has received most service so far. Thus, the instantaneous departure rate of a job depends on its copy that has currently attained most service.

FCFS Policy. With FCFS, the eldest job in the system will be served at all of its compatible servers. A job later in the queue will be served at its compatible servers that are not engaged by earlier jobs in the queue.

The stability condition for the redundancy- d system with FCFS and exponentially distributed service times is characterized in Anton et al. [4], through the average departure rate per type in the so-called *saturated system*. The latter assumes an infinite backlog of jobs waiting for service. The long-run time-average number of jobs in service in the saturated system is denoted by $\bar{\ell}$. A detailed description of the saturated system and the characterization of $\bar{\ell}$ can be found in [4].

Proposition 4 ([4]). *For the redundancy- d system under FCFS with exponentially distributed service times and identical copies, the system is stable if $\lambda < \bar{\ell}\mu$ and unstable if $\lambda > \bar{\ell}\mu$.*

The value of $\bar{\ell}$, and hence the stability region, can be numerically obtained by solving the balance equations of the saturated system, see [4] for more details. We note that the instantaneous departure rate in the saturated system strongly depends on the types in service. As a consequence, no expression has been derived so far for $\bar{\ell}$ for general K and d values.

Note that the stability condition can equivalently be written as $\frac{\lambda}{K\mu} < \frac{\bar{\ell}}{K}$, where $\frac{\lambda}{K\mu}$ is the traffic load. In Figure 3 (originally in [4]), we provide numerical values for $\frac{\bar{\ell}}{K}$, that is, the traffic supported by the system. The table (*left*) shows $\bar{\ell}/K$ for small values of K and the figure (*right*) plots the value of $\bar{\ell}/K$ as K grows large. To obtain the value of $\bar{\ell}$ for $d \neq 1, K - 2, K - 1, K$, the authors simulate the saturated system, rather than solving the balance equations¹. It was proven in [4] that $\bar{\ell}/K$, hence the

¹When $d = K - 1$, there are d servers that process copies of one job, and the remaining $K - d = 1$ server serves one additional job, hence, $\bar{\ell} = 2$. When instead $d = 1$, there is no redundancy and each server serves one job in the saturated system, i.e., $\bar{\ell} = K$. When $d = K$, the system behaves as a single server with capacity μ , that is, $\bar{\ell} = 1$.

$\bar{\ell}/K$	$K = 2$	$K = 3$	$K = 4$	$K = 5$	$K = 6$	$K = 7$	$K = 8$
$d = 1$	1	1	1	1	1	1	1
$d = 2$	0.5	0.66	0.71	0.74	0.76	0.77	0.77
$d = 3$		0.33	0.5	0.54	0.57	0.58	0.60
$d = 4$			0.25	0.4	0.43	0.46	0.47
$d = 5$				0.2	0.33	0.36	0.38
$d = 6$					0.16	0.28	0.31
$d = 7$						0.14	0.25

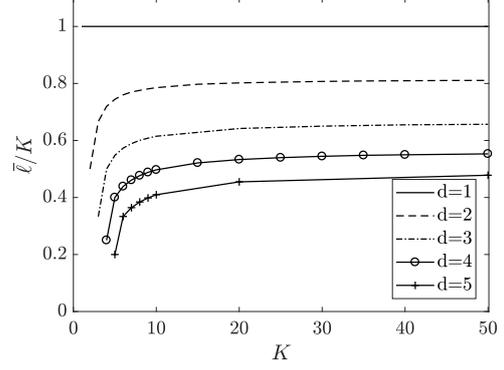


Figure 3: The table and figure show the values of $\bar{\ell}/K$ for different values of d and K .

amount of supported traffic, increases when the number of servers (K) grows large, a property that can be observed in Figure 3.

Processor Sharing Policy. Under PS and identical copies, the stability condition is characterized in [5]. There it is shown that the stability condition coincides with that of a K parallel server system where each type- c job is only dispatched to its so-called least-loaded servers. In order to state this result, we first need to define several sets of servers and customer types. The first subsystem includes all servers, that is $S_1 = S$. We denote by \mathcal{L}_1 the set of least-loaded servers in the system $S_1 = S$. Thus,

$$\mathcal{L}_1 = \left\{ s \in S_1 : s = \arg \min_{\bar{s} \in S_1} \left\{ \frac{1}{\mu_{\bar{s}}} \sum_{c \in \mathcal{C}(\bar{s})} p_c \right\} \right\}.$$

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

$$\begin{aligned} S_i &:= S \setminus \bigcup_{l=1}^{i-1} \mathcal{L}_l, \\ \mathcal{C}_i &:= \{c \in \mathcal{C} : c \subset S_i\}, \\ \mathcal{C}_i(s) &:= \mathcal{C}_i \cap \mathcal{C}(s), \\ \mathcal{L}_i &:= \left\{ s \in S_i : s = \arg \min_{\bar{s} \in S_i} \left\{ \frac{1}{\mu_{\bar{s}}} \sum_{c \in \mathcal{C}_i(\bar{s})} p_c \right\} \right\}. \end{aligned}$$

The S_i -subsystem refers to the system consisting of the servers in S_i , with only jobs of types in the set \mathcal{C}_i . The $\mathcal{C}_i(s)$ is the subset of types that are served in server s in the S_i -subsystem. We let $\mathcal{C}_1 = \mathcal{C}$. The set \mathcal{L}_i represents the set of least-loaded servers in the S_i -subsystem. Finally, we denote by $i^* := \arg \max_{i=1, \dots, K} \{\mathcal{C}_i \neq \emptyset\}$ the last index i for which the subsystem S_i is not empty of job types.

The stability condition is now characterized in [5] by the least-loaded servers that can serve each job type.

Proposition 5 ([5]). *Assume that the service time distribution is such that it has no atoms and is light-tailed in the following sense,*

$$\lim_{r \rightarrow \infty} \sup_{a \geq 0} \mathbf{E}[(X - a)1_{\{X - a > r\}} | X > a] = 0. \quad (2)$$

For a redundancy system with a general topology under PS with identical copies, the system is stable if $\lambda \sum_{c \in \mathcal{C}_i(s)} p_c < \mu_s$, for all $s \in \mathcal{L}_i$, $i = 1, \dots, i^$. The redundancy system is unstable if there exists $\iota \leq i^*$ and $s \in \mathcal{L}_\iota$ such that $\lambda \sum_{c \in \mathcal{C}_\iota(s)} p_c > \mu_s$.*

It can be seen (as observed in [33]) that the light-tailed condition in (2) also implies

$$\sup_{a \geq 0} \mathbf{E}[(X - a) | X > a] \leq \Phi < \infty, \quad (3)$$

which is a usual light-tailed condition (see [14]). Hence, (2) and (3) exclude heavy tailed distributions like Pareto, but include large sets of distributions such as phase type (which are dense in the set of all distributions on \mathbb{R}^+), exponential and hyper-exponential distributions, as well as distributions with bounded support.

For the redundancy- d model, the above stability condition simplifies into $\lambda < K\mu/d$. The latter coincides with the stability condition of a system where all the copies need to be served, that is, the worst possible stability condition.

ROS Policy. When ROS is implemented in the servers, it was shown in [4] that the stability condition is not reduced when adding redundant copies. This was proved for exponentially distributed service times and identical copies for the redundancy- d model. However, as stated in Section 6, we believe that this holds true for any redundancy structure and any correlation structure.

Proposition 6 ([4]). *For the redundancy- d model under ROS with exponentially distributed service times and identical copies, the system is stable if $\lambda < K\mu$.*

The intuition behind the above result is as follows. Whenever there are many jobs in a server, the probability that this server serves a copy of a job that has also a copy elsewhere in service will be close to zero. Hence, with a probability close to 1, all highly-loaded servers are serving copies of different jobs and their instantaneous departure rate equals the sum of their capacities.

4.2 Generally Correlated Copies

In this section, we consider redundancy models where the service times of the copies of each job are correlated according to some general structure.

For FCFS, Raaijmakers et al. [34] consider a general workload model, which subsumes the S&X model, introduced in [17]. The main difference is that in [34] the server capacities are not fixed, but each job samples server capacities from a discrete and finite distribution. The authors assume that the server speed variations (slowdowns) are either distributed according to New-Better-than-Used (NBU) or New-Worse-than-Used (NWU). See [37] for more details on NBU and NWU distributions².

Depending on the random variation in the server speed, the authors prove that either no replication ($d = 1$) or full replication ($d = K$) provides a larger stability region. Note that here the stability region refers to a wider concept than what we considered before. That is, it refers to the set of arrival rates such that there exists a static assignment rule that makes the system stable.

Proposition 7 ([34]). *Consider the following model. Each job is routed to d servers according to some static probabilistic assignment. Servers implement FCFS. Every time a server starts serving a new copy, it samples a speed variation, which is independent across servers. The type of a job is determined by the capacities it would obtain in each server. A job has a generally distributed service time.*

- *If the probabilistic assignment can depend on the job type, and the speed variation follows an NBU distribution, then the stability region for $d = 1$ is larger or equal than that for $d > 1$.*
- *If the probabilistic assignment does not depend on the job type, and the speed variation follows an NWU distribution, then the stability region for $d = K$ is larger or equal than that for $d = 1$.*

² X is said to be *New-Better-than-Used (NBU)* if for all $t_1, t_2 \in \mathbb{R}$, $\bar{F}_X(t_1 + t_2) \leq \bar{F}_X(t_1)\bar{F}_X(t_2)$. X is said to be *New-Worse-than-Used (NWU)* if for all $t_1, t_2 \in \mathbb{R}$, $\bar{F}_X(t_1 + t_2) \geq \bar{F}_X(t_1)\bar{F}_X(t_2)$. A sufficient condition for X to be NBU (NWU) is to have an increasing (a decreasing) hazard rate, i.e., $r(x)$ is increasing (decreasing) in x .

From the above we observe that the optimal redundancy degree does not depend on the job size distributions, but rather on the random variation in the server speeds for a given job among the servers.

A sufficient stability condition for the redundancy- d model with FCFS has been obtained in Mendelson [32]. He considers that the service times of the copies X_1, \dots, X_d are identically distributed with mean 1 and sampled from a joint distribution $F(x_1, \dots, x_d)$.

Proposition 8 ([32]). *Consider the redundancy- d model where FCFS is implemented and the service times of the copies are sampled from a general joint distribution $F(x_1, \dots, x_d)$. Then, $\lambda < \lambda_{lb}$ is a sufficient stability condition, where*

$$\lambda_{lb} := \frac{\mu K}{\sum_{m=0}^d \left(\sum_{j=1}^{d-m} E[\min(X_1, \dots, X_j)] + mE[\min(X_1, \dots, X_d)] \right) P_m},$$

and $P_m = \binom{K-d}{d-m} \binom{d}{m} / \binom{K}{d}$.

For the special cases $d = 1$ and $d = K$, the sufficient condition simplifies to $\lambda < \lambda_{lb} = K\mu$ and $\lambda < \lambda_{lb} = \mu/E[\min(X_1, \dots, X_d)]$, respectively, which are in fact also the necessary stability conditions.

We now consider the redundancy- d model where PS is implemented. Raaijmakers et al. [36] characterize the stability condition under any service time distribution through the minimum of the service times of the copies of a job. The latter can be heuristically explained as follows: assume that all servers are equally loaded. Then, due to PS, the copy that completes first is the one with the smallest service time among all copies of the job.

Proposition 9 ([36]). *For the redundancy- d model under PS where the service times of the copies are sampled from a general joint distribution $F(x_1, \dots, x_d)$, a necessary stability condition is $\lambda d E[\min(X_1, \dots, X_d)] < K\mu$.*

In the particular case where copies are identical, the authors in [36] prove that Proposition 9 gives a sufficient and necessary stability condition, which is given by $\lambda d < K\mu$. We note that the latter coincides with the stability condition for light-tailed service times distributions provided in Proposition 5. Moreover, [36] shows that the stability condition under NWU service time distributions, respectively NBU service time distributions, is larger, respectively smaller, than that for exponential service times.

5 Related Work

In this section, we briefly overview relevant papers on redundancy. Even though the results do not deal directly with stability, they are important pointers for the reader who wishes to work on redundancy.

5.1 Response Time

The response time (a.k.a. delay) measures the time elapsed between arrival and departure. It is together with stability the main performance measure, and it has received considerable attention. The first performance analysis of a redundancy model was for *cancel-on-complete* (*c.o.c.*) with exponentially distributed service times, independent and identically distributed (i.i.d.) copies and FCFS. As discussed in Section 3.1, Gardner et al. [17, 20] and Bonald and Comte [8] exploit the link between this redundancy system and the Order Independent queue [28], in order to show that the steady-state distribution has a product form. The paper [17] showed that the mean response time in the system reduces as the redundancy degree d increases. Redundancy *c.o.s.* with FCFS and exponentially distributed job sizes has been analyzed in Ayesta et al. [7], where it was shown that the steady-state distribution also has

a product form. This was achieved by showing that this model fits within the framework of multi-type jobs and multi-type servers studied in Visschers et al. [40]. The above results have motivated researchers to develop unifying frameworks to explain the emergence of product form distributions in redundancy models. This is done in Ayesta et al. [6] and Gardner and Righter [19] by extending the frameworks of Visschers et al. [40] and Order Independent queues [28], respectively.

Comte and Dorsman [10] introduce the Pass-and-Swap queue, not included in the above unifying frameworks, but for which the product-form of the steady-state distribution is preserved. The authors provide several examples that fall into this framework, including a loss variant of the *c.o.s.* redundancy model.

The response time has also been studied in limiting regimes such as heavy traffic and mean field. Cardinaels et al. [9] consider both *c.o.c.* and *c.o.s.* and establish that in heavy traffic the joint distribution of the number of jobs of the various types converges to the product of an exponentially distributed random variable times a deterministic vector, a phenomenon known as state-space collapse. Hellemans et al. [24, 25] consider the mean-field regime and characterize the stationary workload distribution of *c.o.c.* with FCFS, general service times and both identical and i.i.d. copies. In Hellemans et al. [22] the authors generalize the previous result to other redundancy scheduling implementations such as replication if above certain threshold, delayed replication policy or replicate small jobs. Another mean-field result can be found in Hellemans et al. [23] where the authors analyze the stationary response time and workload distributions of JSW(d), JSQ(d) and redundancy- d under FCFS and general service times.

5.2 Optimizing Redundancy

The stability results presented in this survey show that both the scheduling policy and the degree of redundancy can have a big impact on the stability region and hence on the performance of the system. Motivated by this, researchers have aimed at *i)* characterizing what is the optimal scheduling policy in the servers and *ii)* determining what is the optimum number of copies that should be created.

One of the first papers studying redundancy was by Koole and Righter [27], which considered a system where jobs can dispatch i.i.d. copies to any subset of servers in the system. The authors showed that with FCFS and NWU service time distributions, the best policy is to replicate to all the servers.

Several optimality results have been derived for the Least-Redundant-First (LRF) scheduling policy, which serves jobs in lower priority as their number of copies increases (jobs with the same number of copies are served according to FCFS). In particular, Gardner et al. [15, 18] consider nested redundancy models with exponential service times and i.i.d. copies, and show that the mean response time is minimized under LRF. We note that a redundancy model is nested if for all $c, c' \in \mathcal{C}$, either *i)* $c \subset c'$ or *ii)* $c' \subset c$ or *iii)* $c \cap c' = \emptyset$.

Akgun et al. [1] consider a two-server system in which each server has dedicated traffic, that is, each server is a unique compatible server for one job type. The authors consider the DCF (Dedicated-Customers-First) scheduling policy and analyze the efficiency and fairness for both dedicated and redundant jobs.

Sun et al. [39] consider various low-complexity redundancy scheduling techniques for systems where jobs have i.i.d. copies, and investigate when these are delay-optimal (or nearly-delay optimal) with respect to the stochastic ordering. These new scheduling techniques are based on job replication and job cancellation decision features. For instance, the authors show that the *fewest unassigned task first with low-priority replication* and *earliest due date first with replication* policies are nearly delay-optimal with NBU and NWU distributions, respectively.

5.3 Related Models

Redundancy as considered in this chapter is closely related to the (n, k) fork-join system. In the latter, there exist n servers each one receiving one of the blocks, and the job is completed once $k < n$ blocks

are served. If $k = 1$, this model becomes equivalent to the redundancy- n model with *c.o.c.*.

For the (n, k) fork-join model, Lee et al. [30] provide sequences of systems that upper and lower bound the original one, and that converge to the original system. Through these bounds, the authors characterize the mean response time of the system. Li et al. [31] derive that in the mean-field regime, coding always improves the mean response time compared to the redundancy model, i.e., $(n, 1)$.

In [26], the authors consider the (n, r, k) partial fork-join system, where the job is sent to r out of n servers uniformly chosen at random and waits for the first $k \leq r$ to complete. They study effective replication strategies for various scenarios. The authors show that both latency and cost are minimized when r increases for log-convex (high variable) service time distributions. Duffy et al. [13] compare the tail response time of the (n, r, k) model to that of the redundancy- d model (with batch arrivals of size r). The authors show that the tail distribution of the response time under (n, r, k) partial fork-join is smaller than under the redundancy- d model as long as $r - k \geq d$, as the number of servers tend to infinity.

In a recent paper, Zubeldia [43] considers the $S\&X$ model where the slowdown experienced by each copy in service is independent across servers, but not necessarily independent from the job's service time. The author provides a lower-bound on the mean delay for the (n, r, k) partial fork-join system, and shows that when slowdowns are exponentially distributed and independent of the service time of the job, the expected delay is minimized in the mean-field limit for a constant r that only depends on the arrival rate and mean slowdown.

6 Conclusions, Open Problems and Conjectures

The literature on the stability analysis of redundancy systems is recent and growing. However, there are many important cases that have not been analyzed yet. In this section, we address some of the open problems related to stability, and state several conjectures that are based on our intuitive understanding of the system. It is our hope that this survey might encourage more research on this relevant and timely topic.

6.1 I.i.d. Copies.

As shown in Proposition 1, FCFS is maximum stable with exponential service times and i.i.d. copies. We believe that this result should remain valid for any work-conserving scheduling policy with non-preferential treatment across types, for instance PS, ROS, LCFS, LAS and SRPT. The reason for this is that the i.i.d assumption combined with the non-preferential treatment across types permits to take advantage of diversity when the system is close to saturation.

Conjecture 1. Consider a redundancy system with a general topology with exponentially distributed service times and i.i.d. copies. For any work-conserving non-preferential scheduling policy, the system is stable if for all $C \subseteq \mathcal{C}$,

$$\lambda \sum_{c \in C} p_c < \sum_{s \in S(C)} \mu_s,$$

where $S(C) = \bigcup_{c \in C} \{s \in c\}$.

Open Problem 1. If we relax the exponential service times to general service time distribution, the stability condition is unknown.

6.2 FCFS Scheduling Policy with Identical Copies

In Section 4.1, we saw that $\lambda/\mu K < \bar{\ell}/K$ is the stability condition of the redundancy- d system where jobs have identical copies and exponential service times.

Open Problem 2. If we relax the redundancy- d structure to general topologies, or the exponential service times to general service times, the stability condition is unknown.

For exponential service times with the redundancy- d structure, we observed in Figure 3 that for a given number of copies d , $\lim_{K \rightarrow \infty} \bar{\ell}/K < 1$. Note that $\lambda/\mu K < 1$ is the stability condition for a system with no redundancy. Hence, if it can be proved that $\lim_{K \rightarrow \infty} \bar{\ell}/K < 1$, this would imply that as the number of servers grows large, the traffic load that a redundancy system can support is smaller than if no redundancy was implemented.

Conjecture 2. Consider the redundancy- d model where FCFS is implemented and jobs have exponentially distributed service times and identical copies. Then, for fixed d , $\lim_{K \rightarrow \infty} \bar{\ell}/K < 1$.

The limit should coincide with the stability condition given in [24], where the authors develop a numerical method to derive the stability condition in the mean-field limit.

We also observed the following monotonicity property in the number of redundant copies. More precisely, we conjecture that as the degree of redundancy increases, the stability region becomes smaller.

Conjecture 3. Consider the redundancy- d model where FCFS is implemented and jobs have exponentially distributed service times and identical copies. Then, for fixed K , $\bar{\ell}$ is decreasing in d , and hence, the stability region is decreasing in d .

6.3 ROS Scheduling Policy with Generic Correlation Structure.

In the particular case of ROS, we believe that Conjecture 1 will remain valid even if copies follow a general correlation structure, including identical copies. So far, this was only proved for the redundancy- d model with exponential distributed service times with identical copies, see Proposition 6.

Conjecture 4. Consider a redundancy system with a general topology with exponentially distributed service times and an arbitrary correlation structure among copies. ROS is stable if for all $C \subseteq \mathcal{C}$,

$$\lambda \sum_{c \in C} p_c < \sum_{s \in S(C)} \mu_s,$$

where $S(C) = \bigcup_{c \in C} \{s \in c\}$.

The intuition would be the following. In principle, multiple copies of the same job could be served simultaneously at various of its compatible servers. Due to the heterogeneous capacities and the correlation among the copies, the departure rate of that job depends on the residual service time of each copy. However, when the number of jobs in the system grows large, the probability that more than one copy of the same job is simultaneously in service goes to zero. Using fluid-limit techniques, as done in [4], one then obtains that the fluid limit of the system equals that of the system where jobs have i.i.d. copies. Hence, if Conjecture 1 is valid, this would imply that Conjecture 4 is true as well.

6.4 Redundancy-Aware Scheduling

Another interesting, and so far unexplored area, is the impact of redundancy-aware scheduling policies on the stability region and the performance of the system. By redundancy-aware we refer to policies like LRF or Most-Redundant-First that can use information on the number of copies when choosing which copy to serve in a server. As discussed in Section 5.2, the authors of [15, 18] consider the nested model

with exponentially distributed service times and i.i.d. copies and show that LRF minimizes the mean response time. It would be interesting to explore this further for more general redundancy settings.

7 Acknowledgement

Research of E. Anton supported and research of M. Jonckheere partially supported by the French “Agence Nationale de la Recherche (ANR)” through the project ANR-15-CE25-0004 (ANR JCJC RACON). U. Ayesta has received funding from the Department of Education of the Basque Government through the Consolidated Research Group MATHMODE (IT1294-19).

References

- [1] Akgun, O., Richter, R., Wolff, R.: Partial flexibility in routing and scheduling. *Adv. Appl. Probab.* **45(3)**, 673–691 (2013)
- [2] Ananthanarayanan, G., Ghodsi, A., Shenker, S., Stoica, I.: Why let resources idle? Aggressive cloning of jobs with dolly. In: *Proc. of the 4th USENIX Conf. on Hot Topics in Cloud Computing, HotCloud’ 12*, Article 17, 6 pp. (2012)
- [3] Ananthanarayanan, G., Ghodsi, A., Shenker, S., Stoica, I.: Effective straggler mitigation: attack of the clones. In: *Proc. of the 10th USENIX Conf. on Networked Systems Design and Implementation* **13**, 185–198 (2013)
- [4] Anton, E., Ayesta, U., Jonckheere, M., Verloop, I.M.: On the stability of redundancy models. *Oper. Res.*, *to appear* (2021)
- [5] Anton, E., Ayesta, U., Jonckheere, M., Verloop, I.M.: Improving the performance of heterogeneous data centers through redundancy. In: *Proc. of the ACM on Measurement and Analysis of Computing Systems – SIGMETRICS* **4(3)**, Article 48, 29 pp. (2020)
- [6] Ayesta, U., Bodas, T., Dorsman, J., Verloop, I.M.: A token-based central queue with order-independent service rates. *Oper. Res.*, *to appear* (2021)
- [7] Ayesta, U., Bodas, T., Verloop, I.M.: On a unifying product form framework for redundancy models. *Performance Evaluation* **127-128**, 93–119 (2018)
- [8] Bonald, T., Comte, C.: Balanced fair resource sharing in computer clusters. *Performance Evaluation* **116**, 70–83 (2017)
- [9] Cardinaels, E., Borst, S.C., van Leeuwen, J.S.H.: Redundancy scheduling with locally stable compatibility graphs. *arXiv:2005.14566* (2020)
- [10] Comte, C., Dorsman, J.: Pass-and-swap queues. *arXiv:2009.12299* (2020)
- [11] Dean, J.: Achieving rapid response times in large online services. Google Research 2012, <http://research.google.com/people/jeff/latency.html>, (2012)
- [12] Dean, J., Barroso, L.A.: The tail at scale. *Communications of the ACM* **56**, 74–80 (2013)
- [13] Duffy, K.R., Shneer, S.: MDS coding is better than replication for job completion times. *arXiv:1907.11052* (2019)
- [14] Foss, S., Korshunov, D., Zachary, S.: *An Introduction to Heavy-Tailed and Subexponential Distributions*. Springer, NY (2013)

- [15] Gardner, K., Harchol-Balter, M., Hyytiä, E., Richter, R.: Scheduling for efficiency and fairness in systems with redundancy. *Performance Evaluation* **116**, 1-25 (2017)
- [16] Gardner, K., Harchol-Balter, M., Scheller-Wolf, A., van Houdt, B.: A better model for job redundancy: decoupling server slowdown and job size. *IEEE/ACM Trans. on Networking* **25(6)**, 3353–3367 (2017)
- [17] Gardner, K., Harchol-Balter, M., Scheller-Wolf, A., Velednitsky, M., Zbarsky, S.: Redundancy-d: the power of d choices for redundancy. *Oper. Res.* **65**, 1078–1094 (2017)
- [18] Gardner, K., Hyytiä, E., Richter, R.: A little redundancy goes a long way: convexity in redundancy systems. *Performance Evaluation* **131**, 22-42 (2019)
- [19] Gardner, K., Richter, R.: Product forms for FCFS queueing models with arbitrary server-job compatibilities: an overview. *Queueing Systems* **96(1)**, 3–51 (2020)
- [20] Gardner, K., Zbarsky, S., Doroudi, S., Harchol-Balter, M., Hyytiä, E., Scheller-Wolf, A.: Queueing with redundant requests: exact analysis. *Queueing Systems* **83(3-4)**, 227–259 (2016)
- [21] Harchol-Balter, M.: *Performance Modeling and Design of Computer Systems: Queueing Theory in Action*. Cambridge University Press, NY (2013)
- [22] Hellemans, T., Bodas, T., van Houdt, B.: Performance analysis of workload dependent load balancing policies. In: *Intern. Conf. on Measurement and Modeling of Computer Systems* **3(2)**, Article 35, 35 pp. (2019)
- [23] Hellemans, T., van Houdt, B.: On the Power-of-d-choices with least loaded server selection. In: *Proc. of the ACM on Measurement and Analysis of Computing Systems – SIGMETRICS* **2(2)**, Article 27, 22 pp. (2018)
- [24] Hellemans, T., van Houdt, B.: Analysis of redundancy(d) with identical replicas. *ACM Sigmetrics Performance Evaluation Review* **46(3)**, 74–79 (2018)
- [25] Hellemans, T., van Houdt, B.: Performance of redundancy(d) with identical/independent replicas. In: *ACM Trans. on Modeling and Performance Evaluation of Computing Systems (TOMPECS)* **4(2)**, Article 9, 28 pp. (2019)
- [26] Joshi, G., Soljanin, E., Wornell, G.: Efficient redundancy techniques for latency reduction in cloud systems. In: *ACM Trans. on Modeling and Performance Evaluation of Computing Systems (TOMPECS)* **2(2)**, Article 12, 30 pp. (2017)
- [27] Koole, G., Richter, R.: Resource allocation in grid computing. *J. of Scheduling* **11**, 163–173 (2007)
- [28] Krzesinski, A.E.: Order independent queues. In: Boucherie, R.J., van Dijk, N.M. (eds.) *Queueing Networks: a Fundamental Approach*, pp. 85–120. Springer, Boston, MA (2011)
- [29] Lee, K., Shah, N.B., Huang, L., Ramchandran, K.: When do redundant requests reduce latency? *IEEE Trans. on Communications* **64(2)**, 715–722 (2016)
- [30] Lee, K., Shah, N.B., Huang, L., Ramchandran, K.: The mds queue: analysing the latency performance of erasure codes. *IEEE Trans. on Information Theory* **63(5)**, 2822–2842 (2017)
- [31] Li, B., Ramamoorthy, A., Srikant, R.: Mean-field-analysis of coding versus replication in cloud storage systems. In: *IEEE INFOCOM 2016 - The 35th Annual IEEE Intern. Conf. on Computer Communications*, pp. 1–9 (2016)

- [32] Mendelson, G.: A lower bound on the stability region of redundancy-d with FIFO service discipline. *Oper. Res. Lett.* **49**(1), 113–120 (2021)
- [33] Paganini, F., Tang, A., Ferragut, A., Andrew, L.: Network stability under alpha fair bandwidth allocation with general file size distribution. *IEEE Trans. on Automatic Control* **57**, 579-591 (2012)
- [34] Raaijmakers, Y., Borst, S.C.: Achievable stability in redundancy systems. In: *Proc. of the ACM on Measurement and Analysis of Computing Systems – SIGMETRICS* **4**(3), Article 46, 21 pp. (2020)
- [35] Raaijmakers, Y., Borst, S.C., Boxma, O.: Redundancy scheduling with scaled Bernoulli service requirements. *Queueing Systems* **93**, 67–82 (2019)
- [36] Raaijmakers, Y., Borst, S.C., Boxma, O.: Stability of redundancy systems with processor sharing. In: *Proc. of the 13th EAI Intern. Conf. on Performance Evaluation Methodologies and Tools, Valuetools 20*, pp. 120–127 (2020)
- [37] Ross, S.M.: *Stochastic Processes*. Wiley & Sons, NY (1996)
- [38] Sieber, C., Blenk, A., Hinteregger, M., Kellerer, W.: The cost of aggressive http adaptive streaming: quantifying youtube’s redundant traffic. In: *2015 IFIP/IEEE Intern. Symp. on Integrated Network Management (IM)*, pp. 1261–1267 (2015)
- [39] Sun, Y., Koksals, C.E., Shroff, N.B.: On delay-optimal scheduling in queueing systems with replications. *arXiv:1603.07322* (2016)
- [40] Visschers, J., Adan, I., Weiss, G.: A product form solution to a system with multi-type jobs and multi-type servers. *Queueing Systems* **70**, 269–298 (2012)
- [41] Vulimiri, A., Godfrey, P.B., Mittal, R., Sherry, J., Ratnasamy, S., Shenker, S.: Low latency via redundancy. In: *Proc. of the ACM Conf. on Emerging Networking Experiments and Technologies*, pp. 283–294 (2013)
- [42] Vulimiri, A., Michel, O., Godfrey, P.V., Shenker, S.: More is less: reducing latency via redundancy. In: *Proc. of the 11th ACM Workshop on Hot Topics in Networks, HotNets’11* **11**, pp. 13–18 (2012)
- [43] Zubeldia, M.: Delay-optimal policies in partial fork-join systems with redundancy and random slowdowns. In: *Proc. of the ACM on Measurement and Analysis of Computing Systems – SIGMETRICS* **4**(1), Article 2, 49 pp. (2020)