

# **Exact Analysis of the M/M/k/setup Class of Markov Chains via Recursive Renewal Reward**

**Anshul Gandhi      Sherwin Doroudi\***  
**Mor Harchol-Balter      Alan Scheller-Wolf\***

April 2013  
CMU-CS-13-105

School of Computer Science  
Carnegie Mellon University  
Pittsburgh, PA 15213

\*Tepper School of Business, Carnegie Mellon University

This research was sponsored by the National Science Foundation under grant number NSF-CSR-116282 and Intel Science and Technology Center on Cloud Computing. The views and conclusions contained in this document are those of the author and should not be interpreted as representing the official policies, either expressed or implied, of any sponsoring institution, the U.S. government or any other entity.

**Keywords:** Queuing Theory; Setup Times; Renewal Reward

## Abstract

The M/M/k/setup model, where there is a penalty for turning servers on, is common in data centers, call centers and manufacturing systems. Setup costs take the form of a time delay, and sometimes there is additionally a power penalty, as in the case of data centers. While the M/M/1/setup was exactly analyzed in 1964, no exact analysis exists to date for the M/M/k/setup with  $k > 1$ .

In this paper we provide the first exact, closed-form analysis for the M/M/k/setup and some of its important variants including systems in which: (i) idle servers delay for a period of time before turning off, or (ii) idle servers can either be turned off or put to sleep. Our analysis is made possible by our development of a new technique, Recursive Renewal Reward (RRR), for solving Markov chains with a repeating structure. RRR uses ideas from renewal reward theory and busy period analysis to obtain closed-form expressions for metrics of interest such as the transform of time in system and the transform of power consumed by the system. The simplicity, intuitiveness, and versatility of RRR makes it useful for analyzing Markov chains far beyond the M/M/k/setup. In general, RRR can be used to analyze any 2-dimensional Markov chain which is finite in one dimension and possibly infinite (with a repeating structure) in the other dimension.



# 1 Introduction

Setup times (a.k.a. exceptional first service) are a fundamental component of computer systems and manufacturing systems, and therefore they have always played an important role in queueing theoretic analysis. In manufacturing systems it is very common for a job that finds a server idle to wait for the server to “warm up” before the server is initiated. In retail and hospitals, the arrival of customers may necessitate bringing in an additional human server, which requires a setup time for the server to arrive. In computer systems, setup times are once again at the forefront of research, as they are the key issue in dynamic capacity provisioning for data centers.

In data centers, it is desirable to turn idle servers off, or reallocate the servers, to save power. This is because idle servers burn power at 60–70% of the peak rate, so leaving servers on and idle is wasteful [4]. Unfortunately, most companies are hesitant to turn off idle servers because the setup time needed to restart these servers is very costly; the typical setup times for servers is 200 seconds, while a job’s service requirement is typically less than 1 second [15, 6]. Not only is the setup time prohibitive, but power is also burned at peak rate during the entire setup period, although the server is still not functional. Thus it is not at all obvious that turning off idle servers is advantageous.

Many ideas have been proposed to minimize the number of times that servers in a data center must undergo setup. One major line of research involves load prediction techniques [15, 20, 5, 11]. In the case where load is unpredictable, research has turned to looking at policies such as delayedoff, which delay turning off an idle server for some fixed amount of time, in anticipation of a new arrival [13, 10, 8]. Another line of research involves reducing setup times by developing low power sleep modes [10, 18].

Surprisingly, for all the importance of setup times, very little is known about their analysis. The  $M/G/1$  with setup times was analyzed in 1964 by Welch [24]. The analysis of an  $M/M/k$  system with setup times, which we refer to as  **$M/M/k/setup$** , however, has remained elusive, owing largely to the complexity of the underlying Markov chain. (Fig. 1 shows an  $M/M/k/setup$  with exponentially distributed setup times.) In 2010, various analytical approximations for the  $M/M/k/setup$  were proposed in [9]. These approximations work well provided that either load is low or the setup time is low. The  $M/M/\infty/setup$  was also analyzed in [9] and found to exhibit product form. Other than the above, no progress has been made on the  $M/M/k/setup$ . Even less is known about the  **$M/M/k/setup/delayedoff$** , where idle servers delay for a finite amount of time before turning off, or the  **$M/M/k/setup/sleep$** , where idle servers can either be turned off (high setup time, zero power) or put to sleep (lower setup time, low power). Section 3 describes these models in greater detail. Section 2 describes related prior work, including existing methods for solving general Markov chains with a repeating structure.

This paper is the first to derive an exact, closed-form solution for the  $M/M/k/setup$ , the  $M/M/k/setup/delayedoff$ , and the  $M/M/k/setup/sleep$ . We obtain the Laplace transform of response time, the  $z$ -transform of power consumption, and other important metrics for all of the above models.

Our solution is made possible by our development of a new technique for solving Markov

chains with a repeating structure – one of the key contributions of our work. The technique is based on using renewal reward theory to obtain the metrics of interest, while utilizing certain recursion theorems about the chain. We call this technique Recursive Renewal Reward (RRR). Unlike matrix-analytic methods [16], RRR does not require finding the “rate” matrix. Another feature of RRR is that it is simple enough to be taught in an elementary stochastic processes course.

In general, RRR should be able to reduce the analysis of any 2-dimensional Markov chain which is finite in one dimension and possibly infinite (with repeating structure) in the other to the problem of solving a system of polynomial equations. Further, if all transitions in the repeating portion of the Markov chain are skip-free and all up/down arrows are unidirectional, the resulting system of equations will be at most quadratic, yielding a closed-form solution (see Section 11 and Fig. 6 for more details). We thus anticipate that RRR will prove useful to other researchers in analyzing many new problems.

## 2 Prior work

The few papers that have looked at the M/M/k/setup are discussed in Section 1. For the M/M/k/setup/delayedoff, only iterative matrix-analytic approaches have been used [8]. No analysis exists for M/M/k/setup/sleep. We now discuss papers that have considered repeating Markov chains and have proposed techniques for solving these. We then comment on how these techniques might or might not apply to the M/M/k/setup.

### 2.1 Matrix-analytic based approaches

Matrix-analytic methods are a common approach for analyzing Markov chains with repeating structure. Such approaches are typically numerical, generally involving iteration to find the rate matrix,  $R$ . These approaches do not, in general, lead to closed forms or to any intuition, but are very useful for evaluating chains under different parameters.

There are cases where it is known that the  $R$  matrix can be stated explicitly, as described in [16]. This typically involves using a combinatorial interpretation for the  $R$  matrix. As described in [16], the class of chains for which the combinatorial view is tractable is narrow. However, in [23], the authors show that the combinatorial interpretation extends to a broader class of chains. Their class does not include the M/M/k/setup, however, which is more complicated because the transition (setup) rates are not independent of the number of jobs in system. Much research has been done on improving matrix-analytic methods to make the iteration faster. An example is [22], which develops a fast iterative procedure for finding the rate matrix for a broader class of chains than that in [23]. The authors in [22] also provide an explicit solution for the rate matrix in terms of infinite sums.

## 2.2 Generating function based approaches

Generating functions have also been applied to solve chains with a repeating structure. Like matrix-analytic methods these are not intuitive: Generating function approaches involve guessing the form of the solution and then solving for the coefficients of the guess, often leading to long computations. In theory, they can be used to solve very general chains (see for example [1]). We initially tried applying a generating function approach to the M/M/2/setup and found it to be incredibly complex and without intuition. This led us to seek a simpler and more intuitive approach.

## 2.3 M/M/k with vacations

Many papers have been written about the M/M/k system with vacations, see for example [26, 25, 21, 17]. While the Markov chain for the M/M/k with vacations looks similar to the M/M/k/setup, the dynamics of the two systems are very different. A server takes a vacation as soon as it is idle and there are no jobs in the queue. By contrast, a setup time is initiated by jobs arriving to the queue. In almost all of the papers involving vacations, the vacation model is severely restricted, allowing only a fixed group of servers to go on vacation at once. This is very different from our system in which any number of servers may be in setup at any time. The model in [17] comes closest to our model, although the authors use generating functions and assume that *all* idle servers are on vacation, rather than one server being in setup for each job in queue, which makes the transitions in their chain independent of the number of jobs.

## 2.4 Restricted models of M/M/k with setup

There have been a few papers [2, 3, 9] that consider a very restricted version of the M/M/k/setup, wherein at most one server can be in setup at a time. There has also been prior work [19] that considers an M/M/k system wherein a fixed subset of servers can be turned on and off based on load. The underlying Markov chains for all of these restricted systems are analytically tractable and lead to very simple closed-form expressions, since the rate at which servers turn on is always fixed. Our M/M/k/setup system is more general, allowing any number of servers to be in setup. This makes our problem much more challenging.

## 2.5 How our work differs from all of the above

To the best of our knowledge, we are the first to derive exact closed-form results for the M/M/k/setup problem, with  $k > 1$ . Our solution was made possible by our new RRR technique. The RRR technique results in exact solutions, does not require any iteration, and does not involve infinite sums. Importantly, the RRR technique is highly intuitive and very easy to apply. Using the RRR technique, we go much further than the M/M/k setup, deriving exact closed-form results for important




Figure 1: Markov chain for the M/M/k/setup system. Each state is denoted by the pair  $(i, j)$ , where  $i$  is the number of on servers, and  $j$  is the number of jobs in the system. The number of servers in setup is  $\min\{j - i, k - i\}$ .

variants, such as the M/M/k/setup/delayedoff and the M/M/k with multiple types of setups, neither of which has been solved analytically.

### 3 Model

In our model jobs arrive according to a Poisson process with rate  $\lambda$  and are served at rate  $\mu = \frac{1}{E[S]}$ , where  $S$  denotes the job size and is exponentially distributed. For stability, we assume that  $k \cdot \mu > \lambda$ , where  $k$  is the number of servers in the system.

#### 3.1 M/M/k/setup

In the M/M/k/setup system, each of the  $k$  servers is in one of three states: **off**, **on** (being used to serve a job), or **setup**. When a server is on or in setup, it consumes peak power of  $P_{peak}$  watts. When a server is off, it consumes zero power. Thus, when servers are not in use, they are immediately turned off to save power. Every arriving job that comes into the system picks an off server, if one exists, and puts it into setup mode; the job then joins the queue. We use  $I$  to denote the setup times, with  $E[I] = \frac{1}{\alpha}$ . Unless stated otherwise, we assume that setup times are exponentially distributed. When a job completes service at a server, say server  $s_1$ , and there are no remaining jobs left in the queue, then server  $s_1$  is immediately turned off. However, if the queue is not empty, then server  $s_1$  is not turned off, and the job at the head of the queue is directed to server  $s_1$ . Note that if the job at the head of the queue was already waiting on another server, say



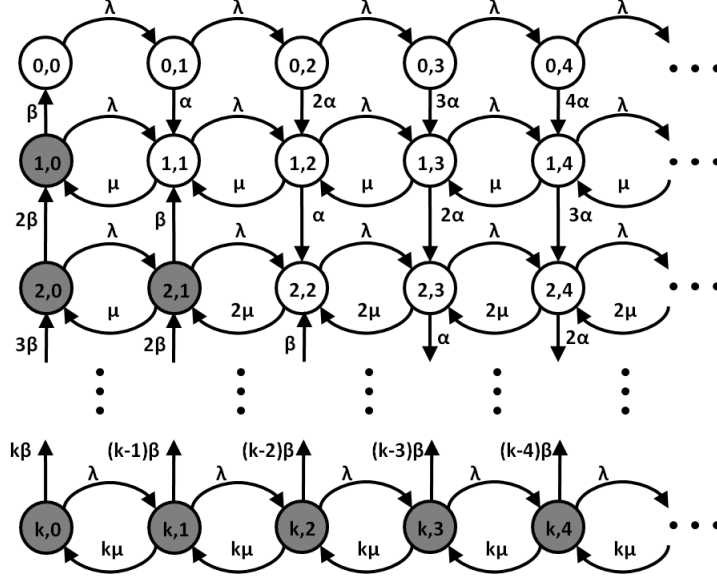


Figure 2: Markov chain for the  $M/M/k/setup/delayedoff$  system. Each state is denoted by the pair  $(i, j)$ , where  $i$  is the number of on or idle servers, and  $j$  is the number of jobs in the system. If  $i < j$ , then the number of servers in setup is  $\min\{j - i, k - i\}$ , and there are no idle servers. If  $i > j$  (gray shaded states), the number of idle servers is  $(i - j)$ , and there are no servers in setup. If  $i = j$ , no servers are idle or in setup.

server  $s_2$ , in setup mode, the job at the head of the queue is still directed to server  $s_1$ . At this point, if there is a job in the queue that did not setup an off server on arrival (because there were no off servers), then server  $s_2$  continues to be in setup for this job. If no such job exists in the queue, then server  $s_2$  is turned off.

The Markov chain for the  $M/M/k/setup$  system is shown in Fig. 1. Each state is denoted by the pair  $(i, j)$ , where  $i$  is the number of on servers, and  $j$  is the number of jobs in the system. Thus, the number of servers in setup is  $\min\{j - i, k - i\}$ . Note that the Markov chain is infinite in one dimension.

### 3.2 M/M/k/setup/delayedoff

The  $M/M/k/setup/delayedoff$  system is the same as the  $M/M/k/setup$  system, except that idle servers are not immediately turned off. Specifically, when a job completes service at a server, say server  $s_1$ , and there are no remaining jobs in the queue,  $s_1$  remains waiting in the idle state for an exponentially distributed amount of time with mean  $t_{wait} = \frac{1}{\beta}$ . If a new job arrives while server  $s_1$  is waiting, the job is immediately directed to  $s_1$ , which is already on. However, if no jobs arrive during server  $s_1$ 's waiting period, then server  $s_1$  is turned off. Intuitively, a higher  $t_{wait}$  results in lower response time, since servers are on longer, but may also increase power usage, since idle servers consume significant power.




Figure 3: *Markov chain for the M/M/k/setup/sleep system. Each state is denoted by the pair  $(i, j)$ , where  $i$  is the number of on servers, and  $j$  is the number of jobs in the system. The number of servers in fast setup is  $s$ , which in this case is  $s = 2$ . The number of servers in setup is  $\min\{j - i, k - i\}$ . If  $i < s$  (gray shaded states), the first  $(s - i)$  servers setting up have a fast setup rate,  $\omega$ , while the other servers in setup have a slow setup rate,  $\alpha$ .*

The Markov chain for the M/M/k/setup/delayedoff system is shown in Fig. 2. The chain is the same as that for M/M/k/setup, except for the new gray shaded states which represent states with idle servers. As before, each state is denoted by the pair  $(i, j)$ , where  $i$  is the number of on or idle servers, and  $j$  is the number of jobs in the system. For the M/M/k/setup/delayedoff system, each server can be in one of four states: off, on (busy), idle, or setup. If  $i < j$ , then the number of servers in setup is  $\min\{j - i, k - i\}$ , and there are no idle servers. If  $i > j$  (gray shaded states), the number of idle servers is  $(i - j)$ , and there are no servers in setup. If  $i = j$ , no servers are idle or in setup.

### 3.3 M/M/k/setup/sleep

The M/M/k/setup/sleep is motivated by servers with sleep modes [10, 18], which allow an idle server to either be turned off or put to sleep. When a server is turned off, it consumes zero power. However, turning on an off server requires an exponentially distributed setup time, with rate  $\alpha$ . By contrast, when a server is sleeping, it consumes some non-zero power,  $P_{sleep}$  watts, which is usually [10, 18] much smaller than the idle power,  $P_{idle}$  watts. When a sleeping server is turned on, it requires an exponentially distributed setup time, with rate  $\omega > \alpha$ . Thus, there is a trade-off between turning off an idle server vs putting it to sleep.

The Markov chain for the M/M/k/setup/sleep is the same as that for the M/M/k/setup, except

that in the case of M/M/k/setup/sleep servers that are setting up can either have a fast setup rate of  $\omega$  or a slow setup rate of  $\alpha$ . In our setting  $s \leq k$  servers have a fast setup rate and  $(k - s)$  servers have a slow setup rate, where  $s$  is a decision variable. In order to leverage the benefits of fast setup times, given all servers are initially off, the first  $s$  servers setting up have a setup rate of  $\omega$  and the rest of the servers in setup have a setup rate of  $\alpha$ .

For ease of analysis, we make the following assumptions about the M/M/k/setup/sleep model: (i) If we have a group of servers in setup, with some servers in fast setup, we assume that the first server to complete setting up is one that had a fast setup rate ( $\omega$ ). Thus, if we are in state  $(i, j)$  with  $i < s$  (gray shaded states in Fig. 3), the first  $(s - i)$  servers in setup will have a fast setup rate. Note that the  $i$  servers already on in state  $(i, j)$ , with  $i \leq s$ , are those that had a fast setup rate. Thus, when we have  $i \leq s$  servers busy, and a server is no longer in use, we put the server to sleep (as opposed to turning it off). (ii) If we have  $i > s$  servers busy, and a server is no longer in use, we turn the server off (as opposed to putting it to sleep). This assumption allows us to save a lot of power when load goes down since off servers consume zero power. The above two assumptions are primarily for tractability of the M/M/k/setup/sleep Markov chain. We simulated an M/M/k/setup/sleep system with and without the above two assumptions and found the results to be qualitatively unchanged.

## 4 The Recursive Renewal Reward technique

In this section we provide a high-level description of our new Recursive Renewal Reward (RRR) technique, which yields exact, closed-form solutions for a range of Markov chains, including the M/M/k/setup (see Sections 5, 6 and 7), the M/M/k/setup/delayedoff (see Section 9) and the M/M/k/setup/sleep (see Section 10).

The RRR technique works by deriving the *expected “reward” earned per unit time in a Markov chain*, where the reward could be any quantity of interest. In the context of our M/M/k/setup problem, the reward earned at time  $t$ ,  $R(t)$ , could be the number of jobs in system at time  $t$ , the square of the number of jobs in system, the current power usage, the number of servers that are on, or any other reward that can be expressed as a function of the states of the Markov chain.

To analyze the average rate of earning reward, we designate a *renewal state*, say  $(0, 0)$ ,<sup>1</sup> which we call the *home state*, and then consider a *renewal cycle* to be the process of moving from the home state back to the home state. By renewal-reward theory, the overall rate of earning reward is the same as the mean reward earned over a renewal cycle, which we denote by  $\mathcal{R}$ , divided by the mean length of the renewal cycle, denoted by  $\mathcal{T}$ .

$$\text{Average rate of earning} = \frac{\mathcal{R}}{\mathcal{T}} = \frac{E \left[ \int_{\text{cycle}} R(t) dt \right]}{E \left[ \int_{\text{cycle}} 1 dt \right]}$$

---

<sup>1</sup>In principle any state can be chosen as the renewal state, but some states allow for an easier (or shorter) analysis.




Figure 4: Markov chain for the M/M/1/setup. The repeating portion is highlighted in gray and the border states are shaded black.

For example, if the goal is to find the mean number of jobs,  $E[N]$ , for our chain, we simply define  $R(t)$  to be the number of jobs at time  $t$ , which can be obtained from the state of the Markov chain at time  $t$ .

It turns out that the quantities,  $\mathcal{T}$  and  $\mathcal{R}$ , are very easy to compute! Consider a Markov chain, such as that in Fig. 4 or Fig. 5. The *repeating portion* of the chain is shown in gray. There are a finite number of *border states* which sit at the edge of the repeating chain and are colored black. We will see that computing  $\mathcal{T}$  and  $\mathcal{R}$  basically reduces to writing one equation for each border state. For the case of  $\mathcal{T}$ , we will need the mean time to move one step left from each border state. For the case of  $\mathcal{R}$ , we will need the mean reward earned when moving one step left from each border state. Computing these border state quantities is made very easy via some neat recursion theorems. We demonstrate this process in the examples below. There are a few details which we will defer until after these examples. For instance, in general, it is necessary to also add equations for the non-repeating portion of the Markov chain. See Sections 7 and 11 for more details on the RRR technique.

## 5 M/M/1/setup

In this section we illustrate the RRR technique by applying it to the simple M/M/1/setup system, whose Markov chain is shown in Fig. 4. Here, the state of the system is represented as  $(i, j)$ , where  $i \in \{0, 1\}$  is the number of servers on and  $0 \leq j < \infty$  is the number of jobs in the system. In general,  $i$  represents the *depth* (or row number) of the state, and  $j$  represents the *level* (or column number) of the state. We start by deriving  $E[N]$ , the mean number of jobs, and then move to more complex metrics. We choose the renewal state to be  $(0, 0)$  and we define the reward earned at time  $t$ ,  $R(t)$ , to be  $N(t)$ , the number of jobs in the system at time  $t$ . As explained in Section 4, all we need is  $\mathcal{T}$  and  $\mathcal{R}$ .

## 5.1 Deriving $\mathcal{T}$ via $T_{0,1}^L$ and $T_{1,1}^L$

$\mathcal{T}$  is the mean time to get from our home state  $(0, 0)$  back to  $(0, 0)$ . This can be viewed as  $\frac{1}{\lambda}$ , the mean time until we leave  $(0, 0)$  (which takes us to  $(0, 1)$ ) plus the mean time to get home from  $(0, 1)$ . We make the further observation that the mean time to get home from  $(0, 1)$  is equal to  $T_{0,1}^L$  (using notation from Table 1), the mean time to move left one level from  $(0, 1)$  (since moving left can only put us in  $(0, 0)$ ). We thus have:

$$\mathcal{T} = \frac{1}{\lambda} + T_{0,1}^L \quad (1)$$

We now need an equation for  $T_{0,1}^L$  for the border state  $(0, 1)$ , which will require looking at the other border state,  $(1, 1)$ , as well. Starting with border state  $(1, 1)$ , it is clear that  $T_{1,1}^L$  is simply the mean length of an M/M/1 busy period,  $B_1$ . Thus, we have:

$$T_{1,1}^L = B_1 = \frac{1}{\mu - \lambda} \quad (2)$$

$T_{0,1}^L$  involves waiting in state  $(0, 1)$  for expected time  $\frac{1}{\alpha + \lambda}$ , before conditioning on where we transition to next. If we go to state  $(1, 1)$  we need an additional  $T_{1,1}^L$ . However if we go to state  $(0, 2)$  we need to add on the time to move one step left from  $(0, 2)$  (which by Fig. 4 takes us to  $(1, 1)$ ) and then an additional  $T_{1,1}^L$ . That is:

$$T_{0,1}^L = \frac{1}{\lambda + \alpha} + \frac{\alpha}{\lambda + \alpha} \cdot T_{1,1}^L + \frac{\lambda}{\lambda + \alpha} (T_{0,2}^L + T_{1,1}^L) \quad (3)$$

It is now time to invoke one of our recursion theorems, which holds for any M/M/k/setup chain:

**Theorem 1** (Recursion theorem for mean time).

*For the M/M/k/setup, the mean time to move one step left from state  $(i, j)$ ,  $T_{i,j}^L$ , is the same for all  $j \geq k$ .*

Thm. 1 follows from the fact that the repeating portion of the Markov chain is identical for all states in a given row. The full proof of Thm. 1 (along with the proofs of all other theorems appearing in this paper) is presented in Appendix A.

Using Thm. 1, we can replace  $T_{0,2}^L$  in Eq. (3) with  $T_{0,1}^L$  to get:

$$T_{0,1}^L = \frac{1}{\lambda + \alpha} + \frac{\alpha}{\lambda + \alpha} \cdot T_{1,1}^L + \frac{\lambda}{\lambda + \alpha} (T_{0,1}^L + T_{1,1}^L) \quad (4)$$

Finally, noting that  $T_{1,1}^L = B_1$  from Eq. (2), we have that:

$$\begin{aligned} T_{0,1}^L &= \frac{1}{\lambda + \alpha} + \frac{\alpha}{\lambda + \alpha} \cdot B_1 + \frac{\lambda}{\lambda + \alpha} (T_{0,1}^L + B_1) \\ \implies T_{0j}^L = T_{0,1}^L &= \frac{1 + (\lambda + \alpha)B_1}{\alpha} \end{aligned} \quad (5)$$

This completes the derivation of  $\mathcal{T}$  via Eq. (1).

| Variable                | Description   |
|-------------------------|---|
| $\mathcal{T}$           | Mean length of the renewal cycle  |
| $\mathcal{R}$           | Mean reward earned during a renewal cycle   |
| $T_{i,j}^L$             | Mean time until we first move one level left of $(i, j)$ , starting from $(i, j)$             |
| $R_{i,j}^L$             | Mean reward earned until we first move one level left of $(i, j)$ , starting from $(i, j)$    |
| $p_{i \rightarrow d}^L$ | Probability that after we first move one level left from state $(i, j)$ , we are at depth $d$ |
| $B_k$                   | Mean length of an $M/M/k$ busy period   |

Table 1: Description of the variables used in our analysis of  $E[N]$ .

## 5.2 Deriving $\mathcal{R}$ via $R_{0,1}^L$ and $R_{1,1}^L$

$\mathcal{R}$  denotes the reward earned in moving from  $(0, 0)$  back to  $(0, 0)$ . Observing that we earn 0 reward in state  $(0, 0)$  (because there are no jobs in the system in that state), and observing that from state  $(0, 0)$  we can only next move to  $(0, 1)$ , we have (using notation from Table 1):

$$\mathcal{R} = R_{0,1}^L \quad (6)$$

It now remains to compute the reward earned in moving one step left from  $(0, 1)$ , which will require looking at the other border state,  $(1, 1)$ , as well.

To do this, we invoke another recursion theorem, which again holds for any  $M/M/k$ /setup system:

**Theorem 2** (Recursion theorem for mean reward).

For the  $M/M/k$ /setup, the mean reward earned in moving one step left from state  $(i, j + 1)$ ,  $R_{i,j+1}^L$ , satisfies  $R_{i,j+1}^L = R_{i,j}^L + T_{i,j}^L$  for all  $j \geq k$ , where the reward tracks the number of jobs in the system.

Applying Thm. 2 to the Markov chain in Fig. 4, we have:

$$\begin{aligned} R_{1,1}^L &= \frac{1}{\lambda + \mu} \cdot 1 + \frac{\mu}{\lambda + \mu} \cdot 0 + \frac{\lambda}{\lambda + \mu} (R_{1,2}^L + R_{1,1}^L) \\ &= \frac{1}{\lambda + \mu} + \frac{\lambda}{\lambda + \mu} ((R_{1,1}^L + T_{1,1}^L) + R_{1,1}^L) \\ &= \frac{1}{\lambda + \mu} + \frac{\lambda}{\lambda + \mu} ((R_{1,1}^L + B_1) + R_{1,1}^L) \quad (\text{from Eq. (2)}) \end{aligned} \quad (7)$$

$$\implies R_{1,1}^L = \frac{1 + \lambda B_1}{\mu - \lambda} \quad (8)$$

| Variable            | Description   |
|---------------------|---|
| $\dot{\mathcal{R}}$ | Mean reward earned (for z-transform) during a renewal cycle   |
| $\dot{\mathcal{E}}$ | Mean reward earned (for transform of power) during a renewal cycle  |
| $\dot{R}_{i,j}^L$   | Mean reward earned (for z-transform) until we first move one level left of $(i, j)$ , starting from $(i, j)$          |
| $\dot{E}_{i,j}^L$   | Mean reward earned (for z-transform of power) until we first move one level left of $(i, j)$ , starting from $(i, j)$ |

Table 2: *Description of the variables used in our transform analyses.*

Similarly, for border state  $(0, 1)$ , we have:

$$\begin{aligned}
R_{0,1}^L &= \frac{1}{\lambda + \alpha} \cdot 1 + \frac{\alpha}{\lambda + \alpha} \cdot R_{1,1}^L + \frac{\lambda}{\lambda + \alpha} (R_{0,2}^L + R_{1,1}^L) \\
&= \frac{1}{\lambda + \alpha} + \frac{\alpha}{\lambda + \alpha} \cdot R_{1,1}^L + \frac{\lambda}{\lambda + \alpha} ((R_{0,1}^L + T_{0,1}^L) + R_{1,1}^L) \quad (\text{from Thm. 2}) \\
\implies R_{0,1}^L &= \frac{1 + \lambda T_{0,1}^L + (\lambda + \alpha) R_{1,1}^L}{\alpha}. \tag{9}
\end{aligned}$$

This completes the derivation of  $\mathcal{R}$  via Eq. (6).

### 5.3 Deriving $E[N]$

Since  $E[N] = \frac{\mathcal{R}}{\mathcal{T}}$ , combining Eq. (1) and Eq. (6), we get:

$$E[N] = \frac{\mathcal{R}}{\mathcal{T}} = \frac{\lambda}{\alpha} + \frac{\lambda}{\mu - \lambda} \tag{10}$$

The second term in the Right Hand Side of Eq. (10) can be identified [14] as the mean number of jobs in an M/M/1 system (without setup). Thus, Eq. (10) is consistent with the known decomposition property for the M/M/1/setup system [24].

### 5.4 Deriving $\widehat{N}(z)$ and $\widetilde{T}(s)$

Deriving the z-transform of the number of jobs,  $\widehat{N}(z) = E[z^N]$ , is just as easy as deriving  $E[N]$ . The only difference is that our reward function is now  $R(t) = z^{N(t)}$ , where  $N(t)$  is again the number of jobs in the system at time  $t$ . Thus

$$\widehat{N}(z) = E[z^N] = \frac{\dot{\mathcal{R}}}{\mathcal{T}},$$

where  $\dot{\mathcal{R}} = E \left[ \int_{\text{cycle}} z^{N(t)} dt \right]$  and  $\mathcal{T}$  is the same as before.

We will again invoke a recursion theorem which applies to any M/M/k/setup (using notation from Table 2):

**Theorem 3** (Recursion theorem for transform of reward).

For the M/M/k/setup,  $\dot{R}_{i,j+1}^L = z \cdot \dot{R}_{i,j}^L$ , for all  $j \geq k$ , where  $\dot{R}$  tracks the  $z$ -transform of the number of jobs in the system.

Let us now express  $\dot{\mathcal{R}}$  by conditioning on the first step from  $(0, 0)$ :

$$\dot{\mathcal{R}} = \frac{1}{\lambda} + \dot{R}_{0,1}^L \quad (11)$$

We again need one equation per border state:

$$\begin{aligned} \dot{R}_{1,1}^L &= \frac{1}{\lambda + \mu} \cdot z + \frac{\lambda}{\lambda + \mu} \left( z \cdot \dot{R}_{1,1}^L + \dot{R}_{1,1}^L \right) \\ \dot{R}_{0,1}^L &= \frac{1}{\lambda + \alpha} \cdot z + \frac{\alpha}{\lambda + \alpha} \cdot \dot{R}_{1,1}^L + \frac{\lambda}{\lambda + \alpha} \left( z \cdot \dot{R}_{0,1}^L + \dot{R}_{1,1}^L \right) \end{aligned}$$

Solving the above system and substituting  $\dot{R}_{0,1}^L$  into Eq. (11) allows us to express  $\dot{\mathcal{R}}$  in closed form. This gives us  $\widehat{N}(z)$ , after some algebra, as follows:

$$\widehat{N}(z) = E[z^N] = \frac{\dot{\mathcal{R}}}{\mathcal{T}} = \frac{\alpha(\mu - \lambda)}{(\mu - \lambda z)(\alpha + \lambda - \lambda z)} \quad (12)$$

To get the Laplace transform of response time,  $\widetilde{T}(s)$ , we use the distributional Little's Law [12] (since M/M/1/setup is a First-In-First-Out system):

$$\widetilde{T}(s) = \widehat{N} \left( 1 - \frac{s}{\lambda} \right) = \frac{\alpha(\mu - \lambda)}{(s + \alpha)(\mu + s - \lambda)} \quad (13)$$

## 5.5 Deriving $\widehat{P}(z)$

We now derive  $\widehat{P}(z)$ , the  $z$ -transform of the power consumed for the M/M/1/setup. The server consumes zero power when it is off, but consumes peak power,  $P_{peak}$  watts, when it is on or in setup. This time, the reward is simply the transform of the energy consumed over the renewal cycle,  $\dot{\mathcal{E}} = E \left[ \int_{\text{cycle}} z^{P(t)} dt \right]$ , where  $P(t)$  is the power consumed at time  $t$ . We begin with the recursive theorem for  $\dot{E}_{i,j}^L$ , just like we had Thm. 3 for  $\dot{R}_{i,j}^L$ .

**Theorem 4** (Recursion theorem for transform of power).

For the M/M/k/setup,  $\dot{E}_{i,j+1}^L = \dot{E}_{i,j}^L = T_{i,j}^L \cdot z^{k \cdot P_{peak}}$ , for all  $j \geq k$ .



Thm. 4 gives us  $\dot{E}_{i,j}^L$  in closed form, in terms of  $T_{i,j}^L$ . Following the usual renewal-reward approach, we get:

$$\widehat{P}(z) = E[z^P] = \frac{\dot{\mathcal{E}}}{\mathcal{T}} = \frac{\alpha(\mu - \lambda) + \lambda(\mu + \alpha)z^{P_{peak}}}{\mu(\lambda + \alpha)} \quad (14)$$

## 6 M/M/2/setup

The M/M/2/setup chain shown in Fig. 5 is analyzed similarly to the M/M/1/setup, except that there are now three border states,  $(0, 2)$ ,  $(1, 2)$ , and  $(2, 2)$ . The only complication is that when moving one level left from a given state, the resulting row is non-deterministic. For example, when moving left from  $(1, 3)$  in Fig. 5, we may end up in row 1 at  $(1, 2)$  or row 2 at  $(2, 2)$ . We use  $p_{i \rightarrow d}^L$  to denote the probability that once we move one level left from  $(i, j)$ , we will be at depth  $d$ .<sup>2</sup> The following theorem proves that  $p_{i \rightarrow d}^L$  is independent of  $j$  for all states  $(i, j)$  in the repeating portion.

**Theorem 5** (Recursion theorem for probability).

*For the M/M/k/setup, for each  $0 \leq d \leq k$  and for each  $0 \leq i \leq k$ ,  $p_{i \rightarrow d}^L$  is the same for all  $j \geq k$ .*

Thus, it suffices to compute  $p_{i \rightarrow d}^L$  for the border states. These probabilities are used in Section 6.2.

### 6.1 Deriving $p_{i \rightarrow d}^L$

Solving for the  $p_{i \rightarrow d}^L$  is easiest “bottom-up” (starting from the greatest depth,  $i$ ). For  $i = 2$ , we have  $p_{2 \rightarrow 2}^L = 1$  for all  $j > 2$ , since we stay at depth 2 after moving left. For  $i = 1$  and  $i = 0$ , we follow the same approach of conditioning on the first step and using recursion theorems:

$$p_{1 \rightarrow 1}^L = \frac{\mu}{\lambda + \mu + \alpha} + \frac{\lambda}{\lambda + \mu + \alpha} (p_{1 \rightarrow 1}^L)^2 \quad (15)$$

$$p_{0 \rightarrow 1}^L = \frac{2\alpha}{\lambda + 2\alpha} (p_{1 \rightarrow 1}^L) + \frac{\lambda}{\lambda + 2\alpha} (p_{0 \rightarrow 1}^L) (p_{1 \rightarrow 1}^L) \quad (16)$$

Eqs. (15) and (16) can now be solved in closed form since they are of degree at most 2. Note that  $p_{1 \rightarrow 2}^L = 1 - p_{1 \rightarrow 1}^L$  and  $p_{0 \rightarrow 2}^L = 1 - p_{0 \rightarrow 1}^L$ .

### 6.2 Deriving $\widehat{N}(z)$ via $\dot{R}_{0,2}^L$ , $\dot{R}_{1,2}^L$ , and $\dot{R}_{2,2}^L$

To derive  $\widehat{N}(z) = E[z^N]$ , we again need to find  $\dot{\mathcal{R}}$  and  $\mathcal{T}$ , where  $\dot{\mathcal{R}} = E \left[ \int_{\text{cycle}} z^{N(t)} dt \right]$ , and  $\mathcal{T} = E \left[ \int_{\text{cycle}} 1 dt \right] = \dot{\mathcal{R}} \Big|_{z=1}$ . Using  $(1, 1)$  as our renewal state and the same arguments as in

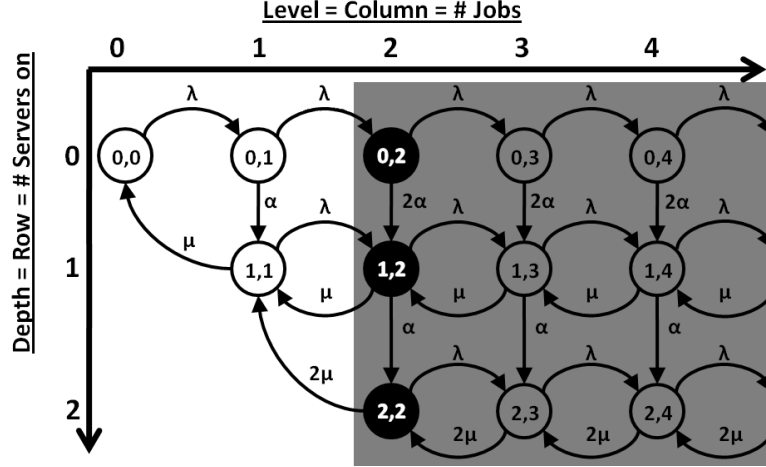


Figure 5: Markov chain for the M/M/2/setup. The repeating portion is highlighted in gray and the border states are shaded black.

Section 5.4, we have:

$$\dot{\mathcal{R}} = \frac{z}{\lambda + \mu} + \frac{\mu}{\lambda + \mu} \left( \frac{1}{\lambda} + \frac{z}{\lambda + \alpha} + \frac{\lambda}{\lambda + \alpha} \cdot \dot{R}_{0,2}^L \right) + \frac{\lambda}{\lambda + \mu} \cdot \dot{R}_{1,2}^L \quad (17)$$

It now remains to compute the reward equations for the border states:  $\dot{R}_{0,2}^L$ ,  $\dot{R}_{1,2}^L$ , and  $\dot{R}_{2,2}^L$ .

$$\dot{R}_{2,2}^L = \frac{z^2}{\lambda + 2\mu} + \frac{\lambda}{\lambda + 2\mu} \left( z \cdot \dot{R}_{2,2}^L + \dot{R}_{2,2}^L \right) \quad (18)$$

$$\dot{R}_{1,2}^L = \frac{z^2}{\lambda + \mu + \alpha} + \frac{\alpha}{\lambda + \mu + \alpha} \cdot \dot{R}_{2,2}^L + \frac{\lambda}{\lambda + \mu + \alpha} \left( z \cdot \dot{R}_{1,2}^L + (p_{1 \rightarrow 1}^L) \dot{R}_{1,2}^L + (1 - p_{1 \rightarrow 1}^L) \dot{R}_{2,2}^L \right) \quad (19)$$

$$\dot{R}_{0,2}^L = \frac{z^2}{\lambda + 2\alpha} + \frac{2\alpha}{\lambda + 2\alpha} \cdot \dot{R}_{1,2}^L + \frac{\lambda}{\lambda + 2\alpha} \left( z \cdot \dot{R}_{0,2}^L + (p_{0 \rightarrow 1}^L) \dot{R}_{1,2}^L + (1 - p_{0 \rightarrow 1}^L) \dot{R}_{2,2}^L \right) \quad (20)$$

Solving the above system of linear equations and substituting  $\dot{R}_{0,2}^L$  and  $\dot{R}_{1,2}^L$  into Eq. (17) allows us to solve for  $\hat{N}(z)$  in closed form as follows:

$$\begin{aligned} \hat{N}(z) = E[z^N] &= \frac{\dot{\mathcal{R}}}{\mathcal{T}} = \frac{\dot{\mathcal{R}}}{\dot{\mathcal{R}}|_{z=1}} \\ &= \frac{\lambda(\lambda + \alpha)(z + \lambda \dot{R}_{1,2}^L) + \mu(\alpha + \lambda(1 + z + \lambda \dot{R}_{0,2}^L))}{\lambda(\lambda + \alpha)(1 + \lambda T_{1j}^L) + \mu(\alpha + \lambda(2 + \lambda T_{0j}^L))} \end{aligned} \quad (21)$$

### 6.3 Deriving $\widetilde{T}(s)$

For the M/M/1/setup system, we were able to derive  $\widetilde{T}(s)$  directly from  $\widehat{N}(z)$  via the distributional Little's Law, since the M/M/1/setup is a FIFO system. Unfortunately, the M/M/2/setup system is not FIFO, since overtaking can occur. However, we can still apply the distributional Little's Law to the queue of the M/M/2/setup since the queue *is* FIFO. The analysis of  $\widehat{N}_Q(z)$  is very similar to that of  $\widehat{N}(z)$  and is thus omitted:

$$\widehat{N}_Q(z) = \frac{\lambda(\lambda + \alpha)(1 + \lambda\dot{R}_{1,2}^L) + \mu(\alpha + \lambda(1 + z + \lambda\dot{R}_{0,2}^L))}{\lambda(\lambda + \alpha)(1 + \lambda T_{1j}^L) + \mu(\alpha + \lambda(2 + \lambda T_{0j}^L))} \quad (22)$$

We now apply the distributional Little's Law to get  $\widetilde{T}_Q(s)$  from  $\widehat{N}_Q(z)$ . Finally, since  $T = T_Q + S$ , where  $S \sim \text{Exp}(\mu)$  is the job size distribution, we have:

$$\begin{aligned} \widetilde{T}(s) &= \widetilde{T}_Q(s) \cdot \frac{\mu}{s + \mu} = \widehat{N}_Q \left(1 - \frac{s}{\lambda}\right) \cdot \frac{\mu}{s + \mu} \\ &= \frac{\mu \left( \lambda(\lambda + \alpha)(1 + \lambda\dot{R}_{1,2}^L) + \mu(\alpha - s + \lambda(2 + \lambda\dot{R}_{0,2}^L)) \right)}{(s + \mu) \left( \lambda(\lambda + \alpha)(1 + \lambda T_{1j}^L) + \mu(\alpha + \lambda(2 + \lambda T_{0j}^L)) \right)} \end{aligned} \quad (23)$$

### 6.4 Deriving $\widehat{P}(z)$

The derivation of  $\widehat{P}(z)$  is similar to that of  $\widehat{N}(z)$  in Section 6.2, and is thus omitted.

$$\begin{aligned} \widehat{P}(z) &= \frac{\mu(\alpha + \lambda) + \lambda(\lambda + \mu + \alpha)z^{P_{peak}}}{\mu(\alpha + \lambda) + \lambda(\lambda + \mu + \alpha) + \lambda^2(\mu T_{0j}^L + (\lambda + \alpha)T_{1j}^L)} \\ &+ \frac{\lambda^2(\mu T_{0j}^L + (\lambda + \alpha)T_{1j}^L)z^{2P_{peak}}}{\mu(\alpha + \lambda) + \lambda(\lambda + \mu + \alpha) + \lambda^2(\mu T_{0j}^L + (\lambda + \alpha)T_{1j}^L)} \end{aligned} \quad (24)$$

## 7 M/M/k/setup

The M/M/k/setup chain shown in Fig. 1 is analyzed similarly to M/M/2/setup. The border states for M/M/k/setup are  $(i, k)$ , with  $0 \leq i \leq k$ . In the M/M/k/setup, the non-repeating portion consists of  $O(k^2)$  states. For  $k = 1$  and  $k = 2$ , we did not have to explicitly write reward equations for the non-repeating states; these were implicitly folded into other equations (see, for example, the term in parentheses in Eq. (17)). However, for arbitrarily large  $k$ , it is necessary to write reward equations for the states in the non-repeating portion. We use  $R_{i,j}^H$  to denote the reward earned until we reach the home state, starting from state  $(i, j)$  in the non-repeating portion. The  $R_{i,j}^H$  equations will be discussed in Section 7.3.

We illustrate the RRR technique for M/M/k/setup by deriving  $\widehat{N}_Q(z)$ , from which we can obtain  $\widetilde{T}(s)$ . One might think that analyzing the M/M/k/setup will require solving a  $k^{\text{th}}$  degree equation. This turns out to be false. Analyzing the M/M/k/setup via RRR only requires solving equations which are, at worst, quadratic.

We choose  $(k-1, k-1)$  to be the renewal state. Using RRR,  $\dot{\mathcal{R}}$  can be expressed as:

$$\dot{\mathcal{R}} = \frac{1 + (k-1)\mu\dot{R}_{k-2,k-2}^H + \lambda\dot{R}_{k-1,k}^L}{\lambda + (k-1)\mu} \quad (25)$$

We now derive the necessary  $\mathbf{p}_{i \rightarrow d}^L$ ,  $\dot{R}_{i,k}^L$ , and  $\dot{R}_{i,j}^H$  for computing  $\dot{\mathcal{R}}$ .

## 7.1 System of equations for $\mathbf{p}_{i \rightarrow d}^L$

The system of equations for  $p_{i \rightarrow d}^L$  is as follows:<sup>2</sup>

$$p_{i \rightarrow i}^L = \frac{\lambda(p_{i \rightarrow i}^L)^2 + i\mu}{\lambda + i\mu + (k-i)\alpha}, \quad (i < k) \quad (26)$$

$$p_{i \rightarrow d}^L = \frac{\lambda \left( \sum_{\ell=i}^d \{ (p_{i \rightarrow \ell}^L)(p_{\ell \rightarrow d}^L) \} \right) + (k-i)\alpha(p_{i+1 \rightarrow d}^L)}{\lambda + i\mu + (k-i)\alpha}, \quad (i < d < k) \quad (27)$$

$$p_{i \rightarrow k}^L = 1 - \sum_{\ell=i}^{k-1} p_{i \rightarrow \ell}^L, \quad (i \leq k) \quad (28)$$

The summation in Eqs. (27) above denotes the possible intermediate depths  $\ell$  through which we can move from initial depth  $i$  to final depth  $d$ . The above system of equations involves linear and quadratic equations (including products of two unlike variables), and can be solved *symbolically* to find  $p_{i \rightarrow d}^L$  in closed form (see Appendix B).

---

<sup>2</sup>The definition given for  $p_{i \rightarrow d}^L$  applies in all cases except when  $j = k$  and  $d \in \{k-1, k\}$ . When  $j = k$ , we can never end in depth  $k$  when moving one step to the left; in this case, we interpret  $p_{i \rightarrow k}^L$  (or  $p_{i \rightarrow k-1}^L$ ) as the probability that we first moved one step left *by transitioning out of a state in depth  $k$  (or  $k-1$ )*.

## 7.2 Deriving $\dot{R}_{i,k}^L$ for the repeating portion

The system of equations for  $\dot{R}_{i,k}^L$  is as follows:

$$\dot{R}_{0,k}^L = \frac{z^k + \lambda \left( z\dot{R}_{0,k}^L + \sum_{\ell=1}^k \left\{ (p_{0 \rightarrow \ell}^L)(\dot{R}_{\ell,k}^L) \right\} \right) + k\alpha\dot{R}_{1,k}^L}{\lambda + k\alpha} \quad (29)$$

$$\dot{R}_{i,k}^L = \frac{z^{k-i} + \lambda \left( z\dot{R}_{i,k}^L + \sum_{\ell=i}^k \left\{ (p_{i \rightarrow \ell}^L)(\dot{R}_{\ell,k}^L) \right\} \right) + (k-i)\alpha\dot{R}_{i+1,k}^L}{\lambda + i\mu + (k-i)\alpha}, \quad (0 < i < k) \quad (30)$$

$$\dot{R}_{k,k}^L = \frac{1 + \lambda(z\dot{R}_{k,k}^L + \dot{R}_{k,k}^L)}{\lambda + k\mu} \quad (31)$$

In the above, we have used the fact that  $\dot{R}_{i,k+1}^L = z\dot{R}_{i,k}^L$  from Thm. 3. The above system of linear equations can be easily solved to find  $\dot{R}_{i,k}^L$  in closed form (see Appendix B).

## 7.3 Deriving $\dot{R}_{i,j}^H$ for the non-repeating portion

The system of equations for  $\dot{R}_{i,j}^H$  is as follows:

$$\dot{R}_{0,j}^H = \frac{z^j + \lambda\dot{R}_{0,j+1}^H + j\alpha\dot{R}_{1,j}^H}{\lambda + j\alpha}, \quad (j < k-1) \quad (32)$$

$$\dot{R}_{i,j}^H = \frac{z^{j-i} + \lambda\dot{R}_{i,j+1}^H + i\mu\dot{R}_{i,j-1}^H + (j-i)\alpha\dot{R}_{i+1,j}^H}{\lambda + i\mu + (j-i)\alpha}, \quad (0 < i < j < k-1) \quad (33)$$

$$\dot{R}_{i,i}^H = \frac{1 + \lambda\dot{R}_{i,i+1}^H + i\mu\dot{R}_{i-1,i-1}^H}{\lambda + i\mu}, \quad (0 < i < k-1) \quad (34)$$

$$\dot{R}_{i,k-1}^H = \frac{z^{k-1-i} + \lambda \left( \dot{R}_{i,k}^L + \sum_{\ell=i}^k \left\{ (p_{i \rightarrow \ell}^L)(\dot{R}_{\ell,k-1}^H) \right\} \right) + i\mu\dot{R}_{i,k-2}^H + (k-1-i)\alpha\dot{R}_{i+1,k-1}^H}{\lambda + i\mu + (k-1-i)\alpha}, \quad (i < k-1) \quad (35)$$

$$\dot{R}_{k-1,k-1}^H = 0 \quad (36)$$

Eqs. (32), (33), and (34), are simply based on the rate transitions in the non-repeating portion of the Markov chain. Eqs. (35) describe the rewards earned when starting in states in the non-repeating portion of the chain that can transition to the repeating portion of the chain via the border states. When we have an arrival in one of these states, we transition to the repeating portion of the chain, and after earning some reward, return to the non-repeating portion of the chain. Finally, Eq. (36)

guarantees that any transition to state  $(k-1, k-1)$  will end the renewal cycle. The above system of linear equations can again be easily solved to find  $\dot{R}_{i,j}^H$  in closed form (see Appendix B).

After solving for  $p_{i \rightarrow d}^L$ ,  $\dot{R}_{i,k}^L$  and  $\dot{R}_{i,j}^H$ , we can derive  $\dot{\mathcal{R}}$ , and consequently  $\widehat{N}_Q(z)$ , via Eq. (25).  $\widetilde{T}(s)$  can then be derived by using the fact  $\widetilde{T}(s) = \widetilde{T}_Q(s) \cdot \frac{\mu}{s+\mu} = \widehat{N}_Q(1 - \frac{s}{\lambda}) \cdot \frac{\mu}{s+\mu}$ .

## 8 M/M/3/setup

In this section we present the systems of equation in Section 7 for the special case  $k = 3$ . Solving this system yields a closed for solution for  $\widehat{N}_Q(z)$  for the M/M/3/setup system.

- The family of equations given in (26)–(28) are as follows:

$$\begin{aligned}
p_{0 \rightarrow 1}^L &= \frac{\lambda(p_{0 \rightarrow 1}^L)(p_{1 \rightarrow 1}^L) + 3\alpha(p_{1 \rightarrow 1}^L)}{\lambda + 3\alpha} \\
p_{0 \rightarrow 2}^L &= \frac{\lambda((p_{0 \rightarrow 1}^L)(p_{1 \rightarrow 2}^L) + (p_{0 \rightarrow 2}^L)(p_{2 \rightarrow 2}^L)) + 3\alpha(p_{1 \rightarrow 2}^L)}{\lambda + 3\alpha} \\
p_{0 \rightarrow 3}^L &= 1 - p_{0 \rightarrow 1}^L - p_{0 \rightarrow 2}^L \\
p_{1 \rightarrow 1}^L &= \frac{\lambda(p_{1 \rightarrow 1}^L)^2 + \mu}{\lambda + \mu + 2\alpha} \\
p_{1 \rightarrow 2}^L &= \frac{\lambda((p_{1 \rightarrow 1}^L)(p_{1 \rightarrow 2}^L) + (p_{1 \rightarrow 2}^L)(p_{2 \rightarrow 2}^L)) + 2\alpha(p_{2 \rightarrow 2}^L)}{\lambda + \mu + 2\alpha} \\
p_{1 \rightarrow 3}^L &= 1 - p_{1 \rightarrow 1}^L - p_{1 \rightarrow 2}^L \\
p_{2 \rightarrow 2}^L &= \frac{\lambda(p_{2 \rightarrow 2}^L)^2 + 2\mu}{\lambda + 2\mu + \alpha} \\
p_{2 \rightarrow 3}^L &= 1 - p_{2 \rightarrow 2}^L \\
p_{3 \rightarrow 3}^L &= 1
\end{aligned}$$

- The family of equations given in (29)–(31) are as follows:

$$\begin{aligned}
\dot{R}_{0,3}^L &= \frac{z^3 + \lambda(z\dot{R}_{0,3}^L + (p_{0 \rightarrow 1}^L)\dot{R}_{1,3}^L + (p_{0 \rightarrow 2}^L)\dot{R}_{2,3}^L + (p_{0 \rightarrow 3}^L)\dot{R}_{3,3}^L) + 3\alpha\dot{R}_{1,3}^L}{\lambda + 3\alpha} \\
\dot{R}_{1,3}^L &= \frac{z^2 + \lambda(z\dot{R}_{1,3}^L + (p_{1 \rightarrow 1}^L)\dot{R}_{1,3}^L + (p_{1 \rightarrow 2}^L)\dot{R}_{2,3}^L + (p_{1 \rightarrow 3}^L)\dot{R}_{3,3}^L) + 2\alpha\dot{R}_{2,3}^L}{\lambda + \mu + 2\alpha} \\
\dot{R}_{2,3}^L &= \frac{z + \lambda(z\dot{R}_{2,3}^L + (p_{2 \rightarrow 2}^L)\dot{R}_{2,3}^L + (p_{2 \rightarrow 3}^L)\dot{R}_{3,3}^L) + \alpha\dot{R}_{3,3}^L}{\lambda + 2\mu + \alpha} \\
\dot{R}_{3,3}^L &= \frac{1 + \lambda(1+z)(\dot{R}_{3,3}^L)}{\lambda + 3\mu}
\end{aligned}$$

- The family of equations given in (32)–(36) are as follows:

$$\begin{aligned}
\dot{R}_{0,0}^H &= \frac{1 + \lambda \dot{R}_{0,1}^H}{\lambda} \\
\dot{R}_{0,1}^H &= \frac{z + \lambda \dot{R}_{0,2}^H + \alpha \dot{R}_{1,1}^H}{\lambda + \alpha} \\
\dot{R}_{1,1}^H &= \frac{1 + \lambda \dot{R}_{1,2}^H + \mu \dot{R}_{0,0}^H}{\lambda + \mu} \\
\dot{R}_{0,2}^H &= \frac{z^2 + \lambda(\dot{R}_{0,3}^L + (p_{0 \rightarrow 1}^L) \dot{R}_{1,2}^H + (p_{0 \rightarrow 2}^L) \dot{R}_{2,2}^H + (p_{0 \rightarrow 3}^L) \dot{R}_{2,2}^H) + 2\alpha \dot{R}_{1,2}^H}{\lambda + 2\alpha} \\
\dot{R}_{1,2}^H &= \frac{z + \lambda(\dot{R}_{1,3}^L + (p_{1 \rightarrow 1}^L) \dot{R}_{1,2}^H + (p_{1 \rightarrow 2}^L) \dot{R}_{2,2}^H + (p_{1 \rightarrow 3}^L) \dot{R}_{2,2}^H) + \mu \dot{R}_{1,1}^H + \alpha \dot{R}_{2,2}^H}{\lambda + \mu + \alpha} \\
\dot{R}_{2,2}^H &= 0
\end{aligned}$$

Note that no equations from family (33) are presented above, as there do not exist integers  $i$  and  $j$  such that  $0 < i < j < k - 1$  when  $k = 3$ ; equations from this family will appear in the system of equations for the M/M/k/setup when  $k \geq 4$ .

- Finally, after solving these equations, we compute

$$\dot{\mathcal{R}} = \frac{1 + \lambda \dot{R}_{2,3}^L + 2\mu \dot{R}_{1,1}^H}{\lambda + 2\mu}$$

This allows us to compute

$$\widehat{N_Q}(z) = \frac{\dot{\mathcal{R}}}{\dot{\mathcal{R}}|_{z \rightarrow 1}}$$

Carrying out the computations in this section *symbolically* with Mathematica, we find a closed-form expression for  $\dot{\mathcal{R}}$  as follows:

$$\begin{aligned}
\dot{\mathcal{R}} &= \frac{1}{\lambda + 2\mu} \left( 1 - \frac{H\lambda(D - 2z^2\lambda - 3\mu + 6z\mu + A)}{z\lambda - 3\mu} - \frac{2}{D} \left( X - \left( (\alpha + \lambda)(B + \lambda + \mu) \left( \frac{X\lambda}{\alpha + \lambda} \right. \right. \right. \right. \\
&\quad + \mu \left( 1 - \mu \left( -\frac{1}{\lambda} - \frac{z}{\alpha + \lambda} - \frac{\lambda}{V} (2z^2 + \lambda(G - z\lambda)^{-1} (6GJ^2Yz + 2z^3 - 6GJWz^2(-4\alpha + B - \lambda + \mu) \right. \right. \right. \\
&\quad \left. \left. \left. + (z\lambda - 3\mu)^{-1}GJ^2(\lambda^2 + 3W(4\alpha - B + \lambda - \mu)(2\alpha^2 + 2\alpha(B + 4\mu - 2A) + (B + \lambda - \mu)(\mu - A)) \right. \right. \right. \\
&\quad \left. \left. \left. - 3Y(D - 3\mu + A) - 3(\alpha^2 + \mu(\lambda + \mu - A) + B(-\mu + A) + \alpha(-2(B + \mu) + A)) \right) \right) \right) \right) / \\
&\quad \left( \lambda \left( B + \lambda + \mu + \frac{\mu(-V + J\mu(3(\alpha + \lambda)\mu - G(10\alpha - 3B + 3\lambda))}{D(2\alpha + \lambda)} \right) \right) \right),
\end{aligned}$$

where we make use of the following auxiliary expressions:

$$\begin{aligned}
D &= \alpha + \lambda + \mu, \\
A &= \sqrt{\mu + (D + \mu)^2 - 8\lambda}, \\
B &= \sqrt{\mu + (D + \alpha)^2 - 4\lambda}, \\
G &= 3\alpha + \lambda, \\
J &= \frac{1}{G - 3\mu}, \\
H &= 1/(\alpha + \lambda - 2z\lambda + 2\mu + A), \\
W &= 1/(2\alpha + B + \lambda - 2z\lambda + \mu), \\
V &= 2(\alpha + \lambda)(2\alpha + \lambda), \\
X &= D \left( z + W\lambda \left( 2z^2 + 2HJz (8\alpha^2 + \alpha(5B + 5\lambda - \mu - 4A) + (B + \lambda - \mu)(\lambda - 2\mu - A)) - \frac{JH}{z\lambda - 3\mu} (6\alpha^3 \right. \right. \\
&\quad \left. \left. + \alpha^2(3B + (11 + 4z)\lambda - 29\mu + 6A) + \alpha(6\lambda^2 - 17B\mu - 11\mu^2 + 3BA + 7\mu A + \lambda(4(1 + z)B + (-23 + 16z)\mu \right. \right. \\
&\quad \left. \left. + (5 - 8z)A) + (B + \lambda - \mu)(\lambda^2 + \mu(2\mu + A) + \lambda((-5 + 2z)\mu + A - 2zA)) \right) \right), \\
Y &= HW (18\alpha^3 + \lambda(B - 2zB + \lambda + (-3 + 2z)\mu)(\lambda - 2\mu - A) + \alpha^2(15B + (15 + 14z)\lambda - 9(\mu + 2A)) \\
&\quad + \alpha(6(1 + z)\lambda^2 - 3(B - 3\mu)(2\mu + A) - \lambda(2(-4 + 5z)B + 2(9 + 5z)\mu + (9 - 2z)A))).
\end{aligned}$$

We numerically verified the correctness of  $E[N_Q]$  obtained via the expression for  $\hat{N}_Q(z)$  (from  $\hat{\mathcal{R}}$  above) with matrix-analytic methods.

## 9 M/M/k/setup/delayedoff

The Markov chain for M/M/k/setup/delayedoff is shown in Fig. 2. Our renewal state this time will be  $(k, k - 1)$ ; thus,  $\dot{\mathcal{R}}$ , the reward earned when going from  $(k, k - 1)$  back to  $(k, k - 1)$  can be expressed as:

$$\dot{\mathcal{R}} = \frac{1 + (k - 1)\mu\dot{R}_{k,k-2}^H + \lambda\dot{R}_{k,k}^L + \beta\dot{R}_{k-1,k-1}^H}{\lambda + (k - 1)\mu + \beta} \quad (37)$$

The analysis for M/M/k/setup/delayedoff via RRR is very similar to that of M/M/k/setup in Section 7 above. In fact, since the repeating portion for the two chains is the same, the system of equations for  $p_{i \rightarrow d}^L$  and  $\dot{R}_{i,j}^L$  is identical, but the non-repeating portion for the two chains is different. We now set up the system of equations for solving  $\dot{R}_{i,j}^H$ .



## 9.1 Deriving $\dot{R}_{i,j}^H$ for the non-repeating portion

The system of equations for  $\dot{R}_{i,j}^H$  is as follows:

$$\dot{R}_{0,j}^H = \frac{z^j + \lambda \dot{R}_{0,j+1}^H + j\alpha \dot{R}_{1,j}^H}{\lambda + j\alpha}, \quad (j < k-1) \quad (38)$$

$$\dot{R}_{i,0}^H = \frac{1 + \lambda \dot{R}_{i,1}^H + i\beta \dot{R}_{i-1,0}^H}{\lambda + i\beta}, \quad (0 < i \leq k) \quad (39)$$

$$\dot{R}_{k,j}^H = \frac{1 + \lambda \dot{R}_{k,j+1}^H + j\mu \dot{R}_{k,j-1}^H + (k-j)\beta \dot{R}_{k-1,j}^H}{\lambda + j\mu + (k-j)\beta}, \quad (1 \leq j < k-1) \quad (40)$$

$$\dot{R}_{i,j}^H = \frac{z^{j-i} + \lambda \dot{R}_{i,j+1}^H + i\mu \dot{R}_{i,j-1}^H + (j-i)\alpha \dot{R}_{i+1,j}^H}{\lambda + i\mu + (j-i)\alpha}, \quad (0 < i < j < k-1) \quad (41)$$

$$\dot{R}_{i,i}^H = \frac{1 + \lambda \dot{R}_{i,i+1}^H + i\mu \dot{R}_{i,i-1}^H}{\lambda + i\mu}, \quad (0 < i < k-1) \quad (42)$$

$$\dot{R}_{i,j}^H = \frac{1 + \lambda \dot{R}_{i,j+1}^H + j\mu \dot{R}_{i,j-1}^H + (i-j)\beta \dot{R}_{i-1,j}^H}{\lambda + j\mu + (i-j)\alpha}, \quad (0 < j < i < k) \quad (43)$$

$$\dot{R}_{i,k-1}^H = \frac{z^{k-1-i} + \lambda \left( \dot{R}_{i,k}^L + \sum_{\ell=i}^k \left\{ (p_{i \rightarrow \ell}^L) (\dot{R}_{\ell,k-1}^H) \right\} \right) + i\mu \dot{R}_{i,k-2}^H + (k-1-i)\alpha \dot{R}_{i+1,k-1}^H}{\lambda + i\mu + (k-1-i)\alpha}, \quad (i \leq k-1) \quad (44)$$

$$\dot{R}_{k,k-1}^H = 0 \quad (45)$$

The above system of linear equations can again be solved to find  $\dot{R}_{i,j}^H$  in closed form. This yields  $\dot{\mathcal{R}}$ , and consequently  $\widehat{N}_Q(z)$ , via Eq. (37).

## 10 M/M/k/setup/sleep

The Markov chain for M/M/k/setup/sleep is shown in Fig. 3. The analysis for M/M/k/setup/sleep via RRR is again similar to that of M/M/k/setup in Section 7. The only difference is in the setup transition rate (downwards transition arrows in the Markov chain): For the M/M/k/setup, the setup rate in state  $(i, j)$  is  $\alpha \cdot \min\{j - i, k - i\}$ . For the M/M/k/setup/sleep, the setup rate in state  $(i, j)$  is more complicated. When  $i \geq s$ , the setup rate is still  $\alpha \cdot \min\{j - i, k - i\}$ . However, if  $i < s$ , the setup rate is  $\omega \cdot (j - i)$  if  $j \leq s$  and  $\omega \cdot (s - i) + \alpha \cdot \min\{j - s, k - s\}$  if  $j > s$ .

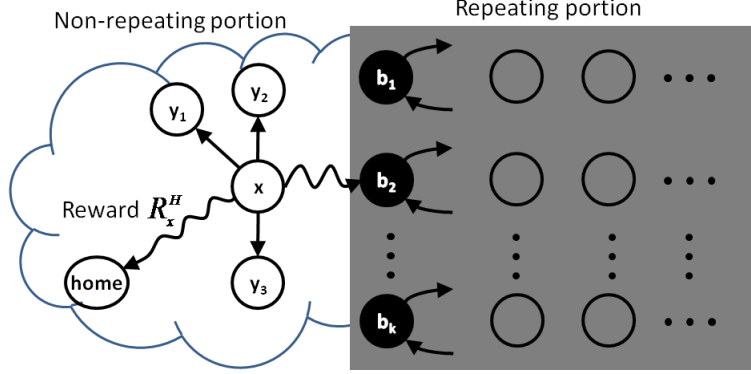


Figure 6: Illustration of the RRR technique for an arbitrary repeating 2-dimensional Markov chain. The repeating portion is highlighted in gray and the border states,  $b_i$ , are shaded black. Note that  $y_i$  are the neighbors of  $x$ .

This can be explained based on the M/M/k/setup/sleep model description in Section 3.3 and the Markov chain in Fig. 3. Based on the above setup rates, we can easily modify the M/M/k/setup sets of equations for  $p_{i \rightarrow d}^L$ ,  $\dot{R}_{i,k}^L$  and  $\dot{R}_{i,j}^H$  from Sections 7.1, 7.2 and 7.3 respectively, to represent the M/M/k/setup/sleep system of equations. The equation for  $\hat{\mathcal{R}}$  will change accordingly.

## 11 The Generalized Recursive Renewal Reward technique

The RRR technique can be applied to solving chains beyond M/M/k/setup and its variants. In this section, we give a brief overview of applying RRR to find the expected rate of earning reward for an infinite Markov chain with a repeating structure.

Fig. 6 illustrates the types of chains that RRR can solve. Within Fig. 6, to get closed forms, it suffices (but is not necessary) that all transitions in the repeating portion are skip-free and all up/down arrows are unidirectional. RRR can be applied to more general structures, but may require solving systems of higher-order polynomials.

We first partition the Markov chain into a finite non-repeating portion and an infinite repeating portion as in Fig. 6; in principle, this partition is not unique. Then, we fix a renewal point, or home state, within the non-repeating portion. For each state,  $x$ , in the non-repeating portion of the chain, we write an equation for the mean reward earned in traveling from  $x$  to the home state. The mean reward,  $R_x^H$ , consists of the mean reward during our residence in  $x$ , and a weighted combination of the rewards  $R_y^H$ , where  $y$  is a neighbor of  $x$ , as in Fig. 6.<sup>3</sup> Observe that all the equations for states in the non-repeating portion are linear (as in Section 7.3). Since the chain is irreducible, at least one state in the non-repeating portion of the chain can transition directly to a state in the repeating portion of the chain. We refer to the states in the repeating portion that are directly accessible from the non-repeating portion as *border states*. These are shown as  $b_i$  in Fig. 6. Before we can solve

<sup>3</sup>If  $y$  is the home state, we replace  $R_y^H$  with 0 in this weighted combination.

the set of linear equations corresponding to each state in the non-repeating portion, we will need to first derive equations for the border states.

We break the mean reward earned when returning home from each border state, say  $b$ , into two parts: (i) the mean reward earned from the time we enter  $b$  until we leave the repeating portion, and (ii) the mean reward earned from when we first exit the repeating portion until returning home. For (i), the reward equation can be expressed in terms of the reward equations for the neighbors of  $b$  in the repeating portion. The fact that the chain has a repeating structure allows us to express the reward from any state in the repeating portion as a function of the rewards of the border states by using “recursion theorems” (as in Thms. 2 and 3). The precise recursion theorems might look different from Thms. 2 and 3 which were for the M/M/k/setup.

In order to compute (ii), we first need the probability distribution over the set of states in the non-repeating portion that we transition to after exiting the repeating portion. The reward in (ii) can then be expressed as a weighted linear combination of the reward earned before returning home from these “exit states.” Solving for these probabilities may also require recursion theorems (as in Thm. 5); in general, the probability equations will require solving a system of polynomial equations (for the M/M/k/setup, these are provably no worse than quadratic as shown in Section 7.1).

Finally, we will have a system of linear equations for the mean rewards earned before returning home from the states in the non-repeating portion and the border states. This yields the mean reward earned during a renewal cycle from home to home; mean time is found analogously.

## 12 Applications

In this section we use our analytical results to evaluate the performance of M/M/k, M/M/k/setup, M/M/k/setup/delayedoff and M/M/k/setup/sleep. In particular, we will be interested in the mean response time,  $E[T]$ , and the mean power consumption,  $E[P]$ , under these policies. Throughout, we assume a load of  $\rho = \frac{\lambda}{k\mu} = 0.3$  (or 30% load), setup times of  $\frac{1}{\alpha} = 100\text{s}$  (when the server is off) and  $\frac{1}{\omega} = 25\text{s}$  (when the server is sleeping), and power consumption values of  $P_{peak} = 200\text{W}$ ,  $P_{idle} = 140\text{W}$ , and  $P_{sleep} = 14\text{W}$ . These parameter values are based on empirical measurements from prior work [4, 10]. We consider job sizes with mean  $E[S] = 1\text{s}$  (typical web workloads [6]),  $E[S] = 10\text{s}$  (database queries or secured transactions), and  $E[S] = 100\text{s}$  (file download or upload), and system sizes ranging from  $k = 5$  to  $k = 100$  servers.

The M/M/k policy keeps  $k$  servers always on. Servers that are not busy serving jobs are left idle. The M/M/k/setup policy (see Section 3.1) immediately turns off idle servers to save power. However, restarting an off server requires a setup time of  $\frac{1}{\alpha} = 100\text{s}$ . The M/M/k/setup/delayedoff policy (see Section 3.2) is the same as the M/M/k/setup policy, except that idle servers wait for an exponentially distributed amount of time with mean  $t_{wait} = \frac{1}{\beta}$  before turning off. The performance of this policy depends on the choice of the  $t_{wait}$  parameter. Finally, the M/M/k/setup/sleep policy (see Section 3.3) is the same as the M/M/k/setup policy, except that  $s$  of the  $k$  servers go to sleep as opposed to turning off, when idle. A sleeping server has a small setup time of  $\frac{1}{\omega} = 25\text{s}$ . The per-

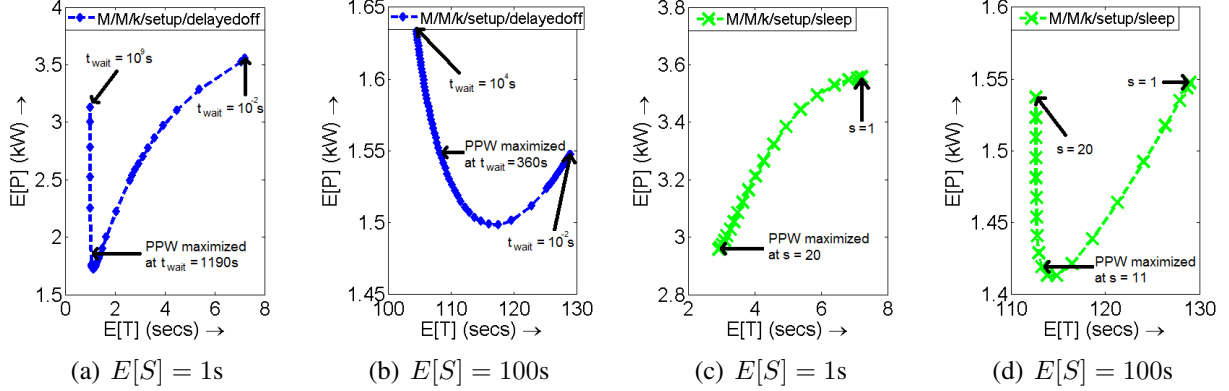


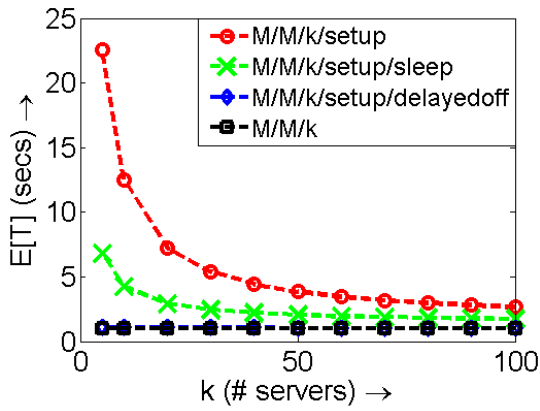
Figure 7:  $E[P]$  versus  $E[T]$  for various values of  $t_{wait}$  and  $s$ .

formance of this policy depends on the choice of the  $s$  parameter. Before comparing the above four policies, we first discuss how we choose the parameter value of  $t_{wait}$  for M/M/k/setup/delayedoff and  $s$  for M/M/k/setup/sleep.


## 12.1 Choosing optimal parameter values

The tradeoff between  $E[P]$  and  $E[T]$  for M/M/k/setup/delayedoff is shown in Figs. 7(a) and 7(b). Each plotted point represents an  $(E[T], E[P])$  pair associated with a specific value of  $t_{wait}$ . Intuitively, as  $t_{wait}$  increases,  $E[T]$  decreases since we avoid setup times. Moreover, before some threshold  $t_{wait}$ ,  $E[P]$  decreases as  $t_{wait}$  increases, because we avoid consuming power at peak rate by repeatedly putting servers in setup. However, beyond this threshold  $t_{wait}$ ,  $E[P]$  starts increasing on account of idle servers. Thus, as  $t_{wait}$  increases, we get the plots in Figs. 7(a) and 7(b), from right to left. We choose the  $t_{wait}$  value that optimizes (i.e., maximizes) the popular *Performance-Per-Watt* metric [10, 7], given by  $PPW = (E[T] \cdot E[P])^{-1}$ . These optimal values are shown in Figs. 7(a) and 7(b). We find that the optimal  $t_{wait}$  value decreases with an increase in  $E[S]$ .

Figs. 7(c) and 7(d) illustrate the tradeoff between  $E[P]$  and  $E[T]$  under M/M/k/setup/sleep for different values of  $s$ . Intuitively, as  $s$  increases,  $E[T]$  decreases since we benefit from faster setup times afforded by sleeping servers. As  $s$  increases,  $E[P]$  first decreases since we avoid the severe power penalty of longer setup times. But beyond a certain  $s$ ,  $E[P]$  increases on account of the sleeping servers. Thus, as  $s$  increases, we get the plots in Figs. 7(c) and 7(d), from right to left. Note that  $E[P]$  monotonically decreases for the case of  $E[S] = 1s$  in Fig. 7(c). This is because  $\frac{1}{\alpha} \gg E[S]$ , and thus, the decrease in power consumption by avoiding power penalties of longer setup times outweighs the increase in power consumption because of  $P_{sleep}$ . We choose the  $s$  value that optimizes the PPW metric, as indicated in Figs. 7(c) and 7(d). We find that the optimal  $s$  value decreases with an increase in  $E[S]$ .

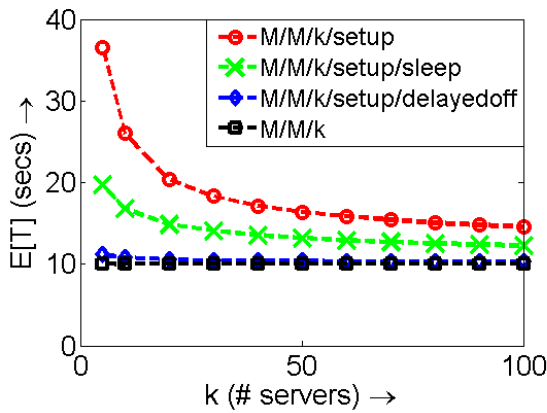


(a)  $E[T]$  versus  $k$




(b)  $E[P]$  versus  $k$

Figure 8: Results when mean job size,  $E[S] = 1$ .

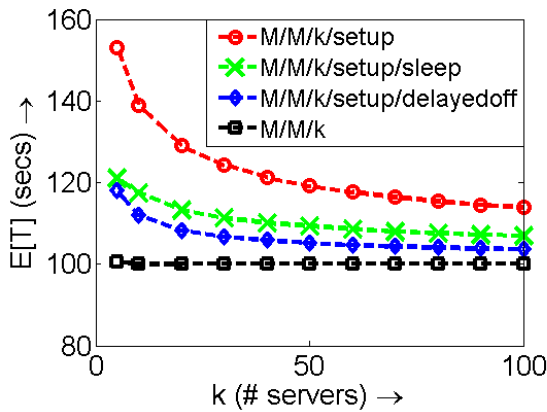


(a)  $E[T]$  versus  $k$




(b)  $E[P]$  versus  $k$

Figure 9: Results when mean job size,  $E[S] = 10$ .



(a)  $E[T]$  versus  $k$



(b)  $E[P]$  versus  $k$

Figure 10: Results when mean job size,  $E[S] = 100$ .




Figure 11:  $\text{Var}(T)$  versus  $k$  under mean job size of  $E[S] = 1\text{s}$ .

## 12.2 Comparison of all policies

Fig. 8 shows our results for  $E[T]$  and  $E[P]$  as a function of  $k$  for the case of  $E[S] = 1\text{s}$ . Comparing M/M/k (squares) and M/M/k/setup (circles), we see that M/M/k/setup has a much higher  $E[T]$ , and only a slightly lower  $E[P]$ . In fact, when  $k$  is low,  $E[P]$  for M/M/k/setup is *higher* than that of M/M/k. This is because of the power penalty involved in the setup cost. Thus, **M/M/k/setup is not a good policy for small job sizes**. The M/M/k/setup/sleep (crosses) has lower  $E[T]$  and lower  $E[P]$  than the M/M/k/setup. Thus, **using sleep modes improves the M/M/k/setup policy**. Finally, we see that M/M/k/setup/delayedoff (diamonds) has  $E[T]$  virtually as low as that of M/M/k, and has the lowest power consumption among all other policies. Thus, **M/M/k/setup/delayedoff is superior to all the other policies for small job sizes**. The reason for lower  $E[P]$  under M/M/k/setup/delayedoff is because of  $t_{wait}$  which avoids unnecessary setups (and the associated power penalties).

Figs. 9 and 10 show our results for the case of  $E[S] = 10\text{s}$  and  $E[S] = 100\text{s}$  respectively. The  $E[T]$  results for these job sizes are qualitatively similar to the results for  $E[S] = 1\text{s}$ . The percentage difference between the  $E[T]$  under different policies goes down as  $E[S]$  goes up. This is because the setup time is not changing as  $E[S]$  goes up, and thus, the queueing delay caused by setup times is not as severe for large  $E[S]$ . Note that the  $E[T]$  under M/M/k/setup/delayedoff actually goes up as  $E[S]$  goes up. This is a side-effect of the optimal  $t_{wait}$  setting which trades off lower  $E[P]$  at the expense of a slightly higher  $E[T]$  for bigger job sizes.

The  $E[P]$  results for different job sizes indicate that  $E[P]$  **under M/M/k/setup and M/M/k/setup/sleep decreases with an increase in job size, and approaches the  $E[P]$  of M/M/k/setup/delayedoff**. This is because an increase in  $E[S]$  necessitates an increase in the inter-arrival time, given fixed load,  $\rho$ . Thus, servers now spend more time in the off or sleep states, and consequently, consume less power. In fact, the M/M/k/setup/sleep has *lower*  $E[P]$  as compared to M/M/k/setup/delayedoff for the case of  $E[S] = 100\text{s}$ . We take a closer look at these two policies in Section 12.3. Note that under M/M/k,  $E[P] = k \cdot \rho \cdot P_{peak} + k \cdot (1 - \rho) \cdot P_{idle}$ , which is linear

in  $k$  and independent of  $E[S]$ .

The  $E[T]$  results for these job sizes are qualitatively similar to the results for  $E[S] = 1\text{s}$ . The percentage difference between the  $E[T]$  under different policies goes down as  $E[S]$  goes up. This is because the setup time is not changing as  $E[S]$  goes up, and thus, the queueing delay caused by setup times is not as severe for large  $E[S]$ . Note that the  $E[T]$  under M/M/k/setup/delayedoff actually goes up as  $E[S]$  goes up. This is a side-effect of the optimal  $t_{wait}$  setting which trades off lower  $E[P]$  at the expense of a slightly higher  $E[T]$  for bigger job sizes.

As mentioned in Section 7, RRR also provides closed-form solutions for higher moments of response time and power. Fig. 11 shows our results for  $Var(T)$ , the variability in response time, for the case of  $E[S] = 1\text{s}$ . We see that  $Var(T)$  follows the same trends as  $E[T]$  in Fig. 8(a). Note that  $Var(T)$  is close to 1 for M/M/k and M/M/k/setup/delayedoff. Also,  $Var(T)$  converges to 1 for all policies for high  $k$ . This is because  $Var(T)$  converges to  $Var(S)$  (no queueing delay) in these cases, and since  $S$  is exponentially distributed with mean  $E[S] = 1\text{s}$ , we get  $Var(T) \rightarrow Var(S) = 1\text{s}^2$ .

All the results above assumed exponential setup times and exponential delay times. However, in real-world scenarios, these times would be deterministic. We use simulations to find  $E[T]$  and  $E[P]$  under deterministic setup times for all the above cases. We find that the relative ordering of the policies and the trends in  $E[T]$  and  $E[P]$  do not change significantly, despite the fact that all values become slightly higher due to the setup rates no longer being additive.

### 12.3 A closer look at M/M/k/setup/delayedoff versus M/M/k/setup/sleep

Fig. 12 shows the tradeoff between  $E[P]$  and  $E[T]$  for M/M/k/setup/delayedoff and M/M/k/setup/sleep for  $E[S] = 100\text{s}$ . These plots are identical to Figs. 7(b) and 7(d). We see that no policy dominates the other. If we are more concerned about reducing  $E[P]$ , M/M/k/setup/sleep is the better choice. However, if we are more concerned about reducing  $E[T]$ , M/M/k/setup/delayedoff is the better choice. Interestingly, by taking a probabilistic mixture of the two policies, we can find additional policies that are superior to the M/M/k/setup/delayedoff and the M/M/k/setup/sleep. The probabilistic mixture can be obtained by taking the convex hull of the two policies, as shown by the dashed line in Fig. 12. This suggests the potential for a policy that combines M/M/k/setup/sleep with delayedoff.

## 13 Conclusion

In this paper we develop a new analysis technique, Recursive Renewal Reward (RRR), which allows us to solve the M/M/k/setup class of Markov chains. RRR is very intuitive, easy to apply, and can be used to analyze many important Markov chains that have a repeating structure. RRR combines renewal reward theory with the development of recursion theorems for the Markov chain

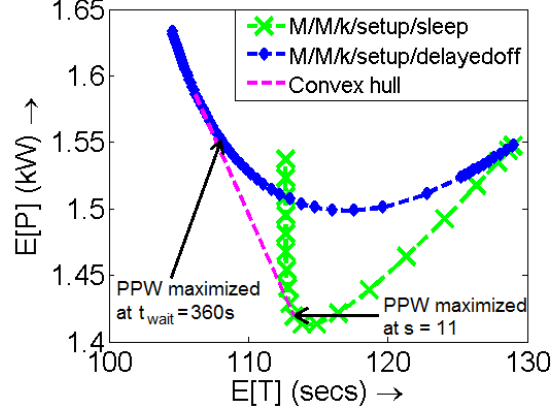


Figure 12:  $E[P]$  versus  $E[T]$  for various values of  $t_{wait}$  and  $s$  under mean job size of  $E[S] = 100s$ .

to yield exact, closed form results for metrics of interest such as the transform of time in system and the transform of power consumed by the system. RRR reduces the solution of the M/M/k/setup chains to solving  $k$  quadratic equations and a system of  $O(k^2)$  linear equations. On an Intel Core i5-based processor machine we found RRR to be almost 5-10 times faster than the iterative matrix-analytic based methods, when using standard MATLAB implementations of both methods.

While we have only considered the M/M/k/setup, the M/M/k/setup/delayedoff, and the M/M/k/setup/sleep in this paper, we have also been able to use RRR for the derivation of exact, closed-form solutions for other important Markov chains with a repeating structure such as: (i) M/M/k/stag [9], wherein at most one server can be in setup, (ii) M/M/k/setup-threshold, wherein servers are turned on based on some threshold for number of jobs in queue, (iii) M/M/k/disasters, wherein the system can empty abruptly due to disasters, and (iv) M/E<sub>2</sub>/k, where the job size distribution is Erlang-2. We have also been able to apply RRR to analyze other Markov chains such as: (i) M<sub>t</sub>/M/1, where the arrival process is Poisson with a time dependent parameter, (ii) M/H<sub>2</sub>/k, where the job size distribution is a 2-phase hyperexponential, and (iii) M<sup>x</sup>/M/k, where there is a Poisson batch arrival process. In the above three cases, RRR reduces the analysis to the problem of solving a polynomial equation, with degree  $> 2$ . In general, RRR should be able to reduce the analysis of any 2-dimensional Markov chain which is finite in one dimension and possibly infinite (with repeating structure) in the other to the problem of solving a polynomial equation, for any reward that can be expressed as a recursive function over the states of the Markov chain.

While not shown in this paper, it is possible to derive an explicit rate matrix for the M/M/k/setup, which leads to closed-form expressions for the limiting probabilities. While RRR does not utilize the rate matrix in any way, it appears that the set of Markov chains that can be solved in closed form via RRR all have an explicit rate matrix.



## References

- [1] ADAN, I., AND RESING, J. A class of Markov processes on a semi-infinite strip. Tech. Rep. 99-03, Eindhoven University of Technology, Department of Mathematics and Computing Sciences, 1999.
- [2] ADAN, I., AND VAN DER WAL, J. Combining make to order and make to stock. *OR Spektrum* 20 (1998), 73–81.
- [3] ARTALEJO, J. R., ECONOMOU, A., AND LOPEZ-HERRERO, M. J. Analysis of a multiserver queue with setup times. *Queueing Syst. Theory Appl.* 51, 1-2 (2005), 53–76.
- [4] BARROSO, L. A., AND HÖLZLE, U. The Case for Energy-Proportional Computing. *IEEE Computer* 40, 12 (2007), 33–37.
- [5] CASTELLANOS, M., CASATI, F., SHAN, M.-C., AND DAYAL, U. iBOM: A platform for intelligent business operation management. In *Proceedings of the 21st International Conference on Data Engineering* (Tokyo, Japan, 2005), ICDE '05, pp. 1084–1095.
- [6] DECANDIA, G., HASTORUN, D., JAMPANI, M., KAKULAPATI, G., LAKSHMAN, A., PILCHIN, A., SIVASUBRAMANIAN, S., VOSSHALL, P., AND VOGELS, W. Dynamo: Amazon’s highly available key-value store. In *Proceedings of twenty-first ACM SIGOPS Symposium on Operating Systems Principles* (Stevenson, WA, 2007), SOSP '07, pp. 205–220.
- [7] GANDHI, A., GUPTA, V., HARCHOL-BALTER, M., AND KOZUCH, M. Optimality Analysis of Energy-Performance Trade-off for Server Farm Management. *Performance Evaluation* 67 (2010), 1155–1171.
- [8] GANDHI, A., AND HARCHOL-BALTER, M. How Data Center Size Impacts the Effectiveness of Dynamic Power Management. *49th Annual Allerton Conference on Communication, Control, and Computing* (2011).
- [9] GANDHI, A., HARCHOL-BALTER, M., AND ADAN, I. Server farms with setup costs. *Performance Evaluation* 67 (2010), 1123–1138.
- [10] GANDHI, A., HARCHOL-BALTER, M., AND KOZUCH, M. Are Sleep States Effective in Data Centers? *3rd IEEE International Green Computing Conference* (2012).
- [11] HORVATH, T., AND SKADRON, K. Multi-mode energy management for multi-tier server clusters. In *Proceedings of the 17th International Conference on Parallel Architectures and Compilation Techniques* (Toronto, Ontario, Canada, 2008), PACT '08, pp. 270–279.
- [12] KEILSON, J., AND SERVI, L. A distributional form of little’s law. *Operations Research Letters* 7, 5 (1988), 223 – 227.

- [13] KIM, J., AND ROSING, T. S. Power-aware resource management techniques for low-power embedded systems. In *Handbook of Real-Time and Embedded Systems*, S. H. Son, I. Lee, and J. Y.-T. Leung, Eds. Taylor-Francis Group LLC, 2006.
- [14] KLEINROCK, L. *Queueing Systems, Volume I: Theory*. Wiley-Interscience, 1975.
- [15] KRIOUKOV, A., MOHAN, P., ALSPAUGH, S., KEYS, L., CULLER, D., AND KATZ, R. NapSAC: Design and implementation of a power-proportional web cluster. In *Proceedings of the First ACM SIGCOMM Workshop on Green Networking* (New Delhi, India, 2010), Green Networking '10, pp. 15–22.
- [16] LATOUCHE, G., AND RAMASWAMI, V. *Introduction to Matrix Analytic Methods in Stochastic Modeling*. ASA-SIAM, Philadelphia, 1999.
- [17] LEVY, Y., AND YECHIALI, U. An M/M/s queue with servers' vacations. *INFOR 14* (1976), 153–163.
- [18] MEISNER, D., GOLD, B. T., AND WENISCH, T. F. PowerNap: Eliminating server idle power. In *Proceeding of the 14th international conference on Architectural Support for Programming Languages and Operating Systems* (Washington, DC, 2009), ASPLOS '09, pp. 205–216.
- [19] MITRANI, I. Managing performance and power consumption in a server farm. *Annals of Operations Research* (2012), 1–14.
- [20] QIN, W., AND WANG, Q. Modeling and control design for performance management of web servers via an IPV approach. *IEEE Transactions on Control Systems Technology 15*, 2 (March 2007), 259–275.
- [21] TIAN, N., LI, Q.-L., AND GAO, J. Conditional stochastic decompositions in the M/M/c queue with server vacations. *Stochastic Models 15*, 2 (1999), 367–377.
- [22] VAN HOUTD, B., AND VAN LEEUWAARDEN, J. Triangular M/G/1-Type and Tree-Like Quasi-Birth-Death Markov Chains. *INFORMS Journal on Computing 23*, 1 (2011), 165–171.
- [23] VAN LEEUWAARDEN, J., AND WINANDS, E. Quasi-birth-and-death processes with an explicit rate matrix. *Stochastic models 22*, 1 (2006), 77–98.
- [24] WELCH, P. On a generalized M/G/1 queueing process in which the first customer of each busy period receives exceptional service. *Operations Research 12* (1964), 736–752.
- [25] XU, X., AND TIAN, N. The M/M/c Queue with (e, d) Setup Time. *Journal of Systems Science and Complexity 21* (2008), 446–455.
- [26] ZHANG, Z. G., AND TIAN, N. Analysis on queueing systems with synchronous vacations of partial servers. *Performance Evaluation 52*, 4 (2003), 269 – 282.

## A Recursion theorems

**Theorem 1** (Recursion theorem for mean time).

For the M/M/k/setup, the mean time to move one step left from state  $(i, j)$ ,  $T_{i,j}^L$ , is the same for all  $j \geq k$ .

*Proof.* For any  $j \geq k$ , observe that when moving one step left from any state  $(i, j)$ , we only visit states with level  $j$  or greater, until the final transition to level  $j - 1$ . Hence,  $T_{i,j}^L$  depends only on the structure of the “subchain” of the M/M/k/setup consisting of levels  $\{j, j + 1, \dots\}$ , including transition rates to level  $j - 1$ . Now consider the subchain for each  $j \geq k$ ; these subchains are isomorphic, by the fact that the chain is repeating from level  $k$  onward. Hence, the time to move one step left is the same regardless of the initial level  $j \geq k$ .  $\square$

**Theorem 2** (Recursion theorem for mean reward).

For the M/M/k/setup, the mean reward earned in moving one step left from state  $(i, j + 1)$ ,  $R_{i,j+1}^L$ , satisfies  $R_{i,j+1}^L = R_{i,j}^L + T_{i,j}^L$  for all  $j \geq k$ , where the reward tracks the number of jobs in the system.

*Proof.* Consider the process of moving one step left from a given state  $(i, j)$  where  $j \geq k$ . At the same time, consider the same process where everything is shifted over one level to the right, so that the initial state is  $(i, j + 1)$ . At any point in time, the number of jobs seen by the second process is exactly one greater than that seen by the first process. Therefore, the total number of jobs accumulated (total reward) during the second process is  $T_{i,j}^L$  greater than that of the first process, since the duration of both processes is  $T_{i,j}^L$  by Theorem 1.  $\square$

**Theorem 3** (Recursion theorem for transform of reward).

For the M/M/k/setup,  $\hat{R}_{i,j+1}^L = z \cdot \hat{R}_{i,j}^L$ , for all  $j \geq k$ , where  $\hat{R}$  tracks the  $z$ -transform of the number of jobs in the system.

*Proof.* The proof is identical to that of Theorem 2, except that in any moment in time the second process (starting in level  $(i, j + 1)$ ) earns  $z$  times as much reward as the first process (starting at  $(i, j)$ ).  $\square$

**Theorem 4** (Recursion theorem for transform of power).

For the M/M/k/setup,  $\hat{E}_{i,j+1}^L = \hat{E}_{i,j}^L = T_{i,j}^L \cdot z^{k \cdot P_{peak}}$ , for all  $j \geq k$ .

*Proof.* When  $j \geq k$ , all  $k$  servers are either on or in setup, putting power consumption at  $k \cdot P_{peak}$ . So the transform of power usage is  $\hat{E}_{i,j}^L = T_{i,j}^L \cdot z^{k \cdot P_{peak}}$ . It then follows immediately from Theorem 1 that  $\hat{E}_{i,j+1}^L = \hat{E}_{i,j}^L$ .  $\square$

**Theorem 5** (Recursion theorem for probability).

For the M/M/k/setup, for each  $0 \leq d \leq k$  and for each  $0 \leq i \leq k$ ,  $p_{i \rightarrow d}^L$  is the same for all  $j \geq k$ .

*Proof.* Recall that  $p_{i \rightarrow d}^L$  is the probability that, given that we start at depth  $i$ , we end at depth  $d$  when moving one step to the left, except when  $j = k$  and  $d \in \{k - 1, k\}$ ; in these cases we interpret  $p_{i \rightarrow k}^L$  (or  $p_{i \rightarrow k-1}^L$ ) as the probabilities that we first moved one step left by *transitioning out of a state* in depth  $k$  (or  $k - 1$ ).

As with  $T_{i,j}^L$ ,  $p_{i \rightarrow d}^L$  depends only on the structure of the “subchain” consisting of levels  $\{j, j + 1, \dots\}$ , including transition rates to level  $j - 1$ . Since for all  $j \geq k$  the resulting subchains are isomorphic,  $p_{i \rightarrow d}^L$  must be the same for all  $j \geq k$ .  $\square$

## B Solution of the system of equations for M/M/k/setup

The steps below illustrate how to solve the system of equations for M/M/k/setup. All of the operations in the steps below can be performed *symbolically* to obtain closed-form results.

### B.1 Solving for $p_{i \rightarrow d}^L$

The system of equations for  $p_{i \rightarrow d}^L$  consists of equation sets (26), (27) and (28). Eqs. (26) are  $k$  quadratic equations, each in one variable:  $p_{0 \rightarrow 0}^L, p_{1 \rightarrow 1}^L, \dots, p_{k-2 \rightarrow k-2}^L, p_{k-1 \rightarrow k-1}^L$ . Thus, we can solve each equation easily using the quadratic formula. It can be easily shown that among the two roots of each equation, the greater root exceeds 1, and is thus disregarded. The lesser root can be shown to lie in the interval  $[0, 1)$ , making it the unique solution of interest to the quadratic equation. Note that  $p_{0 \rightarrow 0}^L = 0$ , as expected (we cannot move to the left when we have no servers on).

The set of equations (27) is a collection of  $O(k^2)$  equations involving linear terms and products of two unlike variables. However, the structure of this system of equations reduces solving the system to solving a set of linear equations via back substitution. Consider solving this set of equations for the unknown values of  $p_{i \rightarrow d}^L$  in this order:

$$p_{k-1 \rightarrow k-1}^L, p_{k-2 \rightarrow k-2}^L, p_{k-2 \rightarrow k-1}^L, \dots, p_{0 \rightarrow 1}^L p_{0 \rightarrow 2}^L, \dots, p_{0 \rightarrow k-1}^L$$

That is, solving from greatest  $(k - 1)$  to least  $(0)$  “original depth,” but within each original depth, solving from least to greatest “target depth.” Solving in this order, each equation we solve will only have one unknown, as all other variables will already have been solved for in an earlier step (including the  $p_{i \rightarrow i}^L$  from Equations (26)), so these variables can be viewed as coefficients and constant terms. Once we have solved Equations (27), we can easily solve Equations (28), yielding  $p_{0 \rightarrow k}^L, p_{1 \rightarrow k}^L, \dots, p_{k-1 \rightarrow k}^L$ , by taking complements. It follows that all  $p_{i \rightarrow d}^L$  can be solved in closed forms that are, at worst, linear combinations of radicals (i.e., square roots).

## B.2 Solving for $\dot{R}_{i,k}^L$

The system of equations for  $\dot{R}_{i,k}^L$  consists of Equations (29) and (30), and Eq. (31). This system is a collection of  $(k+1)$  linear equations with  $(k+1)$  unknowns. Although we could solve this system using standard linear algebraic techniques, the structure of this system suggests an even simpler approach using back substitution. Solving for each  $\dot{R}_{i,k}^L$  only requires knowing the  $\dot{R}_{i,\ell}^L$  such that  $\ell \in \{i+1, \dots, k\}$ . Eq. (31) readily gives us  $\dot{R}_{k,k}^L$ . Thus, we can now solve for  $\dot{R}_{k-1,k}^L$ , then  $\dot{R}_{k-2,k}^L$ , and so on. In this way, each  $\dot{R}_{i,k}^L$  is found by solving a linear equation for one unknown variable.

## B.3 Solving for $\dot{R}_{i,j}^H$

The system of equations for  $\dot{R}_{i,j}^H$  consists of Equations (32), (33), (34) and (35), and Eq. (36). This system is a collection of  $O(k^2)$  dependent linear equations with just as many unknowns. Unlike the earlier systems of equations, there is no apparent structure we can exploit, so the system can be solved via standard linear algebraic techniques such as (symbolic) matrix inversion.