

# Queueing with Redundant Requests: First Exact Analysis

Kristen Gardner<sup>1</sup>      Samuel Zbarsky<sup>2</sup>  
Sherwin Doroudi<sup>3</sup>      Mor Harchol-Balter<sup>1</sup>  
Esa Hyytiä<sup>4</sup>      Alan Scheller-Wolf<sup>3</sup>

December 2014  
CMU-CS-14-143

School of Computer Science  
Carnegie Mellon University  
Pittsburgh, PA 15213

<sup>1</sup>School of Computer Science, Carnegie Mellon University, Pittsburgh, PA, USA

<sup>2</sup>Department of Mathematics, Carnegie Mellon University, Pittsburgh, PA, USA

<sup>3</sup>Tepper School of Business, Carnegie Mellon University, Pittsburgh, PA, USA

<sup>4</sup>Aalto University, Finland

**Keywords:** Redundant Requests; Markov chain; Task Assignment

## Abstract

Recent computer systems research has proposed using redundant requests to reduce latency. The idea is to run a single request on multiple servers and only wait for the *first* completion (discarding all remaining instances of the request). However no exact analysis of systems with redundancy exists.

This paper presents the first exact analysis of systems with redundancy. We allow for any number of classes of redundant requests, any number of classes of non-redundant requests, any degree of redundancy, and any number of heterogeneous servers. In all cases we derive the limiting distribution on the state of the system.

In small (two or three server) systems, we derive simple forms for the distribution of response time of both the redundant classes and non-redundant classes, and we quantify the “gain” to redundant classes and “pain” to non-redundant classes caused by redundancy. We find some surprising results. First, in many cases the response time of the redundant class follows a simple Exponential distribution and that of the non-redundant class follows a Generalized Hyperexponential. Second, once a class is fully redundant, it becomes “immune” to any pain caused by other classes becoming redundant.

We also compare redundancy with other approaches for reducing latency, such as optimal probabilistic splitting of a class among servers (Opt-Split) and Join-the-Shortest-Queue (JSQ) routing of a class. We find that redundancy outperforms JSQ and Opt-Split with respect to overall response time, making it an attractive solution.



# 1 Introduction

Reducing latency has always been a primary concern for computer systems designers. Recent papers have proposed a new approach to reducing latency in multi-server systems: using *redundant requests* [1, 8, 30]. The motivation behind this approach comes from the observation that response times at servers can be highly variable. Two servers in the same system often differ in their current loads, their network congestion, and the configuration of their storage systems. Even if both servers are idle, the same request might experience a far lower service time at one server than another because the request could experience a much lower disk seek time at one server than another (seek times often dominate service times) [16]. The solution is to send the same request to the queues at multiple servers simultaneously (i.e., redundantly). When any copy of the request completes service, all remaining copies of the request are killed.

Using redundant requests is not free. First, data must be replicated among the set of servers to which the copies must be sent. Furthermore, using redundant requests adds to the system load. Nonetheless, using redundant requests has been shown to vastly improve latency in distributed systems, e.g., Google’s BigTable service shows a 42-fold improvement in tail latency by using redundant requests [10].

Unfortunately, there is almost no work analyzing the benefits of redundant requests. Even a two-server system with one redundant class and one non-redundant class has not been analyzed. The first attempts to analyze systems with redundancy are as recent as 2014, but this work derives only bounds and approximations [18].

Redundant requests require a new queueing paradigm: there is no longer a single copy of each job, and redundant copies disappear instantly as soon as one copy completes. While redundant jobs bear some resemblance to fork-join systems, the two models are actually quite different because all copies must complete service in a fork-join system, whereas a redundant job only needs one copy to complete. Likewise, while redundant jobs bear some resemblance to coupled-processor systems, they differ in that the redundant copies can occupy multiple servers even when these servers have non-empty queues. Likewise, redundant jobs are not the same as flexible servers (see Section 2 for more details).

The state space required to capture systems with redundant jobs is very complex. It is not enough to know the number of jobs in each queue, or even the number of jobs of each class (redundant or non-redundant) within each queue. Rather, one needs to track the exact position and type of every job in every queue, so that one knows which jobs to delete when a copy of a redundant job completes service.

This paper provides the first closed form exact analysis of redundant queues. We derive the limiting distribution of the full queue state as well as (in many cases) the distribution of response time of redundant and non-redundant jobs. Our analysis assumes Exponential service times and Poisson arrivals. Our result applies to systems with any number of queues,  $k$ , any number of classes of jobs,  $\ell$ , and a general redundancy structure (see Figure 1). Here a *class* is associated with a set of servers that hold replicated data, meaning that the jobs of a class can be run on any of the servers associated with the class.




Figure 1: The general redundancy model. Each server  $j$  provides service at rate  $\mu_j$ . Each class of jobs  $C_i$  arrives to the system as a Poisson process with rate  $\lambda_{C_i}$  and joins the queue at all servers in  $S_{C_i} = \{j \mid \text{server } j \text{ can serve class } C_i\}$ .




Figure 2: (a) The  $\mathbb{N}$  model. Class  $A$  jobs join the queue at server 2 only, while class  $R$  jobs are redundant at both servers. (b) The  $\mathbb{W}$  model. Class  $A$  jobs join the queue at server 1 only, class  $B$  jobs join the queue at server 2 only, and class  $R$  jobs are redundant at both servers. (c) The  $\mathbb{M}$  model. Class  $R_1$  jobs are redundant at servers 1 and 2, and class  $R_2$  jobs are redundant at servers 2 and 3.

We also investigate how the approach of redundant requests compares to other common approaches for job assignment. For example, how does making  $k$  redundant copies of each request compare with optimally probabilistically splitting load among  $k$  queues (Opt-Split), or with joining the shortest of  $k$  queues (JSQ)? Furthermore, while there may be a benefit to the redundant class of jobs, what is the response time penalty to the other jobs in the system? Do other approaches create less of a penalty? Finally, if one class of jobs creates redundant copies of itself, does that class suffer when others “join the redundancy game” and start creating redundant copies of their jobs as well?

We start by investigating these questions in the context of three simple models, shown in Figure 2. In the  $\mathbb{N}$  model (Figure 2(a)), there are two arrival streams of jobs, each with its own server. However, one class of jobs creates redundant requests at both servers. The  $\mathbb{N}$  model illuminates the response time benefits to the redundant class and the pain to the non-redundant class. We derive the exact distribution of response time for both classes, and explore what happens when the non-redundant class decides that it too wants to become redundant. In the  $\mathbb{W}$  model (Figure 2(b)), we imagine that we have a stable system, where each server is serving its own stream of jobs, when a new stream of jobs arrives which can be processed at either server. We ask how to best deal with this new stream: redundancy, splitting, dispatching to the shortest queue? We then turn to the  $\mathbb{M}$  model (Figure 2(c)), where there is a “shared server,” which can be used by all request streams. We ask how to best use this shared resource.

After addressing these questions for small systems, we turn to issues of scale; our exact closed-form analysis allows us to consider systems of any size. We investigate how size affects our earlier messages for the  $\mathbb{N}$ ,  $\mathbb{W}$ , and  $\mathbb{M}$  models, particularly those comparing redundancy to other resource sharing policies.

The remainder of this paper is organized as follows. In Section 2, we describe related work and how it differs from the present work. In Section 3, we formalize our model and state our main theorem for the general system. In Sections 4, 5, and 6, we discuss detailed results for the  $\mathbb{N}$ ,  $\mathbb{M}$ , and  $\mathbb{W}$  models. Section 7 addresses how these models scale as the number of servers increases. Finally, in Section 8, we conclude.

## 2 Prior Work

In this section, we review several models that are related to redundant requests. All of these models differ from ours in critical ways that change both the mathematical techniques available to analyze the system, and the results obtained. Nonetheless, we hope that the results in this paper might shed some light on the problems below, many of which are notoriously difficult.

## Coupled Processor/Cycle Stealing

In a *coupled processor* system, there are two servers with rates  $\mu_1$  and  $\mu_2$  and two classes of jobs,  $A$  and  $B$ . Server 1 works on class  $A$  jobs in FCFS order, and server 2 works on class  $B$  jobs in FCFS order. However, if there are only jobs of one class in the system, the servers “couple” to serve that class at a faster rate: unlike in the redundancy model, class  $A$  jobs only get to use server 2 when the system is empty of  $B$ ’s (and vice-versa). Generating functions for the stationary distribution of the queue lengths in a two-server system with Exponential service times were derived in [11, 21], but this required solving complicated boundary value problems and provided little intuition for the performance of the systems. The stationary distribution of the workload in the two-server system was derived in [9] using a similar approach. In [15], a power-series approach was used to numerically compute the queue-length stationary distribution in systems with more than two servers under Exponential service times. Much of the remaining work on coupled processor models involves deriving bounds and asymptotic results (for example, [5]).

In the donor-beneficiary model (one-way cycle stealing), only one class of jobs (the beneficiary) receives access to both servers, typically only when no jobs of the other class are present. In addition, if there is only one beneficiary job present, one server must idle (the servers do not “couple”). The donor-beneficiary model has been studied, in approximation, in a variety of settings [14, 23]. However, it differs from the redundancy model because a job is never in service at more than one server, and because donor jobs often have full preemptive priority at their server.

## Fork-Join

Another related model is the *fork-join* system, in which each job that enters a system with  $k$  servers splits into  $k$  pieces, one of which goes to each server. The job is considered complete only when all  $k$  pieces have completed service. This is different from the redundancy model because only one redundant request needs to finish service in the redundancy model. Furthermore, a fork-join job sends work to all  $k$  servers, whereas a redundant job of class  $C_i$  only sends copies to the servers in  $S_{C_i}$ , where  $S_{C_i}$  is a subset of the servers. The fork-join model is known to be very difficult to analyze. Many papers have derived bounds and approximations for such a system (for example, [2, 3, 20, 22, 32]). Exact analysis remains an open problem except for the two-server case [12, 13]; see [7] for a more detailed overview.

## Flexible Server Systems

A third related model is the *flexible server* system, in which each class of jobs has its own dedicated queue, and each server is capable of serving some subset of the classes of jobs. The design and performance of flexible server systems has been studied in [4, 26, 27, 28]. In a flexible server system, traditionally, when a server becomes available, it *chooses* the queue from which to take its next job according to some policy. By contrast in redundancy systems, each server has its own FCFS queue and jobs are routed to a subset of servers upon arrival. However, the key difference between flexible server systems and redundancy systems is that in a redundancy system, a job may be in service at multiple servers simultaneously, whereas in a flexible server system, each job may be processed by only one server.

A special case of the flexible server system uses the following policy. When a server becomes available, it chooses the job that arrived the earliest from among those jobs that it can serve. This policy is similar to the redundancy system because each server works in overall FCFS order among the jobs it can serve. However, there are no redundant jobs in this flexible server system; jobs cannot be *in service* at two servers at the same time. For this model, under a specific routing assumption when an arriving job sees multiple idle servers, the stationary distribution that satisfies the balance equations is given [29]. In our redundancy model, no such routing assumption is required, because an arriving redundant job will enter service at *all* idle servers. Finally, mean response times are lower in a redundant system than in an FCFS flexible server system; our exact analysis allows us to quantify this performance gap.

## Redundancy Models

Recently, in 2012, the  $(n, k, r)$  system was proposed [24], where there are  $n$  servers, and each arriving job sends a request to  $k \leq n$  of these servers. As soon as  $r \leq k$  requests complete, the job is considered finished. If we view the  $k$  requests as  $k$  “redundant” copies of a job, then this problem can be seen as similar to ours, although in our model, jobs have different levels of redundancy and can only be redundant at certain servers. Various bounds and approximations have been derived for the  $(n, k, r)$  model [17, 18, 24], and the optimal value of  $k$  has been determined for different system loads and costs of deleting unnecessary redundant requests [25]. Additionally, other variants have been proposed where a job might wait before issuing redundant requests [31]. Unfortunately, the only exact analysis of the  $(n, k, r)$  system is for a highly simplified model in which each server is actually an  $M/M/\infty$  queue, so there is no queueing [17].

## 3 Model

We consider a system with  $k$  servers, denoted  $1, 2, \dots, k$ , and  $\ell$  classes of jobs, denoted  $C_1, C_2, \dots, C_\ell$ . (see Figure 1). The service time at server  $j$  is distributed Exponentially with rate  $\mu_j$  for all  $1 \leq j \leq k$ , and each server processes the jobs in its queue in FCFS order. Each class of jobs  $C_i$  arrives to the system as a Poisson process with rate  $\lambda_{C_i}$ , and replicates itself by joining the queue at all servers in the subset  $S_{C_i} = \{j \mid \text{server } j \text{ can serve class } C_i\}$ . Jobs in class  $C_i$  cannot join the queue at any server  $j \notin S_{C_i}$ . A job may be in service at multiple servers at the same time; if a job is in service at both servers  $i$  and  $j$ , it receives service at combined rate  $\mu_i + \mu_j$ .

Looking at Figure 1, it is difficult to figure out an appropriate state space. One might think that you could keep track of the number of jobs of each class at each queue, but this state space is missing information about which specific jobs are in multiple queues. Furthermore, servers are not independent, and job classes are not independent, so typical product-form type state spaces and solutions are unlikely to work.

The key insight that allows us to model this system is that we can view the system as having a single central queue in which all jobs wait in the order that they arrived (see Figure 3). Each server processes jobs from this central queue in FCFS order, skipping over those jobs it cannot serve. For example, in Figure 3, server 3 will skip over job  $A^{(1)}$  and move to job  $B^{(1)}$  when choosing its next job. We can write the state of the system as  $(c_n, c_{n-1}, \dots, c_1)$ , where there are  $n$  jobs in the system, and  $c_i$  is the class of the  $i^{\text{th}}$  job in this central queue;  $c_1$  is the class of the job at the head of the queue, which is also in service at all servers in  $S_{C_1}$ .




Figure 3: Let  $A$  and  $B$  be two job classes (see left), where  $A^{(i)}$  is the  $i^{\text{th}}$  arrival of class  $A$ . We can view the general redundancy system (left) as having a single central queue from which each server works in FCFS order, skipping over those jobs it cannot serve (right). The central queue is an interleaving of the individual servers’ queues, where each job appears only once.



**Theorem 1.** *The limiting probability of being in state  $(c_n, c_{n-1}, \dots, c_1)$  is*

$$\pi_{(c_n, \dots, c_1)} = C \prod_{i=1}^n \frac{\lambda_{c_i}}{\sum_{\substack{m \in \bigcup \\ j \leq i} S_{c_j}} \mu_m},$$

where  $C$  is a normalizing constant.

*Proof.* See Appendix. □

Although  $\pi_{(c_n, c_{n-1}, \dots, c_1)}$  looks like a product-form solution, it is not; we cannot write the limiting probabilities as a product of independent marginal per-server terms, or as a product of independent marginal per-class terms.

**Example 1.** *Consider the system shown in Figure 3. The current state in this system is  $(B, B, A)$ , where the head of the queue is at the right, job  $A^{(1)}$  is currently in service at servers 1 and 2, and job  $B^{(1)}$  is currently in service at server 3. From Theorem 1, we find that the limiting probability of this state is*

$$\pi_{(B, B, A)} = C \frac{\lambda_A}{\mu_1 + \mu_2} \left( \frac{\lambda_B}{\mu_1 + \mu_2 + \mu_3} \right)^2.$$

**Corollary 1.** *The general system with redundant requests is stable when  $\forall C \subseteq \{C_1, \dots, C_\ell\}$ ,*

$$\sum_{C \in \mathcal{C}} \lambda_C < \sum_{\substack{m \in \bigcup \\ C \in \mathcal{C}} S_C} \mu_m. \tag{1}$$

*Proof.* All limiting probabilities  $\pi_{(c_n, \dots, c_1)}$  from Theorem 1 have values in  $(0, 1)$  when  $\forall C \subseteq \{C_1, \dots, C_\ell\}$ , (1) holds. □

In Sections 4, 5, and 6, we use the result of Theorem 1 to study the  $\mathbb{N}$ ,  $\mathbb{W}$ , and  $\mathbb{M}$  models, defined below.

## $\mathbb{N}$ Model

The  $\mathbb{N}$  model is the simplest non-trivial example of a redundancy system where there are both redundant and non-redundant classes. In an  $\mathbb{N}$  model there are two servers running at rates  $\mu_1$  and  $\mu_2$  and two classes of jobs (see Figure 2(a)). Class  $A$  jobs are non-redundant; they arrive with rate  $\lambda_A$  and join the queue at server 2 only ( $S_A = \{2\}$ ). Class  $R$  jobs are redundant; they arrive with rate  $\lambda_R$  and join the queue at both servers ( $S_R = \{1, 2\}$ ).

## $\mathbb{W}$ Model

Consider a two-server, two-class system in which each class of jobs has its own dedicated server (no redundancy). Now suppose that a third class of jobs enters the system and chooses to be redundant at both servers. The  $\mathbb{W}$  model helps us understand how the presence of this redundant class affects the existing non-redundant classes. In a  $\mathbb{W}$  model, there are two servers running at rates  $\mu_1$  and  $\mu_2$  and three classes of jobs (see Figure 2(b)). Class  $A$  jobs arrive with rate  $\lambda_A$  and are served at server 1 only ( $S_A = \{1\}$ ), class  $B$  jobs arrive with rate  $\lambda_B$  and are served at server 2 only ( $S_B = \{2\}$ ), and class  $R$  jobs arrive with rate  $\lambda_R$  and may be served at both servers ( $S_R = \{1, 2\}$ ).

## M Model

Again consider the two-server, two-class system in which each class of jobs has its own dedicated server. Now suppose that a new shared server is added to the system, and all classes of jobs issue redundant requests at this server. The M model helps us understand how best to use this new server. In an M model, there are three servers running at rates  $\mu_1$ ,  $\mu_2$ , and  $\mu_3$  and two classes of jobs (see Figure 2(c)). Class  $R_1$  jobs arrive with rate  $\lambda_{R_1}$  and join the queue at servers 1 and 2 ( $S_A = \{1, 2\}$ ), and class  $R_2$  jobs arrive with rate  $\lambda_{R_2}$  and have  $S_{R_2} = \{2, 3\}$ .

## 4 N Model

We first turn our attention to the N model (Figure 2(a)). An immediate consequence of Theorem 1 is Lemma 1, which gives the limiting distribution of the N model.

**Lemma 1.** *In the N model, the limiting probability of being in state  $(c_n, c_{n-1}, \dots, c_1)$  is:*

$$\pi_{(c_n, \dots, c_1)} = C_{\mathbb{N}} \left( \frac{\lambda_A}{\mu_2} \right)^{a_0} \left( \frac{\lambda_R}{\mu_1 + \mu_2} \right)^r \left( \frac{\lambda_A}{\mu_1 + \mu_2} \right)^{a_1}, \quad (2)$$

where  $a_0$  is the number of class A jobs before the first class R job,  $a_1$  is the number of class A jobs after the first class R job,  $r$  is the total number of class R jobs in the queue, and  $C_{\mathbb{N}} = \frac{(\mu_2 - \lambda_A)(\mu_1 + \mu_2 - \lambda_A - \lambda_R)}{\mu_2(\mu_1 + \mu_2 - \lambda_A)}$  is a normalizing constant.

We use this result to find (Theorem 2) that for the redundant class (class R), response time is Exponentially distributed, which is pleasantly surprising because the system is not an M/M/1. Specifically, the distribution of response time is the same as that in an M/M/1, where the arrival rate is  $\lambda_R$  and the service rate is  $\mu' = \mu_1 + \mu_2 - \lambda_A$ . Note that  $\mu'$  can be viewed as giving the R jobs the full  $\mu_1$ , and the *portion* of  $\mu_2$  that is not appropriated for the class A jobs ( $\mu_2 - \lambda_A$ ). Equivalently, this is the response time in an M/M/1 with arrival rate  $\lambda_A + \lambda_R$  and service rate  $\mu_1 + \mu_2$ .<sup>1</sup>

**Theorem 2.** *In the N model,*

1. *The number of class R jobs in the system,  $N_R$ , is distributed Geometric( $1 - \rho$ ) - 1, where  $\rho = \frac{\lambda_R}{\mu_1 + \mu_2 - \lambda_A}$ .*
2. *The response time of class R jobs,  $T_R$ , is distributed Exp( $\mu_1 + \mu_2 - \lambda_A - \lambda_R$ ).*

*Proof.* This is a special case of the more general result in Theorem 10, which is proved in Appendix B.  $\square$

In Theorem 3, we find that the distribution of response time for the non-redundant class,  $T_A$ , follows a Generalized Hyperexponential distribution<sup>2</sup>. We can view the mean response time of class A jobs as that of an M/M/1 with arrival rate  $\lambda_A$  and service rate  $\mu_2$ , plus a penalty term that captures the extent to which the redundant jobs hurt the A's (Equation 4).

**Theorem 3.** *In the N model,*

1. *The number of class A jobs in the system,  $N_A$ , has p.m.f.*

$$\Pr\{N_A = n_A\} = \zeta_{\mathbb{N}1} \left( \frac{\lambda_A}{\mu_2} \right)^n + \zeta_{\mathbb{N}2} \left( \frac{\lambda_A}{\mu_1 + \mu_2 - \lambda_R} \right)^n, \quad (3)$$

<sup>1</sup>This result is counterintuitive because as we later see (Lemma 3, item 2), the distribution of response time for the redundant class does not depend on whether class A is redundant or non-redundant.

<sup>2</sup>A Generalized Hyperexponential,  $H_2(\nu_1, \nu_2, \omega)$  is defined as the weighted mixture of two Exponentials with rates  $\nu_1$  and  $\nu_2$ , where the first Exponential is given weight  $\omega$  and the second is given weight  $1 - \omega$ . Note that  $\omega$  can be any real number; it need not be a probability [6].




Figure 4: Comparing mean response time before and after class  $R$  becomes redundant when  $\mu_1 = \mu_2 = 1$  and  $\lambda_A = \lambda_R$  for (a) class  $R$ , and (b) class  $A$ . The mean response time for the overall system is the weighted average of these two classes.

where

$$\zeta_{N1} = C_N \left( \frac{\mu_1}{\mu_1 - \lambda_R} \right),$$

$$\zeta_{N2} = C_N \left( \frac{\lambda_R}{\mu_1 + \mu_2 - \lambda_R} - \frac{\lambda_R}{\mu_2 - \lambda_R} \right),$$

and  $C_N$  is as in Lemma 1.

2. The distribution of response time of class  $A$  jobs is

$$T_A \sim H_2(\nu_{N1}, \nu_{N2}, \omega_N),$$

where

$$\begin{aligned} \nu_{N1} &= \mu_2 - \lambda_A \\ \nu_{N2} &= \mu_1 + \mu_2 - \lambda_A - \lambda_R \\ \nu_{N3} &= \mu_1 + \mu_2 - \lambda_A \\ \omega_N &= \frac{\lambda_R \nu_{N1}}{(\mu_1 - \lambda_R) \nu_{N3}}. \end{aligned}$$

The expected response time of class  $A$  jobs is

$$\mathbb{E}[T_A] = \frac{1}{\nu_{N1}} + \frac{1}{\nu_{N2}} - \frac{1}{\nu_{N3}}. \quad (4)$$

*Proof.* Deferred to the end of the section. □

Figure 4 compares mean response time in the system before class  $R$  jobs become redundant (each class sees its own independent M/M/1), and after class  $R$  jobs become redundant. We hold  $\mu_1 = \mu_2 = 1$  and vary the load by increasing  $\lambda_R = \lambda_A$ . We find that redundancy helps class  $R$  jobs by a factor of two (Figure 4(a)), but can hurt the non-redundant class by up to 50% (Figure 4(b)).

In Lemma 3, we ask what happens if class  $A$  jobs decide they too should be redundant. That is, all arriving jobs can be served at both servers - the system is *fully redundant*. This transforms the system into an M/M/1 with arrival rate  $\lambda_A + \lambda_R$  and service rate  $\mu_1 + \mu_2$  (Lemma 2). Surprisingly, class  $R$  is *immune* to pain when class  $A$  also becomes redundant: as Lemma 3 shows, the distribution of response time for class  $R$  is the same before and after class  $A$  becomes redundant. Of course, when the  $A$ 's become redundant, they receive the benefit of having two servers.

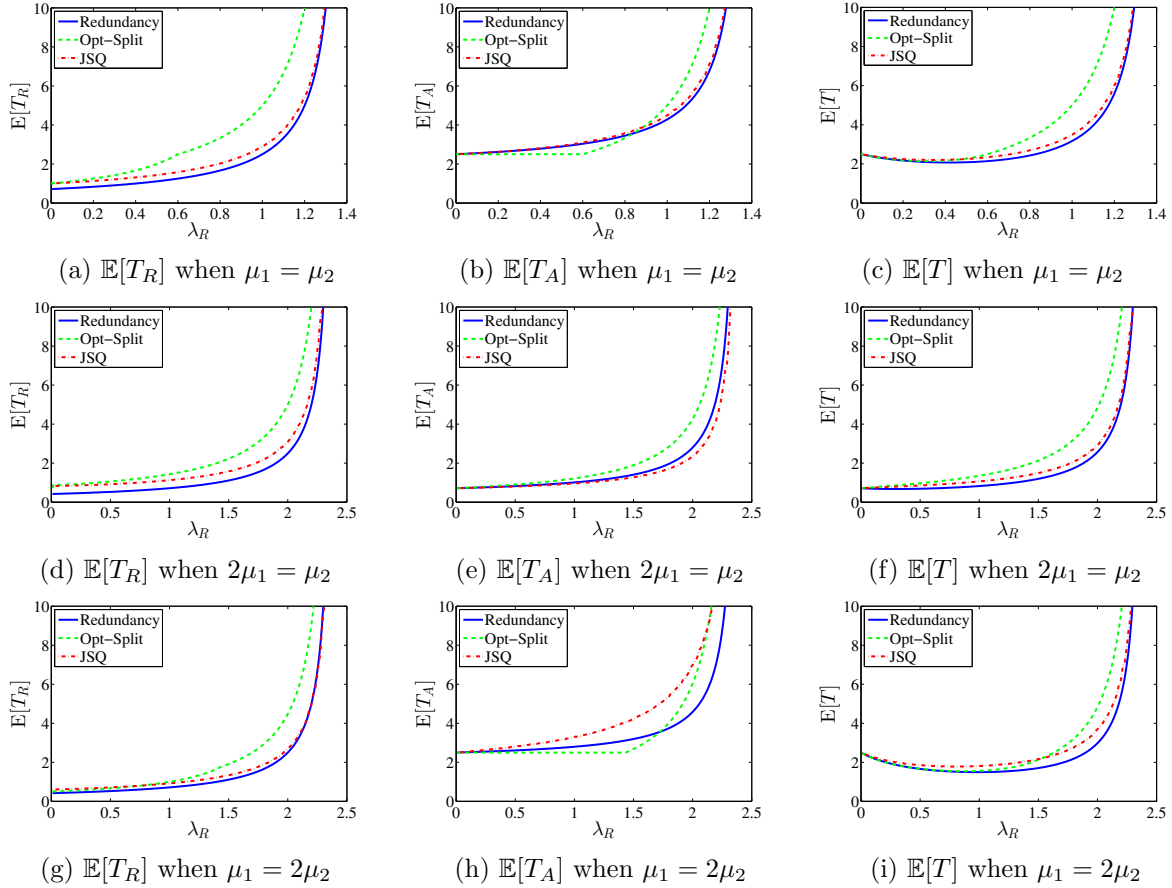


Figure 5: Comparing redundancy, Opt-Split, and JSQ for the  $\mathbb{N}$  model as  $\lambda_R$  increases with  $\lambda_A = 0.6$ . We plot mean response time for the (redundant)  $R$  class (left column), the (non-redundant)  $A$  class (middle column), and the overall system (right column). Rows represent different ratios of server speeds. Redundancy is shown in solid blue, Opt-Split in dashed green, and JSQ in dashed red.

**Lemma 2.** *The fully redundant system, in which both class  $R$  and class  $A$  jobs issue redundant requests at both servers, is equivalent to an  $M/M/1$  with arrival rate  $\lambda_A + \lambda_R$  and service rate  $\mu_1 + \mu_2$ .*

*Proof.* Deferred to the end of the section. □

**Lemma 3.** *With respect to mean response time, both classes of jobs do at least as well in the fully redundant model as in the  $\mathbb{N}$  model. In particular,*

1.  $\mathbb{E}[T_A]^{\text{Fully Redundant}} \leq \mathbb{E}[T_A]^{\text{Redundant}}$
2.  $T_R^{\text{Fully Redundant}} \stackrel{d}{=} T_R^{\text{Redundant}}$ .

*Proof.* Deferred to the end of the section. □

Redundancy clearly helps the redundant class considerably. But there are alternative latency-reducing strategies. For example, each class could optimally probabilistically split jobs among all allowable servers (Opt-Split), or join the shortest queue among allowable servers (JSQ).

In Figure 5, we compare these other options for the  $\mathbb{N}$  model, where the mean response times under Opt-Split are derived analytically (Definition 1), but JSQ is simulated. We find that, for the redundant class  $R$ , redundancy beats JSQ, which beats Opt-Split. Redundancy often is not much better than JSQ, yet they can differ by a factor of 2, depending on the load of class  $R$  and the relative server speeds.

Surprisingly, the non-redundant class,  $A$ , often prefers redundancy of the other class to Opt-Split or JSQ. This is because the non-redundant class wants the redundant class to spend as little time as possible blocking the  $A$  jobs at server 2, and redundancy helps with this.

Note that under Opt-Split we see an inflection point in mean response time for both class  $R$  and class  $A$ . For example, in Figures 5(a) and (b), there is an inflection point at  $\lambda_R = 0.6$ , when  $\lambda_R = \lambda_A$ . This phase change occurs because when  $\lambda_R < \lambda_A$ , no class  $R$  jobs go to server 2 under Opt-Split, but when  $\lambda_R > \lambda_A$  the  $R$ 's compete with the  $A$ 's. Also observe that  $\mathbb{E}[T]$  is not monotonically increasing; this is because as  $\lambda_R$  increases, the redundant class contributes more to the weighted average.

Unsurprisingly, from the overall system's perspective, redundancy is always preferable to Opt-Split and JSQ because it optimizes overall server utilization.

When  $\mu_1 = \mu_2$ , even when non-redundant jobs prefer Opt-Split, redundancy is never more than 50% worse than Opt-Split for the non-redundant jobs (Theorem 4).

**Definition 1.** Under Opt-Split,  $p$  fraction of class  $R$  jobs go to server 2, and  $1 - p$  fraction go to server 1, where  $p$  is chosen to minimize  $\mathbb{E}[T]$ . The mean response times for class  $R$  jobs, class  $A$  jobs, and the system are respectively:

$$\begin{aligned}\mathbb{E}[T_R]^{\text{Opt-Split}} &= \frac{1-p}{\mu_1 - (1-p)\lambda_R} + \frac{p}{\mu_2 - \lambda_A - p\lambda_R} \\ \mathbb{E}[T_A]^{\text{Opt-Split}} &= \frac{1}{\mu_2 - \lambda_A - p\lambda_R} \\ \mathbb{E}[T]^{\text{Opt-Split}} &= \frac{\lambda_A}{\lambda_A + \lambda_R} \mathbb{E}[T_A]^{\text{Opt-Split}} + \frac{\lambda_R}{\lambda_A + \lambda_R} \mathbb{E}[T_R]^{\text{Opt-Split}}.\end{aligned}$$

**Theorem 4.** If  $\mu_1 = \mu_2$ , then the following are true:

1.  $\frac{1}{2} \leq \frac{\mathbb{E}[T_R]^{\text{Redundant}}}{\mathbb{E}[T_R]^{\text{Opt-Split}}} \leq 1$ . If  $\lambda_R > \lambda_A$ , then  $\frac{\mathbb{E}[T_R]^{\text{Redundant}}}{\mathbb{E}[T_R]^{\text{Opt-Split}}} = \frac{1}{2}$ .
2.  $\frac{\mathbb{E}[T_A]^{\text{Redundant}}}{\mathbb{E}[T_A]^{\text{Opt-Split}}} \leq \frac{3}{2}$ .
3.  $\frac{1}{2} \leq \frac{\mathbb{E}[T]^{\text{Redundant}}}{\mathbb{E}[T]^{\text{Opt-Split}}} \leq 1$ .

*Proof.* Deferred to the end of the section. □

## 4.1 Proofs for $\mathbb{N}$ Model

*Proof.* [**Theorem 3**] We wish to compute  $\Pr\{N_A = n_A\}$ . First we consider the case  $n_A = 0$ . Then we have:

$$\Pr\{N_A = 0\} = C \sum_{i=0}^{\infty} \left( \frac{\lambda_R}{\mu_1 + \mu_2} \right)^i = C \frac{1}{1 - \frac{\lambda_R}{\mu_1 + \mu_2}} = C \frac{\mu_1 + \mu_2}{\mu_1 + \mu_2 - \lambda_R}. \quad (5)$$

Now assume  $n_A > 0$ . We consider three cases:

1. There are no  $R$  jobs in the system. Then

$$\Pr\{N_A = n_A \text{ and no } R \text{ jobs in system}\} = C \left( \frac{\lambda_A}{\mu_2} \right)^{n_A}. \quad (6)$$

2. There are both  $R$  and  $A$  jobs in the system, and there is an  $R$  at the head of the queue. Let  $r_0 + 1$  be the number of  $R$  jobs before the first  $A$ , and  $r_i$  be the number of  $R$  jobs following the  $i$ th  $A$ ,  $i > 0$ . Then

$$\begin{aligned}
\Pr\{N_A = n_A, R \text{ at head}\} &= \sum_{r_0=0}^{\infty} \sum_{r_1=0}^{\infty} \cdots \sum_{r_{n_A}=0}^{\infty} C \left( \frac{\lambda_R}{\mu_1 + \mu_2} \right)^{1+r_0+\cdots+r_{n_A}} \left( \frac{\lambda_A}{\mu_1 + \mu_2} \right)^{n_A} \\
&= C \left( \frac{\lambda_A}{\mu_1 + \mu_2} \right)^{n_A} \left( \sum_{r_0=0}^{\infty} \left( \frac{\lambda_R}{\mu_1 + \mu_2} \right)^{r_0} \right) \cdots \left( \sum_{r_{n_A}=0}^{\infty} \left( \frac{\lambda_R}{\mu_1 + \mu_2} \right)^{r_{n_A}} \right) \left( \frac{\lambda_R}{\mu_1 + \mu_2} \right) \\
&= C \left( \frac{\lambda_A}{\mu_1 + \mu_2} \right)^{n_A} \left( \frac{1}{1 - \frac{\lambda_R}{\mu_1 + \mu_2}} \right)^{n_A+1} \left( \frac{\lambda_R}{\mu_1 + \mu_2} \right) \\
&= C \left( \frac{\lambda_A}{\mu_1 + \mu_2} \right)^{n_A} \left( \frac{\mu_1 + \mu_2}{\mu_1 + \mu_2 - \lambda_R} \right)^{n_A+1} \left( \frac{\lambda_R}{\mu_1 + \mu_2} \right) \\
&= C \left( \frac{\lambda_R}{\mu_1 + \mu_2 - \lambda_R} \right) \left( \frac{\lambda_A}{\mu_1 + \mu_2 - \lambda_R} \right)^{n_A} \tag{7}
\end{aligned}$$

3. There are both  $R$  and  $A$  jobs in the system, and there is an  $A$  at the head of the queue. Let  $\ell + 1 = a_0$  be the number of  $A$  jobs before the first  $R$ , let  $r_1 + 1$  be the number of  $R$ 's following these initial  $A$ 's, and let  $r_i$  be the number of  $R$ 's following the  $\ell + i$ th  $A$ ,  $i > 1$ . Then

$$\begin{aligned}
\Pr\{N_A = n_A, A \text{ at head}\} &= \sum_{\ell=0}^{n_A-1} \sum_{r_1=0}^{\infty} \cdots \sum_{r_{n_A-\ell}=0}^{\infty} C \left( \frac{\lambda_A}{\mu_2} \right)^{\ell+1} \left( \frac{\lambda_R}{\mu_1 + \mu_2} \right)^{1+r_1+\cdots+r_{n_A-\ell}} \left( \frac{\lambda_A}{\mu_1 + \mu_2} \right)^{n_A-\ell-1} \\
&= C \left( \frac{\lambda_R}{\mu_2} \right)^{n_A-1} \sum_{\ell=0}^{n_A-1} \left( \frac{\lambda_A}{\mu_2} \right)^{\ell} \left( \frac{\lambda_A}{\mu_1 + \mu_2} \right)^{n_A-\ell} \left( \sum_{r_1=0}^{\infty} \left( \frac{\lambda_R}{\mu_1 + \mu_2} \right)^{r_1} \right) \cdots \left( \sum_{r_{n_A-\ell}=0}^{\infty} \left( \frac{\lambda_R}{\mu_1 + \mu_2} \right)^{r_{n_A-\ell}} \right) \\
&= C \left( \frac{\lambda_R}{\mu_2} \right)^{n_A-1} \sum_{\ell=0}^{n_A-1} \left( \frac{\lambda_A}{\mu_2} \right)^{\ell} \left( \frac{\lambda_A}{\mu_1 + \mu_2} \right)^{n_A-\ell} \left( \frac{1}{1 - \frac{\lambda_R}{\mu_1 + \mu_2}} \right)^{n_A-\ell} \\
&= C \left( \frac{\lambda_R}{\mu_2} \right)^{n_A-1} \sum_{\ell=0}^{n_A-1} \left( \frac{\lambda_A}{\mu_2} \right)^{\ell} \left( \frac{\lambda_A}{\mu_1 + \mu_2 - \lambda_R} \right)^{n_A-\ell} \\
&= C \left( \frac{\lambda_R}{\mu_2} \right) \frac{\mu_2 \left[ \left( \frac{\lambda_A}{\mu_2} \right)^{n_A} - \left( \frac{\lambda_A}{\mu_1 + \mu_2 - \lambda_R} \right)^{n_A} \right]}{\mu_1 - \lambda_R} \\
&= C \left( \frac{\lambda_R}{\mu_1 - \lambda_R} \right) \left[ \left( \frac{\lambda_A}{\mu_2} \right)^{n_A} - \left( \frac{\lambda_A}{\mu_1 + \mu_2 - \lambda_R} \right)^{n_A} \right] \tag{8}
\end{aligned}$$

Finally, we add (6), (7), and (8) to get the result in (3).

We now obtain the Laplace transform of the response time for class  $A$  jobs,  $\tilde{T}_A(s)$ , via distributional Little's Law [19]. First, we find the  $z$ -transform of the number of class  $A$  jobs in the system,  $\hat{N}_A(z)$ :

$$\begin{aligned}
\hat{N}_A(z) &= \sum_{i=0}^{\infty} \Pr\{N_A = i\} z^i \\
&= \zeta_{N1} \sum_{i=0}^{\infty} \left( \frac{\lambda_A}{\mu_2} \right)^i z^i + \zeta_{N2} \sum_{i=0}^{\infty} \left( \frac{\lambda_A}{\mu_1 + \mu_2 - \lambda_R} \right)^i z^i \\
&= \frac{\zeta_{N1} \mu_2}{\mu_2 - \lambda_A z} + \frac{\zeta_{N2} (\mu_1 + \mu_2 - \lambda_R)}{\mu_1 + \mu_2 - \lambda_A - \lambda_R z}.
\end{aligned}$$

Now, observe that class  $A$  jobs depart the system in the same order in which they arrive, so  $A_{T_A}$ , the number of class  $A$  arrivals during a class  $A$  response time, is equivalent to  $N_A$ , the number of jobs seen by an  $A$  departure. Then since  $\widehat{A}_{T_A}(z) = \widetilde{T}_A(\lambda_A - \lambda_A z)$ , we have:

$$\widetilde{T}_A(\lambda_A - \lambda_A z) = \widehat{N}_A(z) = \frac{\zeta_{N1}\mu_2}{\mu_2 - \lambda_A z} + \frac{\zeta_{N2}(\mu_1 + \mu_2 - \lambda_R)}{\mu_1 + \mu_2 - \lambda_A - \lambda_R z}.$$

Let  $s = \lambda_A - \lambda_A z$ , so  $z = 1 - \frac{s}{\lambda_A}$ . Then we have

$$\begin{aligned} \widetilde{T}_A(s) &= \frac{\zeta_{N1}\mu_2}{\mu_2 - \lambda_A(1 - \frac{s}{\lambda_A})} + \frac{\zeta_{N2}(\mu_1 + \mu_2 - \lambda_R)}{\mu_1 + \mu_2 - \lambda_A - \lambda_R(1 - \frac{s}{\lambda_A})} \\ &= (1 + \omega_N) \frac{\nu_{N1}}{\nu_{N1} + s} - \omega_N \frac{\nu_{N2}}{\nu_{N2} + s}. \end{aligned}$$

This is the transform of a Generalized Hyperexponential distribution,  $H_2(\nu_{N1}, \nu_{N2}, \omega_N)$ . Finally,

$$\begin{aligned} \mathbb{E}[T_A] &= -\widetilde{T}'_A(s)|_{s=0} \\ &= - \left[ (1 + \omega_N) \frac{-\nu_{N1}}{(\nu_{N1} + s)^2} + \omega_N \frac{\nu_{N2}}{(\nu_{N2} + s)^2} \right] \Big|_{s=0} \\ &= \frac{1}{\nu_{N1}} + \frac{1}{\nu_{N2}} - \frac{1}{\nu_{N3}}. \quad \square \end{aligned}$$

□

*Proof. [Lemma 2]* In the fully redundant model, all jobs enter the FCFS queue at both servers and depart from both servers immediately upon completion at either server. Hence each job enters service at both servers at exactly the same moment, and receives service rate  $\mu_1 + \mu_2$  for the entire time it is in service. This is exactly an M/M/1 with arrival rate  $\lambda_A + \lambda_B$  and service rate  $\mu_1 + \mu_2$ . □

*Proof. [Lemma 3]* From Theorem 2, we know that in the  $\mathbb{N}$  model,  $T_R \sim \text{Exp}(\mu_1 + \mu_2 - \lambda_A - \lambda_B)$ , which is exactly the distribution of response time in an M/M/1 with arrival rate  $\lambda_A + \lambda_R$  and service rate  $\mu_2 + \mu_2$ .

From Theorem 3, we know that in the  $\mathbb{N}$  model,

$$\mathbb{E}[T_A] = \frac{1}{\mu_1 + \mu_2 - \lambda_A - \lambda_R} + \frac{1}{\mu_2 - \lambda_A} - \frac{1}{\mu_1 + \mu_2 - \lambda_A},$$

which is at least the mean response time in an M/M/1 with arrival rate  $\lambda_A + \lambda_R$  and service rate  $\mu_1 + \mu_2$  since  $\frac{1}{\mu_2 - \lambda_A} - \frac{1}{\mu_1 + \mu_2 - \lambda_A}$  is nonnegative. □

*Proof. [Theorem 4]* Definition 1 gives us  $\mathbb{E}[T_R]^{\text{Opt-Split}}$ ,  $\mathbb{E}[T_A]^{\text{Opt-Split}}$ , and  $\mathbb{E}[T]^{\text{Opt-Split}}$ . We know  $\mathbb{E}[T_A]^{\text{Redundant}}$  from Theorem 3. Theorem 2 tells us that

$$T_R^{\text{Redundant}} \sim \text{Exp}(\mu_1 + \mu_2 - \lambda_A - \lambda_R),$$

so we also know that  $\mathbb{E}[T_R]^{\text{Redundant}} = \frac{1}{\mu_1 + \mu_2 - \lambda_A - \lambda_R}$ . Finally,

$$\mathbb{E}[T]^{\text{Redundant}} = \frac{\lambda_R}{\lambda_A + \lambda_R} \mathbb{E}[T_R]^{\text{Redundant}} + \frac{\lambda_A}{\lambda_A + \lambda_R} \mathbb{E}[T_A]^{\text{Redundant}}.$$

Thus,  $\frac{\mathbb{E}[T_R]^{\text{Redundant}}}{\mathbb{E}[T_R]^{\text{Opt-Split}}}$ ,  $\frac{\mathbb{E}[T_A]^{\text{Redundant}}}{\mathbb{E}[T_A]^{\text{Opt-Split}}}$ , and  $\frac{\mathbb{E}[T]^{\text{Redundant}}}{\mathbb{E}[T]^{\text{Opt-Split}}}$ , and the desired results follow after some minor algebra. □

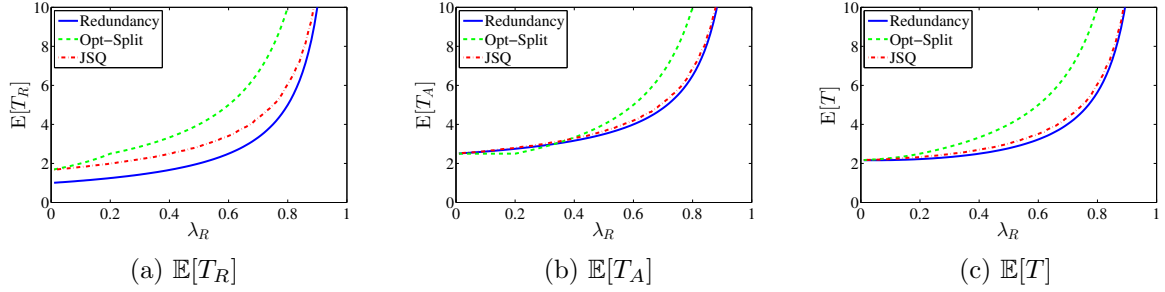


Figure 6: Comparing redundancy, Opt-Split, and JSQ for the  $\mathbb{W}$  model as  $\lambda_R$  increases, for  $\mu_1 = \mu_2 = 1$ ,  $\lambda_A = 0.6$ , and  $\lambda_B = 0.4$ . Lines shown include mean response time under redundancy (solid blue), Opt-Split (dashed green) and JSQ (dashed red) for (a) class  $R$ , (b) class  $A$ , and (c) the system. Results for other values of  $\mu_1$  and  $\mu_2$  are similar.

## 5 $\mathbb{W}$ Model

We now consider the  $\mathbb{W}$  model (see Figure 2(b)). The  $\mathbb{W}$  model has two non-redundant classes,  $A$  and  $B$ , each with its own server. A third class,  $R$ , enters the system and issues redundant requests at both servers. We study how this redundant class affects the performance of the system.

An immediate consequence of Theorem 1 is Lemma 4, which gives the limiting distribution of the  $\mathbb{W}$  model.

**Lemma 4.** *In the  $\mathbb{W}$  model, the limiting probability of being in state  $(c_n, c_{n-1}, \dots, c_1)$  depends on  $c_1$ , as follows:*

$$\begin{aligned}
 \pi_{(c_n, \dots, A)} &= C_{\mathbb{W}} \left( \frac{\lambda_A}{\mu_1} \right)^{a_0} \left( \frac{\lambda_A}{\mu_1 + \mu_2} \right)^{a_1} \left( \frac{\lambda_B}{\mu_1 + \mu_2} \right)^{b_1} \left( \frac{\lambda_R}{\mu_1 + \mu_2} \right)^r \\
 \pi_{(c_n, \dots, B)} &= C_{\mathbb{W}} \left( \frac{\lambda_B}{\mu_2} \right)^{b_0} \left( \frac{\lambda_A}{\mu_1 + \mu_2} \right)^{a_1} \left( \frac{\lambda_B}{\mu_1 + \mu_2} \right)^{b_1} \left( \frac{\lambda_R}{\mu_1 + \mu_2} \right)^r \\
 \pi_{(c_n, \dots, R)} &= C_{\mathbb{W}} \left( \frac{\lambda_A}{\mu_1 + \mu_2} \right)^{a_1} \left( \frac{\lambda_B}{\mu_1 + \mu_2} \right)^{b_1} \left( \frac{\lambda_R}{\mu_1 + \mu_2} \right)^r,
 \end{aligned} \tag{9}$$

where  $a_0$  is the number of class  $A$  jobs before the first class  $B$  or  $R$  job,  $b_0$  is the number of class  $B$  jobs before the first class  $A$  or  $R$  job,  $a_1$  (respectively,  $b_1$ ) is the number of class  $A$  (class  $B$ ) jobs after the first job of class  $R$  or  $B$  ( $A$ ),  $r$  is the total number of class  $R$  jobs, and

$$C_{\mathbb{W}} = \frac{(\mu_1 - \lambda_A)(\mu_2 - \lambda_B)(\mu_1 + \mu_2 - \lambda_A - \lambda_B - \lambda_R)}{\mu_1 \mu_2 (\mu_1 + \mu_2 - \lambda_A - \lambda_B)}$$

is a normalizing constant.

Like in the  $\mathbb{N}$  model, the redundant class (class  $R$ ) has an Exponentially distributed response time (Theorem 5). This is again surprising because the system is not an M/M/1. Nonetheless, the response time for the redundant class is stochastically equivalent to the response time in an M/M/1 with arrival rate  $\lambda_R$  and service rate  $\mu' = \mu_1 + \mu_2 - \lambda_A - \lambda_B$ . We can interpret  $\mu'$  as the remaining service capacity in the system after  $\lambda_A$  and  $\lambda_B$  have been apportioned to classes  $A$  and  $B$  respectively. Alternatively, we can view the response time for the redundant class as that in an M/M/1 with arrival rate  $\lambda_A + \lambda_B + \lambda_R$  and service rate  $\mu_1 + \mu_2$ .

**Theorem 5.** *In the  $\mathbb{W}$  model,*

1. *The number of class  $R$  jobs in the system,  $N_R$ , is distributed Geometric( $1 - \rho$ ) - 1, where  $\rho = \frac{\lambda_R}{\mu_1 + \mu_2 - \lambda_A - \lambda_B}$ .*



2. The response time of class  $R$  jobs,  $T_R$ , is distributed  $\text{Exp}(\mu_1 + \mu_2 - \lambda_A - \lambda_B - \lambda_R)$ .

*Proof.* The proof follows the same approach as that of Theorem 10, and is omitted.  $\square$

In Theorem 6, we derive the distribution of response time for the non-redundant class  $A$  (class  $B$  is symmetric). Again like in the  $\mathbb{N}$  model, we find that  $T_A$  follows a Generalized Hyperexponential distribution. In addition, the mean response time of class  $A$  (or class  $B$ ) jobs can be interpreted as that in an M/M/1 with arrival rate  $\lambda_A$  and service rate  $\mu_1$ , plus a penalty term that captures the extent to which the redundant class hurts the  $A$ 's (or  $B$ 's) (Equation 10). This penalty is the same for class  $A$  and class  $B$  even if they have different loads: the pain caused by the redundant class is shared equally among the non-redundant classes.

**Theorem 6.** *In the  $\mathbb{W}$  model,*

1. The number of class  $A$  jobs in the system,  $N_A$ , has p.m.f.

$$\Pr\{N_A = n_A\} = \zeta_{\mathbb{W}1} \left(\frac{\lambda_A}{\mu_1}\right)^{n_A} + \zeta_{\mathbb{W}2} \left(\frac{\lambda_A}{\mu_1 + \mu_2 - \lambda_B - \lambda_R}\right)^{n_A},$$

where

$$\begin{aligned} \zeta_{\mathbb{W}1} &= \frac{(\mu_1 - \lambda_A)(\mu_2 - \lambda_B)(\mu_1 + \mu_2 - \lambda_A - \lambda_B - \lambda_R)}{\mu_1(\mu_2 - \lambda_B - \lambda_R)(\mu_1 + \mu_2 - \lambda_A - \lambda_B)} \\ \zeta_{\mathbb{W}2} &= \frac{-\lambda_R(\mu_1 - \lambda_A)(\mu_1 + \mu_2 - \lambda_A - \lambda_B - \lambda_R)}{(\mu_2 - \lambda_B - \lambda_R)(\mu_1 + \mu_2 - \lambda_A - \lambda_B)}. \end{aligned}$$

2. The distribution of response time of class  $A$  jobs is

$$T_A \sim H_2(\nu_{\mathbb{W}1}, \nu_{\mathbb{W}2}, \omega_{\mathbb{W}}),$$

where

$$\begin{aligned} \nu_{\mathbb{W}1} &= \mu_1 - \lambda_A \\ \nu_{\mathbb{W}2} &= \mu_1 + \mu_2 - \lambda_A - \lambda_B - \lambda_R \\ \omega_{\mathbb{W}} &= \frac{(\mu_2 - \lambda_B)\nu_{\mathbb{W}2}}{(\mu_2 - \lambda_B - \lambda_R)\nu_{\mathbb{W}3}}. \end{aligned}$$

The mean response time of class  $A$  jobs is

$$\mathbb{E}[T_A] = \frac{1}{\nu_{\mathbb{W}1}} + \frac{1}{\nu_{\mathbb{W}2}} - \frac{1}{\nu_{\mathbb{W}3}}, \quad (10)$$

where

$$\nu_{\mathbb{W}3} = \mu_1 + \mu_2 - \lambda_A - \lambda_B.$$

*Proof.* The proof follows the same approach as that of Theorem 3, and is omitted.  $\square$

The introduction of a new redundant class clearly hurts the existing non-redundant classes, because the new redundant jobs compete for service with the non-redundant jobs. We now ask what would happen if class  $R$  chose which queue(s) to join according to some alternative policy, for example, Opt-Split or JSQ.

In Figure 6, we compare these options, where the mean response time under Opt-Split is derived analytically (Definition 2), but JSQ is simulated. We find that for the redundant class, redundancy outperforms JSQ, which in turn outperforms Opt-Split (Figure 6(a)).

For the non-redundant classes (Figure 6(b)), mean response time is often lower under redundancy than under Opt-Split or JSQ, particularly under higher loads of redundant jobs. This is because even though a larger number of  $R$  jobs compete with class  $A$  at server 1, some of these  $R$  jobs depart the system without ever using server 1 (they complete service at server 2 before entering service at server 1), and some of these  $R$

jobs receive service on both servers at once, thus departing the system faster. As in the  $\mathbb{N}$  model, redundancy is always better for the overall system (Figure 6(c)).

When the servers are homogeneous, in the few cases in which mean response time of class  $A$  or  $B$  is lower under Opt-Split than under redundancy, we show that redundancy is never more than 50% worse for the  $A$  or  $B$  jobs.

**Definition 2.** Under Opt-Split,  $p$  fraction of class  $R$  jobs go to server 1, and  $1 - p$  fraction go to server 2, where  $p$  is chosen to minimize  $\mathbb{E}[T]$ . The mean response times for class  $R$  jobs, class  $A$  jobs, and the overall system are:

$$\begin{aligned}\mathbb{E}[T_R]^{\text{Opt-Split}} &= \frac{p}{\mu_1 - \lambda_A - p\lambda_R} + \frac{1-p}{\mu_2 - \lambda_B - (1-p)\lambda_R} \\ \mathbb{E}[T_A]^{\text{Opt-Split}} &= \frac{1}{\mu_1 - \lambda_A - p\lambda_R} \\ \mathbb{E}[T]^{\text{Opt-Split}} &= \frac{\lambda_A}{\lambda_A + \lambda_B + \lambda_R} \mathbb{E}[T_A]^{\text{Opt-Split}} + \frac{\lambda_B}{\lambda_A + \lambda_B + \lambda_R} \mathbb{E}[T_B]^{\text{Opt-Split}} + \frac{\lambda_R}{\lambda_A + \lambda_B + \lambda_R} \mathbb{E}[T_R]^{\text{Opt-Split}}.\end{aligned}$$

The mean response time for class  $B$  is symmetric to that of class  $A$ .

**Theorem 7.** If  $\mu_1 = \mu_2$ , then the following are true:

1.  $\frac{1}{2} \leq \frac{\mathbb{E}[T_R]^{\text{Redundant}}}{\mathbb{E}[T_R]^{\text{Opt-Split}}} \leq 1$ . If  $\lambda_R \geq |\lambda_A - \lambda_B|$ , then  $\frac{\mathbb{E}[T_R]^{\text{Redundant}}}{\mathbb{E}[T_R]^{\text{Opt-Split}}} = \frac{1}{2}$ .
2.  $\frac{1}{2} \leq \frac{\mathbb{E}[T_A]^{\text{Redundant}}}{\mathbb{E}[T_A]^{\text{Opt-Split}}} \leq \frac{3}{2}$ .
3.  $\frac{1}{2} \leq \frac{\mathbb{E}[T]^{\text{Redundant}}}{\mathbb{E}[T]^{\text{Opt-Split}}} \leq 1$ .

*Proof.* We have  $\mathbb{E}[T_R]^{\text{Opt-Split}}$ ,  $\mathbb{E}[T_A]^{\text{Opt-Split}}$ , and  $\mathbb{E}[T]^{\text{Opt-Split}}$  from Definition 2. We also know  $\mathbb{E}[T_A]^{\text{Redundant}}$  from Theorem 6. Theorem 5 tells us that

$$T_R^{\text{Redundant}} \sim \text{Exp}(\mu_1 + \mu_2 - \lambda_A - \lambda_B - \lambda_R),$$

so

$$\mathbb{E}[T_R]^{\text{Redundant}} = \frac{1}{\mu_1 + \mu_2 - \lambda_A - \lambda_B - \lambda_R}.$$

Finally,

$$\mathbb{E}[T]^{\text{Redundant}} = \frac{\lambda_R}{\lambda_A + \lambda_B + \lambda_R} \mathbb{E}[T_R]^{\text{Redundant}} + \frac{\lambda_A}{\lambda_A + \lambda_B + \lambda_R} \mathbb{E}[T_A]^{\text{Redundant}} + \frac{\lambda_B}{\lambda_A + \lambda_B + \lambda_R} \mathbb{E}[T_B]^{\text{Redundant}}.$$

We use these expressions to find  $\frac{\mathbb{E}[T_R]^{\text{Redundant}}}{\mathbb{E}[T_R]^{\text{Opt-Split}}}$ ,  $\frac{\mathbb{E}[T_A]^{\text{Redundant}}}{\mathbb{E}[T_A]^{\text{Opt-Split}}}$ , and  $\frac{\mathbb{E}[T]^{\text{Redundant}}}{\mathbb{E}[T]^{\text{Opt-Split}}}$ , and the desired results follow after some minor algebra.  $\square$

## 6 $\mathbb{M}$ Model

Finally, we consider the  $\mathbb{M}$  model (Figure 2(c)). Unlike the  $\mathbb{N}$  and  $\mathbb{W}$  models, there are two redundant classes in an  $\mathbb{M}$  model, classes  $R_1$  and  $R_2$ . We study how to best use a shared server at which both classes issue redundant requests. For convenience, throughout the remainder of this section we use the notation

$$\begin{aligned}\mu_{1,2,3} &= \mu_1 + \mu_2 + \mu_3 \\ \mu_{1,2} &= \mu_1 + \mu_2 \\ \mu_{2,3} &= \mu_2 + \mu_3.\end{aligned}$$

An immediate consequence of Theorem 1 is Lemma 5, which gives the limiting distribution of the  $\mathbb{M}$  model.

**Lemma 5.** *In the  $\mathbb{M}$  model, the limiting probability of being in state  $(c_n, c_{n-1}, \dots, c_1)$  depends on  $c_1$ , as follows:*

$$\begin{aligned}\pi_{(c_n, \dots, R_1)} &= C_{\mathbb{M}} \left( \frac{\lambda_{R_1}}{\mu_{1,2}} \right)^{r_{1,0}} \left( \frac{\lambda_{R_1}}{\mu_{1,2,3}} \right)^{r_{1,1}} \left( \frac{\lambda_{R_2}}{\mu_{1,2,3}} \right)^{r_{2,1}} \\ \pi_{(c_n, \dots, R_2)} &= C_{\mathbb{M}} \left( \frac{\lambda_{R_2}}{\mu_{2,3}} \right)^{r_{2,0}} \left( \frac{\lambda_{R_1}}{\mu_{1,2,3}} \right)^{r_{1,1}} \left( \frac{\lambda_{R_2}}{\mu_{1,2,3}} \right)^{r_{2,1}}\end{aligned}\tag{11}$$

where  $r_{1,0}$  (respectively,  $r_{2,0}$ ) is the number of class  $R_1$  ( $R_2$ ) jobs before the first class  $R_2$  ( $R_1$ ) job,  $r_{1,1}$  (respectively,  $r_{2,1}$ ) is the number of class  $R_1$  ( $R_2$ ) jobs after the first  $R_2$  ( $R_1$ ) job, and

$$C_{\mathbb{M}} = \frac{(\mu_{1,2} - \lambda_{R_1})(\mu_{1,2,3} - \lambda_{R_1} - \lambda_{R_2})(\mu_{2,3} - \lambda_{R_2})}{\mu_{1,2}\mu_{2,3}(\mu_{1,2,3} - \lambda_{R_1} - \lambda_{R_2}) + \lambda_{R_1}\lambda_{R_2}\mu_2}$$

is a normalizing constant.

In Theorem 8, we derive the distribution of response time for class  $R_1$  (class  $R_2$  is symmetric). Again, the response time for class  $R_1$  follows a Generalized Hyperexponential distribution. The mean response time for class  $R_1$  can be seen as the mean response time in an M/M/1 with arrival rate  $\lambda_{R_1}$  and service rate  $\mu_{1,2}$ , plus a penalty term because class  $R_1$  has to share server 2 with class  $R_2$ . The penalty incurred by class  $R_1$  and class  $R_2$  differs in proportion to their relative arrival rates.

Note that in the  $\mathbb{N}$  and  $\mathbb{W}$  models, the redundant class had Exponentially distributed response time and the response time distribution for *non-redundant* classes was a Generalized Hyperexponential, whereas in the  $\mathbb{M}$  model, the *redundant* class has a Generalized Hyperexponential response time distribution. We hypothesize that the response time distribution is related to the degree of redundancy: fully redundant classes see Exponentially distributed response time, and partially redundant or non-redundant classes see Generalized Hyperexponentially distributed response times.

**Theorem 8.** *In the  $\mathbb{M}$  model,*

1. *The number of class  $R_1$  jobs,  $N_{R_1}$  has p.m.f.*

$$\Pr\{N_{R_1} = n\} = \zeta_{\mathbb{M}1} \left( \frac{\lambda_{R_1}}{\mu_{1,2}} \right)^n + \zeta_{\mathbb{M}2} \left( \frac{\lambda_{R_1}}{\mu_{1,2,3} - \lambda_{R_2}} \right)^n, \tag{12}$$

where

$$\begin{aligned}\zeta_{\mathbb{M}1} &= C_{\mathbb{M}} \frac{\mu_3}{\mu_3 - \lambda_{R_2}} \\ \zeta_{\mathbb{M}2} &= C_{\mathbb{M}} \left( \frac{\lambda_{R_2}}{\mu_{2,3} - \lambda_{R_2}} - \frac{\lambda_{R_1}}{\mu_{1,2,3} - \lambda_{R_2}} \right).\end{aligned}$$

2. *The distribution of response time of class  $R_1$  jobs is*

$$T_{R_1} \sim H_2(\nu_{\mathbb{M}1}, \nu_{\mathbb{M}2}, \omega_{\mathbb{M}}),$$

where

$$\begin{aligned}\nu_{\mathbb{M}1} &= \mu_{1,2} - \lambda_{R_1} \\ \nu_{\mathbb{M}2} &= \mu_{1,2,3} - \lambda_{R_1} - \lambda_{R_2} \\ \omega_{\mathbb{M}} &= \zeta_{\mathbb{M}1} \frac{\mu_{1,2}}{\mu_{1,2} - \lambda_{R_1}}\end{aligned}$$

*Proof.* The proof follows the same approach as that of Theorem 3, and is omitted.  $\square$




Figure 7: Comparing redundancy, Opt-Split, and JSQ for the  $\mathbb{M}$  model. Lines shown include mean response time for the overall system under redundancy (solid blue), Opt-Split (dashed green), and JSQ with tiebreaking in favor of the faster server (dashed red). Mean response time as a function of (a) increasing  $\lambda_{R_1} = \lambda_{R_2}$ , (b) increasing  $\mu_2$ .

Both classes obviously benefit from issuing redundant requests on a shared server rather than each class having a single dedicated server. However, one might wonder whether mean response time could be further reduced by using some other policy, like Opt-Split or JSQ, instead of redundancy. In Figure 7 we investigate the relative performance of these alternative policies. Mean response time under Opt-Split is derived analytically (Definition 3); JSQ is simulated.

**Definition 3.** Under Opt-Split,  $p$  fraction of class  $R_1$  jobs go to server 2, and  $1 - p$  fraction go to server 1, and  $q$  fraction of class  $R_2$  jobs go to server 2, and  $1 - q$  fraction go to server 3. We choose  $p$  and  $q$  to minimize the overall mean response time, given by

$$\mathbb{E}[T]^{\text{Opt-Split}} = \frac{(1-p)\lambda_{R_1}}{\lambda_{R_1} + \lambda_{R_2}} \frac{1}{\mu_1 - (1-p)\lambda_{R_1}} + \frac{p\lambda_{R_1} + q\lambda_{R_2}}{\lambda_{R_1} + \lambda_{R_2}} \frac{1}{\mu_2 - p\lambda_{R_1} - q\lambda_{R_2}} + \frac{(1-q)\lambda_{R_2}}{\lambda_{R_1} + \lambda_{R_2}} \frac{1}{\mu_3 - (1-q)\lambda_{R_2}}.$$

In all cases, redundancy outperforms both Opt-Split and JSQ. For homogeneous servers (Figure 7(a)), mean response time under JSQ approaches that under redundancy at high load, but at low load, redundancy is better by a factor of 2. For heterogeneous servers (Figure 7(b)), as the service rate of the shared server increases, mean response time under Opt-Split approaches that under redundancy (Theorem 9), but JSQ is worse by a factor of 2. As the system is symmetric, the response times of the individual classes are the same as that of the overall system, and thus are not shown.

We analytically prove performance bounds for the  $\mathbb{M}$  model:

**Theorem 9.** In the  $\mathbb{M}$  model,

1. If  $\mu_1 = \mu_2 = \mu_3$ , then for any  $\lambda_{R_1} = \lambda_{R_2}$  such that the system is stable,  $\frac{1}{3} \leq \frac{\mathbb{E}[T]^{\text{Redundant}}}{\mathbb{E}[T]^{\text{Opt-Split}}} \leq \frac{1}{2}$ .
2. For any  $\mu_1, \mu_3, \lambda_{R_1}$ , and  $\lambda_{R_2}$  such that the system is stable,  $\lim_{\mu_2 \rightarrow \infty} \frac{\mathbb{E}[T]^{\text{Redundant}}}{\mathbb{E}[T]^{\text{Opt-Split}}} = 1$ .
3. For any  $\mu_1, \mu_2, \mu_3, \lambda_{R_1}$ , and  $\lambda_{R_2}$  such that the system is stable,  $\frac{\mathbb{E}[T]^{\text{Redundant}}}{\mathbb{E}[T]^{\text{Opt-Split}}} \leq 1$ .

*Proof.* We know  $\mathbb{E}[T]^{\text{Opt-Split}}$  from Definition 3, and we have

$$\mathbb{E}[T]^{\text{Redundant}} = \frac{\lambda_{R_1}}{\lambda_{R_1} + \lambda_{R_2}} \mathbb{E}[T_{R_1}]^{\text{Redundant}} + \frac{\lambda_{R_2}}{\lambda_{R_1} + \lambda_{R_2}} \mathbb{E}[T_{R_2}]^{\text{Redundant}},$$

where we know  $\mathbb{E}[T_{R_1}]^{\text{Redundant}}$  and  $\mathbb{E}[T_{R_2}]^{\text{Redundant}}$  from Theorem 8. We can then write  $\frac{\mathbb{E}[T]^{\text{Redundant}}}{\mathbb{E}[T]^{\text{Opt-Split}}}$ , and the desired results follow after some minor algebra.  $\square$

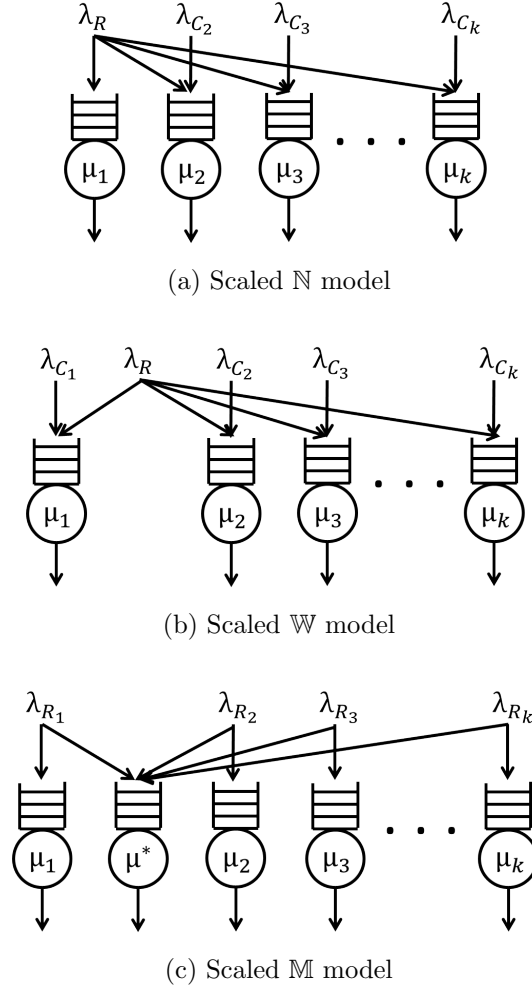


Figure 8: Scaled versions of (a) the  $\mathbb{N}$  model, (b) the  $\mathbb{W}$  model, and (c) the  $\mathbb{M}$  model.

## 7 Scale

Thus far, we only have considered systems with two servers (the  $\mathbb{N}$  and  $\mathbb{W}$  models) and three servers (the  $\mathbb{M}$  model). We now turn our attention to the question of scale.

The scaled  $\mathbb{N}$ ,  $\mathbb{W}$ , and  $\mathbb{M}$  models are shown in Figure 8. In the scaled  $\mathbb{N}$  model, there are  $k$  servers and  $k$  classes of jobs (see Figure 8(a)). Class  $R$  jobs replicate at all servers, while jobs from class  $C_i$  join only the queue at server  $i$  for  $2 \leq i \leq k$ . The scaled  $\mathbb{W}$  model is similar; there are  $k$  servers and  $k + 1$  classes of jobs, with class  $R$  replicating at all servers, and class  $C_i$  going only to server  $i$ ,  $1 \leq i \leq k$  (see Figure 8(b)). In the scaled  $\mathbb{M}$  model, each class  $R_i$ ,  $1 \leq i < k$ , joins the queue at its own dedicated server and at a single server shared by all classes (see Figure 8(c)).

The limiting probabilities derived in Theorem 1 for the general redundancy system apply to the scaled  $\mathbb{N}$ ,  $\mathbb{W}$ , and  $\mathbb{M}$  models. In Theorem 10, we use this result to find that in both the scaled  $\mathbb{N}$  and  $\mathbb{W}$  models, response time for class  $R$  is Exponentially distributed, extending the results of Theorem 2 and Theorem 5 respectively.

**Theorem 10.**




Figure 9: Comparing  $\mathbb{E}[T_R]$  under redundancy (solid blue), Opt-Split (dashed green), and JSQ (dashed red) in scaled systems with homogeneous servers, all with rate 1. (a) The scaled  $\mathbb{N}$  model with  $\lambda_{C_i} = 0.6$  for all non-redundant classes, and  $\lambda_R = 1.2$ . (b) The scaled  $\mathbb{W}$  model with  $\lambda_{C_i} = 0.6$  for all non-redundant classes, and  $\lambda_R = 0.7$ . (c) The scaled  $\mathbb{M}$  model with  $\lambda_{R_i} = 0.6$  for all classes.

1. In the scaled  $\mathbb{N}$  model, the distribution of the number of class  $R$  jobs in the system is

$$N_R \sim \text{Geometric} \left( 1 - \frac{\lambda_R}{\sum_{i=1}^k \mu_i - \sum_{i=2}^k \lambda_i} \right) - 1,$$

and the distribution of the response time of class  $R$  jobs is

$$T_R \sim \text{Exp} \left( \sum_{i=1}^k \mu_i - \sum_{i=2}^k \lambda_i - \lambda_R \right).$$

2. In the scaled  $\mathbb{W}$  model, the distribution of the number of class  $R$  jobs in the system is

$$N_R \sim \text{Geometric} \left( 1 - \frac{\lambda_R}{\sum_{i=1}^k \mu_i - \sum_{i=1}^k \lambda_i} \right) - 1,$$

and the distribution of the response time of class  $R$  jobs is

$$T_R \sim \text{Exp} \left( \sum_{i=1}^k \mu_i - \sum_{i=1}^k \lambda_i - \lambda_R \right).$$

*Proof.* See Appendix. □

For the  $\mathbb{M}$  model and for the non-redundant classes in the  $\mathbb{N}$  and  $\mathbb{W}$  models, the result from Theorem 1 does not easily yield a closed-form expression in the scaled models. The results discussed in the remainder of this section for these classes are obtained via simulation.

For the two-server  $\mathbb{N}$  and  $\mathbb{W}$  models, we saw that the redundant class had lower mean response time under redundancy than under both JSQ and Opt-Split, but often JSQ was very close to redundancy. Here, for scaled models, we investigate whether redundancy enjoys a greater advantage over JSQ and Opt-Split as the number of servers increases.

Indeed, we find that the redundant class sees a much greater benefit under redundancy than under Opt-Split and JSQ as  $k$  increases for the scaled  $\mathbb{N}$  and  $\mathbb{W}$  models (see Figures 9(a) and (b)). In fact, as  $k$  increases, the benefit grows unboundedly because when a class  $R$  job enters a system with high  $k$ , it tends to see many idle servers. Under Opt-Split, this job may not be routed to one of the the idle servers. Under JSQ, the job goes to a *single* idle server  $i$  and receives mean response time  $\frac{1}{\mu_i}$ . Under redundancy, the job gets to use *all* of the idle servers, thereby receiving mean response time  $\frac{1}{\sum_i \mu_i}$ .

In the two-server  $\mathbb{N}$  and  $\mathbb{W}$  models, we saw that the benefit that class  $R$  received from redundancy came at a cost to the non-redundant class  $A$ . In the scaled  $\mathbb{N}$  and  $\mathbb{W}$  models, this cost approaches 0 because the pain caused by the redundant class is spread among all non-redundant classes, so the effect on any one of these classes is minimal; the response time for each non-redundant class  $C_i$  approaches that of an M/M/1 with arrival rate  $\lambda_{C_i}$  and service rate  $\mu_i$ .

In the three-server version of the  $\mathbb{M}$  model (Section 6), we saw that redundancy significantly outperformed Opt-Split and JSQ. In Figure 9(c), we look at the relative performance of the three policies as  $k$  increases. In the scaled  $\mathbb{M}$  model, at low  $k$ , redundancy indeed gives a lower mean response time than Opt-Split and JSQ. However, as  $k$  increases, response time becomes the same under all three policies. As the load on the shared server becomes high, no class benefits from this server; each class experiences an independent M/M/1. Convergence to  $k$  independent M/M/1 queues is slow; for example, at  $k = 200$ , redundancy still provides a 5% lower mean response time than independent M/M/1 queues.

## 8 Conclusion

In this paper we study a multi-server system with *redundant requests*. In such a system, each job that arrives joins the queue at some subset of the servers and departs the system as soon as it completes service at one of these servers. While recent empirical work in computer systems has demonstrated that redundant requests can greatly reduce response time, theoretical analysis of systems with redundancy has proved challenging.

We present the first exact analysis of systems with redundancy, deriving the limiting distribution of the queue state. Our state space is very complex and furthermore yields a non-product form, and non-obvious, limiting distribution. Nonetheless, we find very clean, simple results for response time distributions for both redundant and non-redundant classes. Many of our results are counterintuitive:

1. The redundant class experiences a response time distribution identical to that in an M/M/1, even though the system is not an M/M/1 ( $\mathbb{N}$  and  $\mathbb{W}$  models).
2. Once a class is fully redundant, it is *immune* to additional classes becoming redundant: the distribution of its response time does not change ( $\mathbb{N}$  and  $\mathbb{W}$  model).
3. The non-redundant class often prefers the other class to be redundant as opposed to routing the other class according to Opt-Split or JSQ ( $\mathbb{N}$  and  $\mathbb{W}$  models).
4. When multiple classes share a single server, redundancy can improve mean response time relative to Opt-Split and JSQ by a factor of 2 ( $\mathbb{M}$  model).
5. As the number of servers increases, redundancy gives an even greater benefit to the redundant class while causing less pain to the non-redundant classes (scaled  $\mathbb{N}$  and scaled  $\mathbb{W}$  models).

The redundancy system is closely related to many other queueing models for which exact analysis has long been elusive: Coupled processor systems, fork-join systems, and systems with flexible servers all bear a resemblance to redundancy systems in that they all involve jobs that can be processed by multiple servers. The specific mechanism that determines which jobs run on which servers, and whether jobs can run simultaneously on multiple servers, varies between models, but all of these models share the underlying theme of flexibility. We hope that the new analysis presented in this paper will provide insights on how to analyze these other difficult systems.

While the results presented here apply to a very general redundancy system, we do make some simplifying assumptions, such as Poisson arrivals and Exponential service times. One can ask how our results would change if these assumptions were relaxed. In addition, in real systems there may be some cost to deleting redundant requests once one request has completed. The present work provides the analytical foundation needed to investigate these questions and others, which we leave for future work.

## References

- [1] G. Ananthanarayanan, A. Ghodsi, S. Shenker, and I. Stoica. Why let resources idle? Aggressive cloning of jobs with Dolly. In *HotCloud*, June 2012.
- [2] F. Baccelli and A. Makowski. Simple computable bounds for the fork-join queue. Technical Report RR-0394, Inria, 1985.
- [3] F. Baccelli, A. M. Makowski, and A. Shwartz. The fork-join queue and related systems with synchronization constraints: Stochastic ordering and computable bounds. *Advances in Applied Probability*, 21:629–660, 1989.
- [4] A. Bassamboo, R. S. Randhawa, and J. A. V. Mieghem. A little flexibility is all you need: On the value of flexible resources in queueing systems. *Operations Research*, 60:1423–1435, 2012.
- [5] S. Borst, O. Boxma, and M. V. Uitert. The asymptotic workload behavior of two coupled queues. *Queueing Systems*, 43(1-2):81–102, January 2003.
- [6] R. F. Botta, C. M. Harris, and W. G. Marchal. Characterizations of generalized hyperexponential distribution functions. *Communications in Statistics, Stochastic Models*, 3(1):115–148, 1987.
- [7] O. Boxma, G. Koole, and Z. Liu. Queueing-theoretic solution methods for models of parallel and distributed systems. In *Performance Evaluation of Parallel and Distributed Systems Solution Methods. CWI Tract 105 & 106*, pages 1–24, 1994.
- [8] H. Casanova. Benefits and drawbacks of redundant batch requests. *Journal of Grid Computing*, 5(2):235–250, February 2007.
- [9] J. W. Cohen and O. J. Boxma. *Boundary Value Problems in Queueing System Analysis*. North-Holland Publishing Company, 1983.
- [10] J. Dean. Achieving rapid response times in large online services. Slides from talk given at Berkeley AMPLab Cloud Seminar, March 2012.
- [11] G. Fayolle and R. Iasnogorodski. Two coupled processors: The reduction to a Riemann-Hilbert problem. *Zeitschrift für Wahrscheinlichkeitstheorie und verwandte Gebiete*, 47(3):325–351, 1979.
- [12] L. Flatto. Two parallel queues created by arrivals with two demands II. *SIAM Journal on Applied Mathematics*, 45(5):1159–1166, October 1985.
- [13] L. Flatto and S. Hahn. Two parallel queues created by arrivals with two demands I. *SIAM Journal on Applied Mathematics*, 44(5):250–255, October 1984.
- [14] M. Harchol-Balter, C. Li, T. Osogami, A. Scheller-Wolf, and M. Squillante. Cycle stealing under immediate dispatch task assignment. In *Annual Symposium on Parallel Algorithms and Architectures*, pages 274–285, June 2003.
- [15] G. Hooghiemstra, M. Keane, and S. V. de Ree. Power series for stationary distributions of coupled processor models. *SIAM Journal on Applied Mathematics*, 48(5):861–878, October 1988.
- [16] H. Huang, W. Hung, and K. G. Shin. FS2: dynamic data replication in free disk space for improving disk performance and energy consumption. In *Proc. of SOSP’05*, pages 263–276, December 2005.
- [17] G. Joshi, Y. Liu, and E. Soljanin. Coding for fast content download. In *Allerton Conference’12*, pages 326–333, 2012.
- [18] G. Joshi, Y. Liu, and E. Soljanin. On the delay-storage trade-off in content download from coded distributed storage systems. *IEEE Journal on Selected Areas in Communications*, 32(5):989–997, May 2014.
- [19] J. Keilson and L. Servi. A distributional form of Little’s Law. *Operations Research Letters*, 7(5):223–227, 1988.
- [20] C. Kim and A. K. Agrawala. Analysis of the fork-join queue. *IEEE Transactions on Computers*, 38(2):1041–1053, February 1989.
- [21] A. G. Konheim, I. Meilijson, and A. Melkman. Processor-sharing of two parallel lines. *Journal of Applied Probability*, 18(4):952–956, December 1981.
- [22] R. Nelson and A. N. Tantawi. Approximate analysis of fork/join synchronization in parallel queues. *IEEE Transactions on Computers*, 37(6):739–743, 1988.
- [23] T. Osogami, M. Harchol-Balter, and A. Scheller-Wolf. Analysis of cycle stealing with switching times and thresholds. In *SIGMETRICS*, pages 184–195, June 2003.



- [24] N. B. Shah, K. Lee, and K. Ramchandran. The MDS queue: Analysing latency performance of codes and redundant requests. Technical Report arXiv:1211.5405, November 2012.
- [25] N. B. Shah, K. Lee, and K. Ramchandran. When do redundant requests reduce latency? Technical Report arXiv:1311.2851, June 2013.
- [26] A. L. Stolyar and T. Tezcan. Control of systems with flexible multi-server pools: a shadow routing approach. *Queueing Systems*, 66:1–51, 2010.
- [27] J. Tsitsiklis and K. Xu. On the power of (even a little) resource pooling. *Stochastic Systems*, 2:1–66, 2012.
- [28] J. Tsitsiklis and K. Xu. Queueing system topologies with limited flexibility. In *SIGMETRICS*, 2013.
- [29] J. Visschers, I. Adan, and G. Weiss. A product form solution to a system with multi-type jobs and multi-type servers. *Queueing Systems*, 70:269–298.
- [30] A. Vulimiri, P. B. Godfrey, R. Mittal, J. Sherry, S. Ratnasamy, and S. Shenker. Low latency via redundancy. Technical Report arXiv:1306.3707, November 2013.
- [31] D. Wang, G. Joshi, and G. Wornell. Efficient task replication for fast response times in parallel computation. Technical Report arXiv:1404.1328, April 2014.
- [32] C. Xia, Z. Liu, D. Towsley, and M. Lelarge. Scalability of fork/join queueing networks with blocking. In *SIGMETRICS*, pages 133–144, June 2007.

## A Proof of Theorem 1

**Theorem 1.** *The limiting probability of being in state  $(c_n, c_{n-1}, \dots, c_1)$  is*

$$\pi_{(c_n, \dots, c_1)} = C \prod_{i=1}^n \frac{\lambda_{c_i}}{\sum_{\substack{m \in \bigcup \\ j \leq i} S_{c_j}} \mu_m},$$

where  $C$  is a normalizing constant.

*Proof.* We begin by writing local balance equations for our states. The local balance equations are:

$$\begin{aligned} A &\equiv \text{Rate entering state } (c_n, \dots, c_1) &= \text{Rate leaving state } (c_n, \dots, c_1) &\equiv A' \\ &\text{due to an arrival} && \text{due to a departure} \\ \\ B_c &\equiv \text{Rate entering state } (c_n, \dots, c_1) &= \text{Rate leaving state } (c_n, \dots, c_1) &\equiv B'_c \\ &\text{due to a departure of class } c && \text{due to an arrival of class } c \end{aligned}$$

For an empty system, the state is  $()$ . It is not possible to enter state  $()$  due to an arrival or to leave due to a departure, so we only have one local balance equation of the form  $B_c = B'_c$ :

$$\pi_{()} \lambda_c = \pi_{(c)} \sum_{m \in S_c} \mu_m. \quad (13)$$

For any other state  $(c_n, c_{n-1}, \dots, c_1)$ , we have local balance equations of the form:

$$A = \pi_{(c_{n-1}, \dots, c_1)} \lambda_{c_n} = \pi_{(c_n, \dots, c_1)} \sum_{m \in \bigcup_{j \leq n} S_{c_j}} \mu_m = A' \quad (14)$$

$$B_c = \sum_{i=0}^n \sum_{\substack{m \in S_c, \\ m \notin S_{c_j}, \\ 1 \leq j \leq i}} \pi_{(c_n, \dots, c_{i+1}, c, c_i, \dots, c_1)} \mu_m = \pi_{(c_n, \dots, c_1)} \lambda_c = B'_c. \quad (15)$$

We guess the following form for  $\pi_{(c_n, \dots, c_1)}$ :

$$\pi_{(c_n, \dots, c_1)} = C \prod_{i=1}^n \frac{\lambda_{c_i}}{\sum_{\substack{m \in \bigcup \\ j \leq i} S_{c_j}} \mu_m}. \quad (16)$$

We will prove inductively that our guess satisfies the balance equations. The base case is state  $()$ . Substituting the guess from (16) into the left-hand side of (13), we get:

$$\pi_{(c)} \sum_{m \in S_c} \mu_m = C \frac{\lambda_c}{\sum_{m \in S_c} \mu_m} \sum_{m \in S_c} \mu_m = C \lambda_c = \pi_{()} \lambda_c,$$

which is exactly the right-hand side of (13).

Now, assume that (14) and (15) hold for some  $n-1 \geq 0$ . We will show that both hold for  $n$ .

1.  $A = A'$ . From (14), we have:

$$A = \pi_{(c_{n-1}, \dots, c_1)} \lambda_{c_n} = \prod_{i=1}^{n-1} \frac{\lambda_{c_i}}{\sum_{\substack{m \in \bigcup \\ j \leq i} S_{c_j}} \mu_m} \lambda_{c_n}$$

$$\begin{aligned}
&= \pi_{c_n, \dots, c_1} \frac{\sum_{m \in \bigcup_{j \leq n} S_{c_j}} \mu_m}{\lambda_{c_n}} \lambda_{c_n} \\
&= \pi_{c_n, \dots, c_1} \sum_{m \in \bigcup_{j \leq n} S_{c_j}} \mu_m = A'.
\end{aligned}$$

2.  $\mathbf{B}_c = \mathbf{B}'_c$ . From (15), we have:

$$\begin{aligned}
B_c &= \sum_{i=0}^n \sum_{\substack{m \in S_c, \\ m \notin S_{c_j}, \\ 1 \leq j \leq i}} \pi_{(c_n, \dots, c_{i+1}, c, c_i, \dots, c_1)} \mu_m \\
&= \sum_{i=0}^{n-1} \sum_{m \in S_c \setminus \bigcup_{j \leq i} S_{c_j}} \pi_{(c_n, \dots, c_{i+1}, c, c_i, \dots, c_1)} \mu_m + \sum_{m \in S_c \setminus \bigcup_{j \leq n} S_{c_j}} \pi_{(c, c_n, \dots, c_1)} \mu_m \\
&= \sum_{i=0}^{n-1} \sum_{m \in S_c \setminus \bigcup_{j \leq i} S_{c_j}} \frac{\lambda_{c_n} \pi_{(c_{n-1}, \dots, c_{i+1}, c, c_i, \dots, c_1)}}{\sum_{t \in \bigcup_{j \leq n} S_{c_j} \cup S_c} \mu_t} \mu_m + \frac{\lambda_c \sum_{m \in S_c \setminus \bigcup_{j \leq n} S_{c_j}} \mu_m}{\sum_{m \in \bigcup_{j \leq n} S_{c_j} \cup S_c} \mu_m} \prod_{i=1}^n \frac{\lambda_{c_i}}{\sum_{m \in \bigcup_{j \leq i} S_{c_j}} \mu_m} \\
&= \frac{\lambda_{c_n}}{\sum_{m \in \bigcup_{j \leq n} S_{c_j} \cup S_c} \mu_m} \sum_{i=0}^{n-1} \sum_{m \in S_c \setminus \bigcup_{j \leq i} S_{c_j}} \pi_{(c_{n-1}, \dots, c_{i+1}, c, c_i, \dots, c_1)} \mu_m + \lambda_c \pi_{(c_n, \dots, c_1)} \frac{\sum_{m \in S_c \setminus \bigcup_{j \leq n} S_{c_j}} \mu_m}{\sum_{m \in \bigcup_{j \leq n} S_{c_j} \cup S_c} \mu_m} \\
&= \frac{\lambda_{c_n}}{\sum_{m \in \bigcup_{j \leq n} S_{c_j} \cup S_c} \mu_m} \pi_{(c_{n-1}, \dots, c_1)} \lambda_c + \lambda_c \pi_{(c_n, \dots, c_1)} \frac{\sum_{m \in S_c \setminus \bigcup_{j \leq n} S_{c_j}} \mu_m}{\sum_{m \in \bigcup_{j \leq n} S_{c_j} \cup S_c} \mu_m} \\
&= \lambda_c \frac{\lambda_{c_n}}{\sum_{m \in \bigcup_{j \leq n} S_{c_j} \cup S_c} \mu_m} \prod_{i=1}^{n-1} \frac{\lambda_{c_i}}{\sum_{m \in \bigcup_{j \leq i} S_{c_j}} \mu_m} + \lambda_c \pi_{(c_n, \dots, c_1)} \frac{\sum_{m \in S_c \setminus \bigcup_{j \leq n} S_{c_j}} \mu_m}{\sum_{m \in \bigcup_{j \leq n} S_{c_j} \cup S_c} \mu_m} \\
&= \lambda_c \pi_{(c_n, \dots, c_1)} \frac{\sum_{m \in \bigcup_{j \leq n} S_{c_j}} \mu_m}{\sum_{m \in \bigcup_{j \leq n} S_{c_j} \cup S_c} \mu_m} + \lambda_c \pi_{(c_n, \dots, c_1)} \frac{\sum_{m \in S_c \setminus \bigcup_{j \leq n} S_{c_j}} \mu_m}{\sum_{m \in \bigcup_{j \leq n} S_{c_j} \cup S_c} \mu_m} \\
&= \lambda_c \pi_{(c_n, \dots, c_1)} = B'_c.
\end{aligned}$$

Hence the local balance equations hold for all  $n$ , and so the guess for the limiting probabilities from (16) is correct.  $\square$

## B Proof of Theorem 10

**Theorem 10.** 1. In the scaled  $\mathbb{N}$  model, the distribution of the number of class  $R$  jobs in the system is

$$N_R \sim \text{Geometric} \left( 1 - \frac{\lambda_R}{\sum_{i=1}^k \mu_i - \sum_{i=2}^k \lambda_i} \right) - 1,$$

and the distribution of the response time of class  $R$  jobs is

$$T_R \sim \text{Exp} \left( \sum_{i=1}^k \mu_i - \sum_{i=2}^k \lambda_i - \lambda_R \right).$$

2. In the scaled  $\mathbb{W}$  model, the distribution of the number of class  $R$  jobs in the system is

$$N_R \sim \text{Geometric} \left( 1 - \frac{\lambda_R}{\sum_{i=1}^k \mu_i - \sum_{i=1}^k \lambda_i} \right) - 1,$$

and the distribution of the response time of class  $R$  jobs is

$$T_R \sim \text{Exp} \left( \sum_{i=1}^k \mu_i - \sum_{i=1}^k \lambda_i - \lambda_R \right).$$

*Proof.* To find  $\Pr\{N_R = n_R\}$ , we will consider the non- $R$  jobs in the queue as being split into two pieces: the non- $R$  jobs before the first  $R$  in the queue, and the non- $R$  jobs after the first  $R$  in the queue. We sum over all possible lengths of these two pieces, and all possible classes of these non- $R$  jobs. Let  $x_0$  be the number of non- $R$  jobs before the first  $R$  in the queue, and let  $x_1$  be the number of non- $R$  jobs after the first  $R$  in the queue. Then we have:

$$\begin{aligned} \Pr\{N_R = n_R\} &= \sum_{x_0=0}^{\infty} \sum_{x_1=0}^{\infty} C \eta_R^{n_R} X_0 \binom{x_1 + n_R - 1}{x_1} \prod_{\substack{j \geq x_0 \\ c_j \neq R}} X_1 \\ &= C \eta_R^{n_R} \left( \sum_{x_0=0}^{\infty} X_0 \right) \left( \sum_{x_1=0}^{\infty} X_1^{x_1} \binom{x_1 + n_R - 1}{x_1} \right), \end{aligned}$$

where

$$\begin{aligned} \eta_R &= \frac{\lambda_R}{\sum_{m=1}^k \mu_m}, \\ X_0 &= \prod_{j=1}^{x_0} \frac{\sum_{c_i \neq R} \lambda_{c_i}}{\sum_{\substack{m \in \cup \\ t \leq j}} S_t \mu_m}, \\ X_1 &= \frac{\sum_{c_i \neq R} \lambda_{c_i}}{\sum_{m=1}^k \mu_m}. \end{aligned}$$

The sums in the numerators of  $X_0$  and  $X_1$  take into account all of the possible combinations of classes making up the  $x_0$  and  $x_1$  jobs, respectively.

Now let  $C_1 = \sum_{x_0=0}^{\infty} X_0$  (note that this is a constant with respect to  $n_R$ ). Using the identity  $\sum_{i=0}^{\infty} p^i \binom{i+n-1}{i} = \left( \frac{1}{1-p} \right)^n$  for  $|p| < 1$ , we have:

$$\Pr\{N_R = n_R\} = C C_1 \eta_R^{n_R} \left( \frac{1}{1 - \frac{\sum_{c_i \neq R} \lambda_{c_i}}{\sum_m \mu_m}} \right)^{n_R}$$

$$= CC_1 \left( \frac{\lambda_R}{\sum_m \mu_m - \sum_{c_i \neq R} \lambda_{c_i}} \right)^{n_R}.$$

Using the normalization equation

$$\sum_{n_R=0}^{\infty} \Pr\{N_R = n_R\} = 1,$$

we find

$$CC_1 = 1 - \frac{\lambda_R}{\sum_m \mu_m - \sum_{c_i \neq R} \lambda_{c_i}}.$$

Hence  $N_R \sim \text{Geometric}(1 - \frac{\lambda_R}{\sum_m \mu_m - \sum_{c_i \neq R} \lambda_{c_i}}) - 1$ .

Next, we obtain the Laplace transform of the response time for class  $R$  jobs,  $\tilde{T}_R(s)$ , via distributional Little's Law. First, we consider the  $z$ -transform of the number of class- $R$  Poisson arrivals during  $T$ ,  $\hat{A}_{T_R}(z) = \tilde{T}_R(\lambda_R - \lambda_R z)$ . Class  $R$  jobs depart the system in the same order in which they arrive, so  $A_{T_R}$  is equivalent to  $N_R$ , the number of jobs seen by an  $R$  departure. Hence

$$\tilde{T}_R(\lambda_R - \lambda_R z) = \hat{N}_R(z).$$

We know that  $N_R$  is distributed  $\text{Geometric}(p) - 1$ , where  $p = 1 - \frac{\lambda_R}{\sum_m \mu_m - \sum_{c_i \neq R} \lambda_{c_i}}$ . Hence we have

$$\tilde{T}_R(\lambda_R - \lambda_R z) = \hat{N}_R(z) = \frac{p}{1 - z(1 - p)}.$$

Let  $s = \lambda_R - \lambda_R z$ , so that  $z = 1 - s/\lambda_R$ . Then we have

$$\tilde{T}_R(s) = \frac{p}{1 - (1 - \frac{s}{\lambda_R})(1 - p)},$$

which after some simplification gives

$$\tilde{T}_R(s) = \frac{\sum_m \mu_m - \sum_{c_i \neq R} \lambda_{c_i} - \lambda_R}{\sum_m \mu_m - \sum_{c_i \neq R} \lambda_{c_i} - \lambda_R + s}.$$

Hence  $T_R \sim \text{Exp}(\sum_m \mu_m - \sum_{c_i \neq R} \lambda_{c_i} - \lambda_R)$ .

The derivation for the  $\mathbb{W}$  model is very similar, and is omitted here. □