

Submitted to *Operations Research*

An Upper Bound on the M/M/k Queue With Deterministic Setup Times

(Authors' names blinded for peer review)

Authors are encouraged to submit new papers to INFORMS journals by means of a style file template, which includes the journal title. However, use of a template does not certify that the paper has been accepted for publication in the named journal. INFORMS journal templates are for the exclusive purpose of submitting to an INFORMS journal and are not intended to be a true representation of the article's final published form. Use of this template to distribute papers in print or online or to submit papers to another non-INFORM publication is prohibited.

Abstract. In many systems, servers do not turn on instantly; instead, a *setup time* must pass before a server can begin work. These “setup times” can wreak havoc on a system’s queueing; this is especially true in modern systems, where servers are regularly turned on and off as a way to reduce operating costs (energy, labor, CO_2 , etc.). To design modern systems which are both efficient *and* performant, we need to understand how setup times affect queues.

Unfortunately, despite successes in understanding setup in a single-server system, setup in a multiserver system remains poorly understood. To circumvent the main difficulty in analyzing multiserver setup, all existing results assume that setup times are memoryless, i.e. distributed Exponentially. However, in most practical settings, setup times are close to Deterministic, and the widely used Exponential-setup assumption leads to unrealistic model behavior and a dramatic underestimation of the true harm caused by setup times.

This paper provides a comprehensive characterization of the average waiting time in a multiserver system with *Deterministic* setup times, the M/M/k/Setup-Deterministic. In particular, we derive upper and lower bounds on the average waiting time in this system, and show these bounds are within a multiplicative constant of each other. These bounds are the first closed-form characterization of waiting time in any finite-server system with setup times. Further, we demonstrate how to combine our upper and lower bounds to derive a simple and accurate approximation for the average waiting time. These results are all made possible via a new technique for analyzing random time integrals that we named the Method of Intervening Stopping Times, or MIST.

Key words: queueing; multiserver systems; setup times; Deterministic setup times; exceptional first service

1. Introduction

1.1. What are setup times?

In many systems, servers do not turn on instantly; instead, a *setup time* must pass before a server can begin work (Allahverdi and Soroush 2008). For example, for applications hosted in the cloud, application

replicas must take time to boot up before they can begin fulfilling requests (Rzadca et al. 2020); for overwhelmed hospitals, traveling nurses must wait to have their credentials confirmed before they can begin helping patients (Tuttas 2013); for many businesses, workers must go through a long and expensive recruitment/onboarding process before they can begin serving customers (Behroozi et al. 2020). By thinking about this “initial delay before service” as an abstract *setup time*, we can learn how setup time affects all of these systems simultaneously (Allahverdi and Soroush 2008).

1.2. Why do setup times matter?

Setup times can have a significant impact on a system’s queuing behavior, especially in modern systems. For systems which keep their servers *on* all the time, clearly setup times do not affect their performance. However, in many modern systems, servers are regularly turned on and off. Because servers don’t turn on instantly, jobs in a system with setup times end up delayed compared to their no-setup counterparts. If setup times are long enough, this additional delay can be significant.

Nevertheless, many systems still regularly turn their servers on and off. Why? Because, by doing so, they can save a considerable amount on operating costs (e.g. energy, money, CO_2 , etc.) (Rzadca et al. 2020). However, this cost-saving measure is only a viable option if the additional delay caused by setup times is not too large. Therefore, if we want to design systems which are simultaneously efficient *and* performant, we need a good understanding of how setup times affect queuing performance.

1.3. Prior Art on Understanding the Effect of Setup Times

State of the art. Unfortunately, despite continued academic interest, we still struggle to understand the impact of setup times on customer wait times outside of a few very simple settings. In the single server setting, (Welch 1964) completely characterized the behavior of the waiting time under extremely general conditions, i.e. in the $M/G/1$ queue with generally-distributed setup times. In the more complex multiserver setting, little progress was made until the publication of (Gandhi et al. 2013) almost half a century later. Outside of the fact that the model of (Gandhi et al. 2013) has multiple servers, their model is much simpler than that of (Welch 1964). In particular, the main results of (Gandhi et al. 2013) rely on the assumption that both job service times and server setup times are distributed Exponentially, i.e. the authors study the $M/M/k/Setup$ -Exponential model.

Limitations of the Exponential model. Despite the fact that (Gandhi et al. 2013) represented the first breakthrough in our understanding of the setup effect in 50 years, their results are limited in two significant ways. First, instead of a closed-form formula for the average waiting time, the authors only derive an algorithm for computing the average waiting time. This algorithm is useful in the sense that it bypasses the need to simulate the system, but unfortunately fails to give intuition about how wait times scale with system parameters. Second, one of their simplifying assumptions, that setup times are distributed Exponentially,




Figure 1 Simulation results for the M/M/k/Setup-Deterministic, M/M/k/Setup-Exponential, M/M/k (no setup), with mean service time $\frac{1}{\mu} = 1$ ms, mean setup time $\beta = 1000$ ms, and load kept at a constant $\rho = 0.5$. Note the high separation between the Exponential and Deterministic models at large scales.

turns out to severely limit the utility of their results. We discuss this point in depth within Section 1.4, but, put briefly, this lack of utility stems from the fact that setup times have relatively low variance (Mao and Humphrey 2012), i.e. they are closer to *Deterministic* than Exponential. Accordingly, we analyze the Deterministic version of the (Gandhi et al. 2013) model.

Previous work on the Deterministic model. At the moment, only one other result characterizing the average waiting time in the M/M/k/Setup-Deterministic exists: (Williams et al. 2022). The paper (Williams et al. 2022) proves the first-ever lower bound on the average waiting time in the M/M/k/Setup-Deterministic queue. (Williams et al. 2022) differs from this paper in a few ways. First, (Williams et al. 2022) proves *only* a lower bound on the average waiting time, whereas we prove the first-ever upper bound as well as a nearly matching lower bound. Furthermore, our lower bound is significantly stronger. To be more precise, (Williams et al. 2022) incorporates additional assumptions to prove their result, and the final result there becomes trivial in certain scaling regimes. In contrast, our proofs only assume 1) that setup times are significantly larger than service times and 2) that enough servers are being regularly utilized, and we prove explicitly that our lower bound and upper bound are separated by at most a multiplicative constant.

1.4. Limitations of the Exponential model

The Exponential model is unrealistic. As previously mentioned, the Exponential assumption of (Gandhi et al. 2013) turns out to be extremely problematic; we highlight two particular issues. First, this assumption leads to unrealistic model behavior. To illustrate where the breakdown in realism happens, consider a scenario where only a single server is setting up and compare it to a scenario where 100 servers begin setup at

the same time. In the Deterministic model, all 100 servers perform no work while they are setting up, then all of them turn on simultaneously. By contrast, in the Exponential model, the 100-server system receives its first server on average 100 times faster than the single-server system receives its first server. In other words, in the Exponential model, the *longer* the system's queue is, the *more rapidly* the system's servers turn on to help *drain* that queue. In a sense, the Exponential system can rapidly “react” to increases in queue length. **Note that this “blessing of variability” effect only comes into play when there are multiple servers; in the single server case, as seen in Figure 1, the variability in setup time hurts the waiting time.**

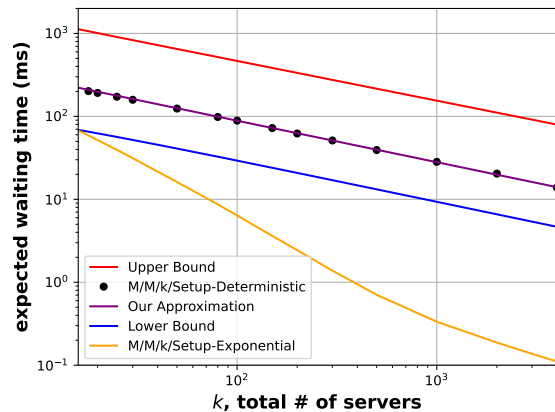
The Exponential model underestimates waiting. This unrealistic “reactivity” phenomenon causes a further, more concerning, problem: in real systems, the Exponential model dramatically underestimates how much waiting actually occurs. To be more precise, in modern systems: 1) average setup times are often larger than average job sizes by two or three orders of magnitude (Gandhi et al. 2012, Mao and Humphrey 2012, Mogul and Kompella 2015, Maheshwari et al. 2018, Hao et al. 2021); and 2), as noted in the state-of-the-art paper (Gandhi et al. 2013), setup times are actually closer to *Deterministic*. When these two criteria are satisfied, as observed in Figure 1, the true waiting time is often orders of magnitude larger than what the Exponential model predicts. Accordingly, in many practical studies of the setup effect (Gandhi et al. 2012, Kara 2017, Hyytiä et al. 2018), setup times are assumed to be Deterministic, e.g. servers take a fixed time of 2 minutes to set up. However, despite its apparent practical limitations, the Exponential setup model remains the *de facto* choice for theoretical analysis, since it allows for the application of a number of existing theoretical techniques.

1.5. Our Results


In this paper, for the first time, we provide the first characterization of the average waiting time in a dynamically-provisioned multiserver system with Deterministic setup times; we plot our results in Figure 2a. To be more specific, we study the Deterministic analogue of the model in (Gandhi et al. 2013), the *M/M/k/Setup-Deterministic*, and derive three main results. Our first result is Approximation 1, a prediction of the average waiting time which appears extremely accurate across a wide variety of parameter settings (see Figure 4 for a detailed visualization). When simplified, this closed-form approximation roughly states that

$$\mathbb{E}[T_Q] \approx \frac{1}{2} \sqrt{\frac{\pi}{2}} \frac{\beta}{\sqrt{k\rho}} + \frac{1}{\mu k(1-\rho)}.$$

In support of this approximation, in our second and third results, we prove upper and lower bounds on the average waiting time; these make up Theorems 1 and 2, respectively. Furthermore, we later show, in Theorem 3, that our results capture the correct order-wise scaling of the average waiting time—they agree, up to multiplication by an absolute constant. Note that, at the moment, no other such characterizations exist for any similar multiserver setup system; not even for the extensively-studied *M/M/k/Setup-Exponential* model.



(a) An illustration of our results.



(b) A provisioning example.

Figure 2 Our theoretical results along with simulation data for the $M/M/k/Setup$ -Deterministic and $M/M/k/Setup$ -Exponential, varying the number of servers k while keeping the mean service time $\frac{1}{\mu} = 1$ ms, the mean setup time $\beta = 1000$ ms, and the load $\rho = 0.5$ fixed. (a) A comparison of our results to the true average waiting time in the $M/M/k/Setup$ -Deterministic. Our results behave like the true average waiting time, while the Exponential model behaves differently. (b) A provisioning example highlighting the differences between the Deterministic and Exponential models. To achieve a target waiting time of 20 ms, our approximation correctly predicts it will take $k \approx 2000$ servers, while the Exponential model predicts that only $k \approx 50$ servers should suffice. See Figure 4 for a more comprehensive evaluation.

1.6. Impact: How understanding setup times helps with capacity provisioning

A common but complex problem which arises in many areas is that of designing a system such that the average waiting time of a customer is below some target waiting time. Historically, we understand this problem well for systems without setup times, e.g. there's a straightforward formula for the average waiting time in the $M/M/k$ without setup. Unfortunately, our understanding of this problem is quite poor for more modern systems, since their average waiting times are affected by setup times. In particular, modern systems dynamically control the number of servers that they keep on, periodically turning servers off in order to save energy. As we mentioned before, previous results on understanding the relationship between setup times and the average waiting time leave much to be desired. Our new results expand on the state-of-the-art Exponential model in two important ways: 1) obtaining the predicted average waiting time is much easier computationally, and 2) the quality of the prediction is much better.

Easier predictions. Compared to the Exponential model, our new Deterministic approximation greatly simplifies the design process. In particular, when predicting the average waiting time in the Exponential model using the state-of-the-art method from (Gandhi et al. 2013), one must solve a system of $O(k^2)$ quadratic equations to find the average waiting time $\mathbb{E}[T_Q]$. Two practical issues arise from this fact. First, the equations change depending on the number of servers k , meaning that the computation must be repeated every time one wishes to test a new number of servers. Second, the opacity of the process makes it difficult

to get intuition about how the average waiting time changes as one alters the system parameters. In contrast, Approximation 1 is a relatively simple function of the relevant parameters. The simplicity of our approximation has, likewise, two benefits: 1) computing the waiting time becomes easy, and 2) our approximation's form makes it clear how and why the waiting time behaves the way it does.

Higher quality predictions. Moreover, when compared to the predictions of the Exponential model, the predictions we obtain using our Deterministic approximation are of a much higher quality. This difference in quality is perhaps best illustrated by looking at a simple example. In Figure 2b, we compare the prediction from the Exponential model to the prediction from our approximation, plotting how the predicted average waiting time changes as one increases the number of servers k while fixing the load $\rho = 0.5$, the average setup time $\beta = 1000$ ms, and the average service time $\frac{1}{\mu} = 1$ ms (note that these are typical relative values in many applications (Mao and Humphrey 2012)). Our goal is to determine how large the number of servers k needs to be before we reach our target waiting time $T^{\text{target}} = 20$ ms. In both models, the average waiting time decreases as the system gets larger. However, the Exponential model predicts that the average waiting time will be small enough once $k = 50$. On the other hand, as captured by our approximation, the Deterministic setup system will only reach the target waiting time once the number of servers $k \approx 2000$ —a full 40 times larger than what the Exponential system predicts! At even a modest number of servers, the Exponential system underestimates the waiting time by orders of magnitude.

1.7. Challenges and Our Approach

What makes the multiserver setting difficult. Although the harm caused by setup times is non-trivial to understand even in single server systems, the setup effect can be especially difficult to understand when multiple servers can set up at the same time. In particular, when servers can set up *simultaneously*, their server states begin to interact; via the speed of their processing, the system's *busy* servers indirectly control the setup behavior of the system's *not-busy* servers. For example, if server A is *on* while server B is *setting up*, then server A might finish all the work in the queue before server B even has a chance to turn *on*. As such, in the multiserver setting, it can sometimes make sense to *cancel* a server's setup process; a situation which would *never* occur in the single server setting. Of course, the opposite can also happen: if the busy servers are working much more slowly than expected, then the queue might grow large enough that we begin *setting up* a server that would otherwise be left *off*. This interaction between departure behavior and setup behavior is exactly what makes the setup effect so much harder to understand in the multiserver setting.

Why Deterministic setup is especially difficult. As such, while using Deterministic setup times might be more realistic, it also comes with a set of unique theoretical challenges. In the Deterministic case, there is no avoiding the complexity of setup: even in simulation, one must track the individual remaining setup time of *every* server that is currently setting up. By contrast, because the Exponential distribution is *memoryless*, in the Exponential case it suffices to track only the *total number* of servers setting up, greatly simplifying

the system state. Moreover, the Exponential model's simple state forms a Continuous-Time Markov Chain, a well-studied class of stochastic processes for which a number of techniques have been developed. For the Deterministic setup model, no such techniques exist.

Our approach. To derive our first-of-their-kind results, we develop a new technique called the Method of Intervening Stopping Times, or MIST, which allows us to more naturally investigate queueing systems with high-dimensional, continuous-valued state spaces. Generally, MIST is a method for bounding the expectation of a random time integral, and it works by dividing up the original time interval into smaller, more manageable pieces using stopping times. By focusing on these smaller pieces, we avoid the issue of tracking the entire high-dimensional system state at every point in time. Instead, MIST allows us to narrow our focus to only the “most important” aspects of the state for a particular integral, and to shift our definition of “most important” based on the particular integral we are analyzing. Since, for a particular integral, we know the most important aspects of the system state, we can derive strong bounds on that integral that hold *regardless* of the remaining state information. Moreover, we can apply this approach recursively, chopping a still-large time interval into a sequence of even smaller time intervals. By chopping time in the right way, eventually we can make the system so well-behaved during an interval that we can adapt powerful techniques like martingale theory and Wald's equation to the analysis of a phase. **To place our approach in a broader context, MIST could be considered a fleshed-out instantiation of a suggestion made by Kingman in Kingman (2009) that “...it may be possible to sew these small martingales, or constant multiples of them, together to form a martingale on the whole line, and then all the apparatus of martingale theory—stopping identities, inequalities, central limit theorems, and the like—becomes available;” for more details on our approach, see Section 5.**

1.8. Potential Mitigation

The delay in serving jobs due to setup times comes from turning servers off and then needing to wait before those servers can turn on. One way to potentially mitigate the impact of setup times is then to turn off servers less aggressively. We analyze such an approach in Appendix H, and our analysis reveals that a natural class of policies have minimal impact on the waiting time. This preliminary analysis highlights the need for further studies on designing policies for turning servers on and off in queueing systems with setup times to achieve a better tradeoff between energy efficiency and performance.

1.9. Outline

We proceed in the following order: In Section 2, we give a detailed description of our model. In Section 3, we discuss the previous work on analyzing setup times in greater detail. In Section 4, we fully state our main results. In Section 5, we discuss the key ideas which enabled us to prove our main results, including




Figure 3 An example of M/M/k/Setup-Deterministic with $k = 4$. The state pictured has $Z(t) = 2$ busy servers, which means there are 2 jobs in service. There is $Q(t) = 1$ job in queue, and thus $N(t) = Z(t) + Q(t) = 4$ jobs in system.

MIST. In Section 6, we give the proof of Theorem 1; we relegate the proof of Theorem 2 to Appendix B. Finally, in Section 7 we make some closing remarks and propose follow-up work based on our findings.

2. Model: The M/M/k/Setup-Deterministic Queue

The system behavior, excluding setup. We now describe our model of interest, the M/M/k/Setup-Deterministic. As in the typical M/M/k queue, jobs arrive in a Poisson process of rate $k\lambda$ into a FCFS queue where jobs wait to be served at one of k servers. The job at the head of queue enters service whenever a server frees up, either from a job completing service or from a server finishing set up. Once a job enters service, it remains in service for $\text{Exp}(\mu)$ time before departing. We assume all the servers have identical service and setup **distributions**. As such, we can assign each server an index from 1 to k , and without loss of generality assume that departures always occur at the busy server with the highest index; i.e., we re-index the servers when a job departs so the server with the newly departed job has the highest index among the busy servers. From here, we define the quantity $Z(t)$ to be the number of busy servers (or jobs in service) at time t , the quantity $Q(t)$ to be the number of jobs waiting in the queue at time t , and the quantity $N(t) = Q(t) + Z(t)$ to be the total number of jobs in our system. Excluding the setup dynamics, one sees that, as promised, the behavior of our model is identical to the M/M/k queue.

The setup dynamics. From here, it suffices to describe precisely how servers will be turned *on* and *off*. We assume that each server is always in one of three states: *on*, *off*, or *in setup*. A given server remains *on* only as long as that server remains busy. In other words, a server turns *off* when it finishes its current job and the queue is empty. Recall that we re-index the servers to ensure that the job departure is from the highest indexed server among the busy servers. On the other hand, server i begins *setup* when a job arrives to the system and there are only $i - 1$ jobs in the system. Server i remains in setup until one of two events occurs: either 1) some fixed quantity β time has passed, or 2) there are fewer than i jobs in the system; accordingly, we refer to β as the *setup time* of a server. In the first case, if β time has passed without $N(t)$ dipping below i , then server i has completed its setup and begins working on the job at the head of the queue. In

the second case, if the number of jobs $N(t)$ dips below i before server i completes setup, then the setup is canceled and server i turns *off*. We use $Y_i(t)$ to denote the detailed state of server i at time t . If server i is *off*, we set $Y_i(t) = \text{OFF}$; if server i is *on*, we set $Y_i(t) = \text{ON}$; if server i is *in setup*, we let $Y_i(t)$ denote the remaining amount of time until server i would finish setup, if left uninterrupted. To be precise, $Y_i(t)$ is set to β when server i first initiates *setup*, and this value decreases at rate 1 until either setup completes or setup is canceled. For convenience, we assume, without loss of generality, that $\text{ON} < s < \text{OFF}$ for every possible remaining setup time $s \in (0, \beta]$; this ensures that the detailed state $Y_i(t)$ is non-decreasing in i . As a shorthand, we use $\mathbf{Y} = (Y_1(t), Y_2(t), \dots, Y_k(t))$ to denote the vector of detailed server states.

A state descriptor. Accordingly, a Markovian state descriptor for our system at time t is $S(t) \triangleq (N(t), \mathbf{Y}(t))$. Note that, since one can recover the number of jobs in service $Z(t)$ from the detailed server states $\mathbf{Y}(t)$, one could also choose the state to be $(Q(t), \mathbf{Y}(t))$. Either suffices in providing a complete description of the forward dynamics of the system. Furthermore, when discussing the steady-state distribution of, say, the number of jobs $N(t)$, we use the notation $N(\infty)$.

Some important constants. We define some system parameters which are critical to system behavior. We use $\rho \triangleq \frac{\lambda}{\mu}$ to refer to the load of our system, i.e., the time-average fraction of servers working on a job. We call the offered load $R \triangleq k\rho$; this is the time-average *number* of busy servers. To enforce stability, we require that $\rho < 1$. As discussed, the symbol β refers to the fixed (Deterministic) setup time of a server.

Busy period notation. Our results can be stated more concisely with two quantities related to a busy period of an M/M/1 queue. We give the notation below. We use $T^{\text{busy}}(n, j)$ to denote the expectation of the random *length* of an M/M/1 busy period with arrival rate $k\lambda$, service rate $k\lambda + \mu j$, and which starts with n jobs in the system. Likewise, we use $I^{\text{busy}}(n, j)$ to denote expectation of the random *time integral of the number of jobs* within the M/M/1 over the same period. Explicitly, we have

$$T^{\text{busy}}(n, j) \triangleq \frac{n}{\mu j} \tag{1}$$

and

$$I^{\text{busy}}(n, j) \triangleq \frac{n}{\mu j} \left[\frac{n+1}{2} + \frac{1}{1 - \frac{k\lambda}{k\lambda + \mu j}} \right] = \frac{n}{\mu j} \left[\frac{n+1}{2} + \frac{R}{j} + 1 \right]. \tag{2}$$

3. Related Work

In Section 1, we discussed two works most directly related to ours, (Gandhi et al. 2013) and (Welch 1964). Here, we briefly describe a few other works studying setup times in queueing systems.

3.1. Single Server Setup

The single server setting. Since the seminal paper of (Welch 1964) on the M/G/1/Setup, various followup work has been devoted to extending this work to service disciplines beyond First-Come-First-Served (Bischof 2001) and arrival processes beyond Poisson (He and Jewkes 1995, Choudhury 1998).

Staggered setup. The first investigations into multiserver setup study the $M/M/k/Setup$ and the $M/G/k/Setup$ with the additional assumption that **only one server can be in setup at a time** (Artalejo et al. 2005, Gandhi and Harchol-Balter 2013a,b). However, in real systems, **servers can undergo the setup process simultaneously** (e.g. **ten virtual machines can be booting up at the same time**), meaning that the strong results obtained for the staggered setup model are of limited practical use.

3.2. Multiserver Setup

M/M/k/Setup-Exponential: Approximations. In (Gandhi et al. 2010b), a precursor to (Gandhi et al. 2013), the authors derive a number of intuitive approximations for the average waiting time. Further, (Pender and Phung-Duc 2016) analyzes an extended version of their model which includes customer abandonment and time-varying arrival rates, providing an approximation for the average queue length at each point in time. Unfortunately, unlike our results, none of these results are in closed-form, nor do they provide a bound on the approximation error.

M/M/k/Setup-Exponential: Exact analysis. In (Gandhi et al. 2013), the authors provide a method for computing the expectation of any function of the steady-state queue length; their method was later shown to be analogous to the matrix analytic method (Phung-Duc 2017). As usual, though, these results are not in closed-form; even now, it remains unknown how the average waiting time in the $M/M/k/Setup-Exponential$ model varies as one varies the system parameters.

Dispatching/Load-balancing. In (Mukherjee et al. 2017), the authors consider a dispatched, finite-buffer version of the $M/M/k/Setup-Exponential$, for which they design TABS, an asymptotically-optimal (in energy waste, as $k \rightarrow \infty$) load-balancing/scaling scheme. In a followup paper, (Mukherjee and Stolyar 2019), the authors extend their work to the infinite buffer case, confirming its asymptotic optimality.

Deterministic setup. Besides the lower bound of (Williams et al. 2022), a few others have investigated multiserver systems with Deterministic setup times. In the control setting, (Hyytiä et al. 2018) considers a dispatching version of the $M/G/2/Setup-Deterministic$ model, building near-optimal policies for the joint control of setup and dispatching. In (Kara 2017), the author provides a simulation-based analysis of the $M/M/k/Setup-Deterministic$ queue, corroborating our findings on the large gap between Exponential and Deterministic setup systems.

Multiserver scheduling. In (Hong and Scully 2023), the authors study the performance of the Gittins policy in the $G/G/k/Setup$ queue and, instead of explicitly bounding the average wait as we do, they bound the difference between Gittins' average waiting time and the optimal average waiting time.

Policies for mitigating the harm caused by setup. Many works have considered the problem of mitigating the energy waste and additional delay caused by setup times, often formalized in terms of the Energy-Response-time-Product (ERP) (Gandhi et al. 2010a), the Normalized-Performance-Per-Watt (NPPW) (Gandhi and Harchol-Balter 2011), or the energy expended given a fixed tail cutoff for response

time (Gandhi et al. 2012). Many works study the *DelayedOff* policy, where a server must idle for a period before it may turn off (Gandhi et al. 2010a,b, 2012, Pender and Phung-Duc 2016). When using *DelayedOff*, the choice of which *idle-but-on* server receives a job matters: (Gandhi et al. 2010a) considers routing the job to the *Most Recently Busy* server; (Gandhi et al. 2012) creates a ranking of all servers, always routing to the lowest-ranked idle server. Other works (Gandhi et al. 2010a, 2011) consider utilizing sleep states, useful intermediates between the “no power, long setup” *off* state and the “full power, no setup” *on* state. **We comment that the setting considered in this paper corresponds to a natural policy. But it is unclear whether this natural policy, or any other policy, achieves an optimal tradeoff between the energy consumption and the additional delay caused by setup times.**

4. Results

Summary. We now state our three main results concerning the average queue length $\mathbb{E}[Q(\infty)]$: an approximation, an upper bound, and a lower bound. Note that, by Little’s Law (Kleinrock 1975), which states that $\mathbb{E}[Q(\infty)] = k\lambda\mathbb{E}[T_Q]$, our results immediately translate to results on the average waiting time $\mathbb{E}[T_Q]$; see Figure 2a for an illustration. Our first main result, Approximation 1, is a simple and accurate approximation to the average queue length $\mathbb{E}[Q(\infty)]$; we demonstrate its accuracy in Figure 4. Our second main result, the upper bound of Theorem 1, is the first-ever upper bound on the average queue length $\mathbb{E}[Q(\infty)]$ in a multiserver system with Deterministic setup times. By contrast, our third main result, the lower bound of Theorem 2, is only the second lower bound on the queue length $\mathbb{E}[Q(\infty)]$; however, improving on the lower bound of (Williams et al. 2022), we show that our bounds differ by at most a multiplicative constant. Notably, all three results are in explicit closed-form, unlikely nearly all other results concerning multiserver systems with setup times.

We begin by discussing Approximation 1 in Section 4.1. We then discuss the assumptions used in the proof of our main theorems in Section 4.2. Afterwards, we state the exact forms of the upper and lower bounds in Theorems 1 and 2 in Sections 4.3 and 4.4, respectively, whose proofs are given in Section 6 and Appendix B, respectively. Finally, we briefly discuss the order of these bounds in Section 4.5.

4.1. Approximation: Statement and Evaluation

Approximation 1 (Approximation of the Average Queue Length) *Let $C_{\text{apx}} \triangleq \sqrt{\frac{\pi}{2}}$. Then, if the offered load $R \triangleq k\rho > 1$, one can approximate the steady state queue length as*

$$\mathbb{E}[Q(\infty)] \approx Q_{\text{apx}} \triangleq \frac{\frac{1}{2}\mu\beta^2 C_{\text{apx}} \sqrt{R} + \frac{\mu\beta C_{\text{apx}} \sqrt{R}}{\mu k(1-\rho)} \left[\frac{\mu\beta C_{\text{apx}} \sqrt{R} + 1}{2} + \frac{1}{1-\rho} \right]}{\beta + \frac{\mu\beta C_{\text{apx}} \sqrt{R}}{\mu k(1-\rho)}}. \quad (3)$$

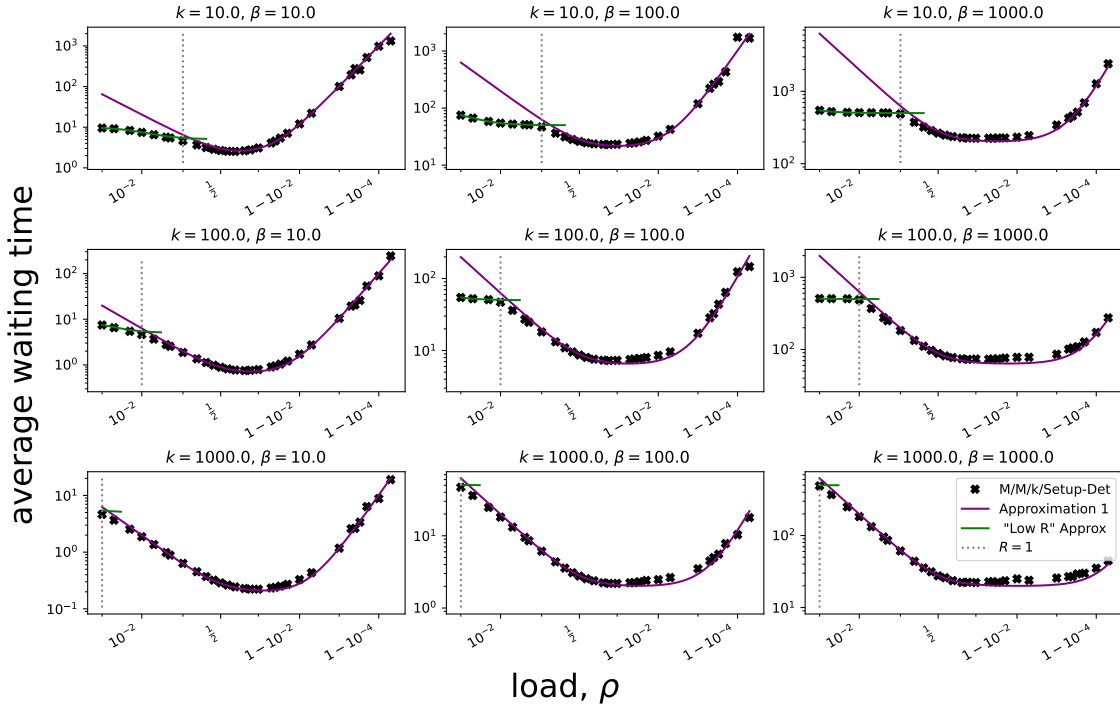


Figure 4 Simulation results demonstrating the high accuracy of Approximation 1. For each of these 9 plots, we plot the behavior of the average waiting time as one varies the load ρ from 0 to 1, holding fixed the total number of servers k as well as the setup time β . In each row, we hold the number of servers k constant while testing increasing values of the setup times β . In each column, we hold the setup time β constant while increasing the number of servers. We plot three quantities: 1) in black, the simulated average waiting time for the M/M/k/Setup-Deterministic; 2) in purple, the predicted average waiting time of Approximation 1; and 3) the predicted average waiting time of as given by the “low R” approximation of (4), a variation on the single-server setup result of Welch (1964). We also include, as a reference, a dotted line illustrating the point at which the offered load $R \triangleq k\rho = 1$. Our approximation works well when the average number of busy servers $R > 1$, and the “low R” approximation works well when $R < 1$.

Evaluation. In Figure 4, we observe that our approximation is accurate across a wide range of system parameters, so long as the offered load $R > 1$, a considerable weakening of the assumptions used in our proofs. In that sense, it serves as a testament to the strength of our approach that our resulting approximation remains accurate all the way down to offered loads $R \approx 1$.

When $R < 1$, we find that the waiting time exhibits a “single-server bottleneck” effect, mirroring the behavior in (Welch 1964). In particular, in the M/M/1/Setup-Deterministic, one finds in (Welch 1964) that

$$\mathbb{E}[T_Q] = \frac{1}{\mu - \lambda} + \frac{\beta}{2} \left[\frac{2 + \lambda\beta}{1 + \lambda\beta} \right].$$

Taking the second term, which measures the additional wait in the system due to setup times, and using R in place of $\frac{\lambda}{\mu}$, we obtain a “low R ” approximation

$$\mathbb{E}[T_Q] \approx \frac{\beta}{2} \left[\frac{2 + \mu R\beta}{1 + \mu R\beta} \right]. \quad (4)$$

We include this “low R” approximation in Figure 4 and note that, when $R < 1$, this “low R” approximation seems to capture the behavior of the waiting time quite well.

4.2. Assumptions.

In our theoretical analysis, in order to more simply characterize the system’s behavior, we make two assumptions. First, we assume that setup times are large compared to service time, i.e. the average setup time $\beta \geq 100\frac{1}{\mu}$; this is often satisfied in practice (Gandhi et al. 2012, Mao and Humphrey 2012, Mogul and Kompella 2015, Maheshwari et al. 2018, Hao et al. 2021). Second, we assume that the system utilizes, on average, at least 100 servers, i.e. the offered load $R \triangleq k\rho \geq 100$.

Note that the specific values within these assumptions are not strictly necessary for our analysis to go through; these assumptions are made predominantly to simplify the expressions which arise in our analysis. For example, in our final bound on the quantity $\mathbb{E}[L]$ at the end of Appendix G.6, we are left with the expression

$$\mathbb{E}[L] \geq \left(1 - \frac{b_1}{\sqrt{\mu\beta R}}\right) \left[\left(1 - \frac{2}{\sqrt{R}}\right) \left(\sqrt{\frac{\pi}{2}} - \frac{1.15}{\sqrt{\mu\beta}} - 2e^{-4}\right) - \frac{1}{2\sqrt{R}}\right] \sqrt{R}.$$

By imposing the above constraints, we are able to reduce the above expression to

$$\mathbb{E}[L] \geq \frac{2}{3} \sqrt{\frac{\pi}{2}} \sqrt{R},$$

which still captures the correct order-wise scaling (and even points towards the correct asymptotic behavior) while being much easier to read.

4.3. Upper Bound: Statement

THEOREM 1 (Upper Bound on Average Queue Length). *For an M/M/k/Setup-Deterministic with an offered load $R \triangleq k\rho \geq 100$ and a setup time $\beta \geq 100\frac{1}{\mu}$, the expected number of jobs in queue in steady state is upper-bounded as*

$$\mathbb{E}[Q(\infty)] \leq 3.6\sqrt{\mu\beta R} + 2.04\frac{\rho}{1-\rho} + \frac{4.05\mu\beta^2\sqrt{R} + g\left(9(\mu\beta)^2R, 3\mu\beta\sqrt{R}, k(1-\rho)\right)}{\beta + \frac{L_1\mu\beta\sqrt{R}}{\mu k(1-\rho)}},$$

where the function $g(x, y, z) \triangleq x\frac{1}{2\mu z} + y\left[\frac{R}{\mu z^2} + \frac{3}{2\mu z}\right]$ and the constant $L_1 = \frac{2}{3}\sqrt{\frac{\pi}{2}}$.

4.4. Lower Bound: Statement

THEOREM 2 (Improved Lower Bound on Average Queue Length). *For an M/M/k/Setup-Deterministic with an offered load $R \triangleq k\rho \geq 100$ and a setup time $\beta \geq 100\frac{1}{\mu}$, the expected number of jobs in queue in steady state is lower-bounded as*

$$\mathbb{E}[Q(\infty)] \geq \frac{L_1\beta^2\sqrt{R} + I^{\text{busy}}\left(\left[L_1\beta\sqrt{R} - (k-R)\right]^+, k-R\right)}{2.08\beta + \frac{1}{\mu}\frac{F_1\beta\sqrt{R}}{k-R} + \frac{1}{\mu}\frac{3}{2}\ln(\beta) + \frac{1}{\mu}\ln(F_1D_1) + \frac{2}{\mu} + \frac{1}{\mu}\left[D_2 + \frac{D_3}{\sqrt{R}}\right] \max\left(\frac{1}{D_1\sqrt{\mu\beta}}, \frac{1}{\sqrt{R}}\right)},$$

where $L_1, F_1, D_1, D_2,$ and D_3 are constants independent of system parameters.

4.5. Discussion of Bounds

After simplification, our bounds become

$$\mathbb{E}[Q(\infty)] =_c \mu\beta\sqrt{R} + \frac{1}{1-\rho};$$

where the operator $=_c$ denotes equality up to multiplicative constants; we show this explicitly in Section C.

This simplified characterization of the queue length gives us insight into how the system parameters govern the system's queuing behavior. The first term, $\mu\beta\sqrt{R}$, scales linearly with the setup time β and captures the effect of turning servers *off* and *on*. The second term, $\frac{1}{1-\rho}$, dominates the first term only when the load ρ is high enough, say in the renowned super-Halfin-Whitt regime (Halfin and Whitt 1981, Liu and Ying 2022) where $\rho = 1 - \gamma k^{-\alpha}$ with $0 < \gamma < 1$ and $\alpha > 0.5$. In this case, it recovers the $\frac{1}{1-\rho}$ scaling seen in the M/M/k without setup times.

5. Key Ideas and Techniques

We now describe our approach to analyzing the average waiting time in the M/M/k/Setup-Deterministic, using the upper bound, Theorem 1, as a case study. To begin, we go through the first few steps in our proof, arriving at three main lemmas. We then give a detailed explanation of the technique we have developed for proving those lemmas, which we call the *Method of Intervening Stopping Times* (MIST). Throughout, we highlight the technical challenges that arise in our approach and how we address those challenges.

5.1. Initial Steps: Applying the Renewal-Reward Theorem

5.1.1. Reduction to analyzing $\mathbb{E}[N(\infty) - R]$. We begin by applying the Renewal Reward theorem. Although it is tempting to apply the theorem directly to the queue length $\mathbb{E}[Q(\infty)]$, it simplifies the analysis considerably if one analyzes the number of jobs $\mathbb{E}[N(\infty) - R]$ instead. To justify this, note that, in steady-state, the number of busy servers $\mathbb{E}[Z(\infty)] = R$, and that the total number of jobs in system $N(t)$ satisfies $N(t) = Q(t) + Z(t)$. It follows that the average queue length

$$\mathbb{E}[Q(\infty)] = \mathbb{E}[N(\infty) - Z(\infty)] = \mathbb{E}[N(\infty)] - R = \mathbb{E}[N(\infty) - R].$$

5.1.2. Applying Renewal-Reward. Applying the Renewal Reward Theorem, for any renewal cycle,

$$\mathbb{E}[N(\infty) - R] = \frac{\mathbb{E}\left[\int_{\text{cycle}} [N(t) - R] dt\right]}{\mathbb{E}[\text{cycle length}]} = \frac{\mathbb{E}\left[\int_0^X [N(t) - R] dt\right]}{\mathbb{E}\left[\int_0^X 1 dt\right]}, \quad (5)$$

where we have set time 0 to be an arbitrary renewal point and time X to be the renewal point which immediately follows it. To upper bound (5), it suffices to upper-bound the right side's numerator and lower-bound the right side's denominator. These bounds constitute our three main lemmas; two lemmas for the numerator (which is harder to bound) and one for the denominator. But before we can state these main lemmas, we must first address our first technical challenge: defining the collection of renewal points \mathcal{X} .

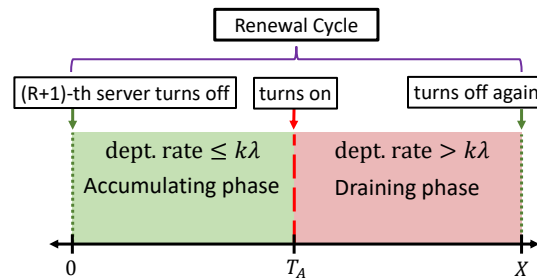


Figure 5 A depiction of our decomposition of a renewal cycle into an accumulating phase and draining phase, described in Section 5.1. During the accumulating phase, the departure rate $\mu Z(t) \leq \mu R = k\lambda$, so that the system is transiently unstable and a queue *accumulates*. During the draining phase, the departure rate $\mu Z(t) > k\lambda$, so that the queue *drains*.

5.1.3. Technical Challenge: Defining the renewal points. To define a “good” collection of renewal points \mathcal{X} is a non-trivial task. There are two important criteria one must consider when selecting the collection of renewal points. First, we need the length of a renewal cycle to have finite expectation. While this is trivially satisfied in many queueing models, in our model, the presence of remaining setup times in our system state makes this criterion non-trivial. **Second, we want the renewal period to be defined in such way that one can straightforwardly reason about the system’s behavior in between renewals.** For example, we could define our renewals as occurring when the system completely empties; this would avoid the continuous-state-space concerns, but then one would need to reason about what happens to the system until the next time it empties —by analogy to the M/M/k without setup, a typical renewal cycle would last exponentially long (in k). Thus, while the “system empties” renewal definition does indeed give a simple analysis for $k = 1$, for even $k = 2$ the analysis becomes much more difficult. We must find another way.

5.1.4. Key Idea: Renewals when the $(R + 1)$ -th server turns off. By defining our collection of renewal points $\mathcal{X} = \{t : Z(t^-) = R + 1, Z(t) = R\}$ as those moments when the $(R + 1)$ -th server is turned off, we satisfy both criteria. By defining our renewals to occur at moments when a server is turned *off*, we avoid the continuous state-space issue: any time a server shuts *off*, we know that there must be no servers in *setup* (since we would cancel any servers in *setup* before canceling an idling server). Furthermore, by specifically analyzing around the moment the $(R + 1)$ -th server turns off, the length of time between renewals becomes well-controlled, in the sense that it is lower-bounded almost surely by a setup time and upper-bounded in expectation by a quantity which we can show is “not too large.”

5.1.5. Further Benefit: A natural decomposition. Another benefit of defining a renewal cycle in this way is that it naturally splits a renewal cycle into two parts: the first part of the cycle *before* the $(R + 1)$ -th server turns on, and the second part of the cycle *after* it turns on; we depict this decomposition in Figure 5. During the beginning of a cycle (when $Z(t) \leq R$), the departure rate $\mu Z(t)$ is smaller than the arrival rate, making the system behave, in a transient sense, like a critically- or over-loaded queue. On the other hand,

after the $(R + 1)$ -th server turns on and until the renewal cycle is over, the departure rate $\mu Z(t)$ is guaranteed to be strictly greater than the arrival rate, making the queue, on the whole, drain over time. This observation turns out to be hugely useful in our analysis. As such, we have special names for each of these special times: we call the time before the $(R + 1)$ -th server turns on the *accumulation period*, we call the moment when the $(R + 1)$ -th server turns on the *accumulation time* T_A , and we call the period from time T_A until the cycle ends the *draining period*. At times, we will also refer to these periods as *phases*.

5.2. Statement of Three Main Lemmas

We may now state our three main lemmas. (Note that the preconditions of Theorem 1 are implicit here.)

LEMMA 1 (Upper Bound on Integral Over Accumulation Period). *Suppose the system begins at time 0 with R jobs in service and no jobs in the queue (and thus no servers in setup), and define the accumulation time $T_A \triangleq \min \{t \geq 0 : Z(t) = R + 1\}$ to be the moment the $(R + 1)$ -th server turns on. Then, taking $B_1 = 3.6$ and $B_2 = 1.04$, one can bound the integral of the reduced number of jobs $[N(t) - R]$ by*

$$\mathbb{E} \left[\int_0^{T_A} [N(t) - R] dt \right] \leq B_1 \sqrt{\mu\beta R} \cdot \mathbb{E}[T_A] + B_2 \beta^2 \mu \sqrt{R}.$$

LEMMA 2 (Upper Bound on Integral Over Draining Period). *Recall that the accumulation time T_A is the first (and only) time the $(R + 1)$ -th server turns on during a renewal cycle, and that the next renewal point $X = \min \{t > T_A : Z(t) = R\}$ is simply the next time the $(R + 1)$ -th server turns off. Then,*

$$\mathbb{E} \left[\int_{T_A}^X [N(t) - R] dt \right] \leq 3.01 \mu \beta^2 \sqrt{R} + 2.04 \beta \frac{\rho}{1 - \rho} + g \left(9(\mu\beta)^2 R, 3\mu\beta\sqrt{R}, k(1 - \rho) \right)$$

where $g(x, y, z) \triangleq x \frac{1}{2\mu z} + y \left[\frac{R}{\mu z^2} + \frac{3}{2\mu z} \right]$.

LEMMA 3 (Lower Bound on Cycle Length). *Taking $L_1 = \frac{2}{3} \sqrt{\frac{\pi}{2}}$, the renewal cycle length $\mathbb{E}[X]$ is*

$$\mathbb{E}[X] \geq \beta + \frac{L_1 \mu \beta \sqrt{R}}{\mu k (1 - \rho)}.$$

5.3. Proof of Theorem 1, assuming the Three Main Lemmas

After proving these lemmas, the result easily follows. First note that, by summing Lemmas 1 and 2,

$$\mathbb{E} \left[\int_0^X [N(t) - R] dt \right] \leq B_1 \sqrt{\mu\beta R} \cdot \mathbb{E}[T_A] + (B_2 + 3.01) \mu \beta^2 \sqrt{R} + 2.04 \beta \frac{\rho}{1 - \rho} + g \left(9(\mu\beta)^2 R, 3\mu\beta\sqrt{R}, k(1 - \rho) \right),$$

where $g(x, y, z) \triangleq x \frac{1}{2\mu z} + y \left[\frac{R}{\mu z^2} + \frac{3}{2\mu z} \right]$.

Upon dividing by Lemma 3 and noting that a setup time $\beta < \mathbb{E}[T_A] < \mathbb{E}[X]$, we obtain

$$\mathbb{E}[Q(\infty)] \leq B_1 \sqrt{\mu\beta R} + 2.04 \frac{\rho}{1 - \rho} + \frac{(B_2 + 3.01) \mu \beta^2 \sqrt{R} + g \left(9(\mu\beta)^2 R, 3\mu\beta\sqrt{R}, k(1 - \rho) \right)}{\beta + \frac{L_1 \mu \beta \sqrt{R}}{\mu k (1 - \rho)}}. \quad \square$$




Figure 6 A depiction of the stopping time decomposition of Lemma 4. In order to bound the integral $\mathbb{E} \left[\int_{T_0}^P Y_t dt \right]$, we bound the expected contribution of each random interval $[T_i, T_{i+1})$, where every T_i is a cleverly defined stopping time. This divide-and-conquer approach features heavily in our proofs.

5.4. Key Technique: The Method of Intervening Stopping Times (MIST)

5.4.1. Technical Challenge: Applying simple bounds to a complex system. We begin by describing the technical challenge which led to the development of the MIST method: attempting to apply simple coupling bounds to a complex, dynamic system. In general, one can easily obtain simple coupling bounds on the behavior of, e.g., the number of jobs in the system $N(t)$ or the number of busy servers $Z(t)$; unfortunately, these simple bounds may only be accurate and/or applicable for a *random* amount of time. For example: Suppose that at time τ the number of busy servers $Z(\tau) = 6$ and the smallest remaining setup time in the system is 10 seconds. Then, until either a server is turned off or 10 seconds have passed, we understand the behavior of the number of jobs $N(t)$ very well — in a sample path sense, it is *precisely* the behavior of an M/M/1 queue, with arrival rate $k\lambda$ and departure rate 6μ . Suppose, though, that by time $\tau + 2$, the system is empty. If we continued to estimate the system behavior via the same M/M/1 queue as before, then we would be wildly over-estimating the true system's departure rate. If we want to keep our couplings accurate, we need to track when the system's behavior changes significantly, and alter our coupling accordingly.

5.4.2. Key Idea: Decompose using state-based stopping times. The Method of Intervening Stopping Times (MIST) allows us to easily and dynamically apply these coupling arguments. Our key idea is to break up our original long time interval into a *random number* of smaller, more manageable pieces. We do this by defining *intervening events*, moments where the system state changes in a way that allows to easily characterize the system's behavior. From there, we can define a “small piece” of time as the time in between our intervening events. For example, in this work, it can often be useful to analyze the system around time points where the number of jobs $N(t)$ gets large. This is useful because, in a system with setup times, if we have *enough* jobs for *long enough*, we can afterwards guarantee that many servers are turned on.

5.4.3. Reduction to three claims. By performing this decomposition of the integral, we reduce our initial bounding problem to showing two or three bounds. First, we must bound the integral of each of the chunks. Often, it is useful to analyze the “initial” chunk separately from the “successive” chunks afterward. Typically, we use martingale arguments combined with worst-case coupling arguments to prove these two integral bounds. Next, we must show that not too many of these “successive chunks” actually occur. For




Figure 7 A depiction of the sequence of “new lows” τ_j and the epochs that they define. In this case, the number of epochs which occur is $n_e = 3$.

this “not too many” condition, it’s often helpful to show some kind of regularity condition, e.g. if the i -th intervening event has occurred, then the $(i + 1)$ -th event occurs with at most constant probability. By formalizing our notion of *intervening events* using stopping times and applying some ideas from Wald’s equation, we obtain the Intervening Stopping Time Lemma, Lemma 4.

5.5. The Intervening Stopping Time Lemma: Statement

We now state the Intervening Stopping Time Lemma, Lemma 4, deferring its proof to Section A. Afterwards, we sketch how we use MIST to prove Lemma 1, the bound on the time integral $\mathbb{E} \left[\int_0^{T_A} [N(t) - R] \right]$.

LEMMA 4 (Intervening Stopping Time Lemma). *Given a starting stopping time T_0 , an ending stopping time P , and a collection of intervening stopping times $(T_i : i \in \mathbb{Z}^+)$, define the random variable F to be such that $T_F \leq P < T_{F+1}$. Now, given some time-varying random variable $Y_t \geq 0$ which is a function of the underlying Markov state of the system $\mathcal{S}(t)$, suppose that:*

1. $\mathbb{E} \left[\int_{T_0}^{\min(T_1, P)} Y_t dt \middle| \mathcal{F}_{T_0} \right] \leq G_0(\mathcal{S}(T_0))$,
2. $\mathbb{E} \left[\int_{T_i}^{\min(T_{i+1}, P)} Y_t dt \middle| \mathcal{F}_{T_i}, F \geq i \right] \leq G_i + B \cdot \mathbb{E} [\min(T_{i+1}, P) - T_i | \mathcal{F}_{T_i}, F \geq i] + C [T_i - T_{i-1}]$,
3. and $\Pr(F \geq i | \mathcal{F}_{T_i}, F \geq i - 1) \leq 1 - p_i$,

where G_0 is also some function of the system state, and the G_i ’s, the p_i ’s, C , and B are all constants (possibly depending on system parameters). Then one can bound the time integral of Y_t from T_0 to P by

$$\mathbb{E} \left[\int_{T_0}^P Y_t dt \right] \leq \mathbb{E}[G_0(\mathcal{S}(T_0))] + \Pr(F > 0) \sum_{j=1}^{\infty} G_j \prod_{i=2}^j (1 - p_i) + (B + C) \cdot \mathbb{E}[P - T_0].$$

5.6. Example: Application to Accumulation Period Upper Bound (Lemma 1)

5.6.1. Application of Lemma 4.

Specification Step. We begin our bound on the integral over the accumulation period with an application of the Intervening Stopping Time Lemma, Lemma 4. To apply our new lemma, we must first specify all the necessary stopping times: the initial time T_0 , the final time P , and the intervening times T_i . In this case, we

of course take $T_0 = 0$ and $P = T_A$. For the intervening times T_i , we choose a sequence which allows us to lower bound the departure rate of our system. In particular, we take $T_j = \tau_j$ where the j -th epoch start τ_j is

$$\tau_j \triangleq \min \{t \geq 0 : N(t) \leq R - j\};$$

in other words, the moments when the number of jobs $N(t)$ reaches a new minimum value or “new low.” We call the period $[\tau_j, \min(\tau_{j+1}, T_A))$ the j -th epoch, and say epoch j occurs whenever $\tau_j < T_A$. We then let the random variable n_e denote the number of epochs that occur, i.e. we take $F = n_e$.

The reduced problem. With this decomposition, understanding the behavior of the accumulation phase reduces to understanding the behavior of each epoch. The beauty in our construction is that the behavior of $N(t)$ within an epoch is much more constrained: within epoch j , the departure rate $\mu Z(t)$ of our system is both lower-bounded by $\mu(R - j)$ and upper-bounded by μR . We use this observation together with various coupling arguments to show the three necessary preconditions of Lemma 4. To prove these preconditions, we actually need to apply Lemma 4 *again*; see Section 6.3 for details.

6. Proof of Theorem 1's Three Main Lemmas

As shown in Section 5.3, to prove Theorem 1 it suffices to prove the following lemmas:

1. Lemma 1, an upper bound on the reward integral over the accumulation phase,
2. Lemma 2, an upper bound on the reward integral over the draining phase, and
3. Lemma 3, a lower bound on the length of our renewal cycle.

Before we continue, we first give a formal system construction, then prove some useful coupling claims.

6.1. Construction

We now discuss how we formally construct this system using Poisson processes; being explicit here will prove useful when we make coupling arguments in the future.

The arrival and departure processes. We take the number of jobs that have arrived at time t to be $\Pi_A(t)$, where Π_A is a Poisson process of rate $k\lambda$. In a slight abuse of notation, we let $\Pi_A([a, b])$ denote the number of arrivals that occur in the interval $[a, b]$; we apply the same extension to all other counting processes mentioned here. We set the potential departure process of, say, server i to be $\Pi_i(t)$, where Π_i is a Poisson process of rate μ . A potential departure from server i only “counts” if server i is busy when that potential departure occurs, i.e., if the number of busy servers $Z(t) \geq i$ at the time. Thus, the total number of departures from our system by time t is, taking integrals with respect to the Poisson processes Π_i as counting processes,

$$D(t) \triangleq \sum_{i=1}^k \int_0^t \mathbf{1}\{Z(s) \geq i\} d\Pi_i(s).$$

The number of busy servers $Z(t)$. To find the number of busy servers $Z(t)$, one could count the number of setup completion events that have occurred so far and the number of server shutoffs that have occurred so far; this description is a bit difficult to work with. Alternatively, one can see from the initial description of setup dynamics that server i is *on* at time t if and only if the total number of jobs $N(s) \geq i$ for all $s \in [t - \beta, t]$, where one should recall that β is the setup time. An easier description of $Z(t)$ follows:

$$Z(t) = \min \left(k, \min_{s \in [t - \beta, t]} N(s) \right).$$

A *departure operator*. We can extend our departure process $D(t)$ to a departure operator $\mathcal{D}[f(s)](\mathcal{I})$ which takes a function $f(s) \in \{0, 1, \dots, k\}$ defined on some interval \mathcal{I} and computes the number of departures that would occur in that interval provided that the number of busy servers $Z(s) = f(s)$, i.e.

$$\mathcal{D}[f(s)]((a, b]) \triangleq \sum_{i=1}^k \int_a^b \mathbf{1}\{f(s) \geq i\} d\Pi_i(s).$$

Note that the total number of departures can now be written as $D(t) = \mathcal{D}[Z(s)]([0, t])$.

6.2. Three Coupling Claims

We now describe three useful claims applied throughout the proof. The first, we will state and prove immediately. The latter two, we prove later, in Section G.

6.2.1. Basic coupling claim: Maintaining an initial relation.

CLAIM 1 (Basic Coupling). *Suppose that we have two processes N_1 and N_2 with an initial relation $N_1(a) \leq N_2(a)$, where the behavior of each process is governed, for all times s from a up to some stopping time τ , by the equation*

$$N_j(s) \triangleq N_j(a) + \Pi_A((a, s]) - \mathcal{D}[Z_j(x)]((a, s]), \text{ for } j \in \{1, 2\}.$$

Furthermore, suppose that the first system's number of busy servers $Z_1(s) \geq Z_2(s)$ for all times $s \in [a, \tau]$. Then, for all $s \in [a, \tau]$, the relation is maintained, i.e. $N_1(s) \leq N_2(s)$.

Proof. We show equivalently that $N_2(s) - N_1(s) \geq 0$. Applying the definitions of N_1 and N_2 ,

$$\begin{aligned} N_2(s) - N_1(s) &= N_2(a) - N_1(a) + [\mathcal{D}[Z_1(x)]((a, s]) - \mathcal{D}[Z_2(x)]((a, s])] \\ &\geq [\mathcal{D}[Z_1(x)]((a, s]) - \mathcal{D}[Z_2(x)]((a, s))] \\ &= \sum_{i=1}^k \int_a^s \mathbf{1}\{Z_1(x) \geq i\} d\Pi_i(x) - \sum_{i=1}^k \int_a^s \mathbf{1}\{Z_2(x) \geq i\} d\Pi_i(x) \\ &= \sum_{i=1}^k \int_a^s \left[\mathbf{1}\{Z_1(x) \geq i\} - \mathbf{1}\{Z_2(x) \geq i\} \right] d\Pi_i(x). \end{aligned}$$

Since $Z_1(x) \geq Z_2(x)$, the integrand $\left[\mathbf{1}\{Z_1(x) \geq i\} - \mathbf{1}\{Z_2(x) \geq i\} \right] \geq 0$; the claim follows. \square

6.2.2. Statement of remaining coupling claims

High-level explanation. This claim leads nicely into a couple more claims. Both are concerned with bounding a quantity involving a general “down-crossing” time. In particular, our analysis will begin at a stopping time τ and will “end” at the down-crossing time d_{gen} , where $d_{\text{gen}} \triangleq \min \{t \geq 0 : N(t + \tau) \leq h\}$ is the length of time it takes for the number of jobs $N(t)$ to become lower than some given threshold h . The first claim, Claim 2, uses a coupling argument to bound the expected integral of $N(t)$ from some arbitrary time τ until $N(t)$ drops below some pre-defined threshold h , provided that one has a lower bound on the number of busy servers $Z(t)$ over that period. The second claim, Claim 3, uses a related argument to bound the probability that $N(t)$ drops below some threshold h within some amount of time ℓ , given that one has bounds on $Z(t)$ over the relevant period. We defer the proof of these claims to Sections E.1 and E.2.

CLAIM 2 (Coupling Integral Bound). *Let τ be some stopping time and d_{gen} be the next down-crossing as described in Section 6.2.2. Suppose that, at time τ , we have a lower bound on the number of busy servers over a period, i.e. we know that the number of busy servers $Z(t) \geq R - j$, for all $t \in [\tau, \tau + \min(\ell, d_{\text{gen}})]$ and for some non-negative j . Then we have the following bound on the integral over this time period:*

$$\mathbb{E} \left[\int_{\tau}^{\tau + \min(d_{\text{gen}}, \ell)} [N(t) - h] dt \middle| \mathcal{F}_{\tau} \right] \leq \ell \cdot [N(\tau) - h]^+ + \frac{1}{2} \mu j \ell^2.$$

CLAIM 3 (Coupling Probability Bound). *Let τ be some stopping time and d_{gen} be the next down-crossing as described in Section 6.2.2. We consider two cases.*

*In the first case, suppose that we have a **lower** bound on the number of busy servers $Z(t)$ over some length ℓ interval starting at time τ , i.e. the busy servers $Z(t) \geq R - j$, for all $t \in [\tau, \tau + \min(\ell, d_{\text{gen}})]$ and for some non-negative j . Then, we can bound the threshold-crossing probability by*

$$\Pr(d_{\text{gen}} < \ell | \mathcal{F}_{\tau}) \geq 2\Phi \left(- \left[\frac{N(\tau) - h + \mu j \ell}{\sqrt{\ell(2k\lambda - \mu j)}} \right] \right) - \frac{2}{3\sqrt{\ell(2k\lambda - \mu j)}}.$$

In particular, if $N(\tau) - h = c_1 \sqrt{\mu\beta R}$, then the probability $\Pr(d_{\text{gen}} < \ell | \mathcal{F}_{\tau}) \geq 2\Phi \left(-\frac{c_1}{\sqrt{2}} \right) - \frac{1}{100}$.

*In the second case, suppose that we instead have the **upper** bound on $Z(t) \leq R$ during this interval instead. Then,*

$$\Pr(d_{\text{gen}} < \ell | \mathcal{F}_{\tau}) \leq 2\Phi \left(- \left[\frac{N(\tau) - h}{\sqrt{2\ell k \lambda}} \right] \right) - \frac{2}{3\sqrt{2k\lambda\ell}}.$$

As before, if $N(\tau) - h = c\sqrt{\mu\beta R}$, then the probability $\Pr(d_{\text{gen}} < \ell | \mathcal{F}_{\tau}) \leq 2\Phi \left(-\frac{c}{\sqrt{2}} \right) + \frac{1}{100}$.

6.3. Proof of Lemma 1, Upper Bound on Integral Over Accumulation Period

We prove this result via two applications of the Intervening Stopping Time Lemma, Lemma 4. To apply this decomposition lemma, there are two broad steps. First, we must specify a starting time (T_0), an ending time (P), a series of intervening stopping times (T_i), the process (Y_t), and an counting variable (F). Second, we must prove that the three preconditions of the lemma hold, given these specifications.

6.3.1. First application of Lemma 4, at the epoch level.

Definition of (τ_j) . We define the sequence of stopping times $(\tau_j : j = 0, 1, \dots, R)$ as $\tau_j \triangleq \min \{t > 0 : N(t) \leq R - j\}$, i.e., τ_j is the first time there are only $R - j$ jobs within the system. Note that, by definition, $\tau_0 = 0$. We call the period $\left[\tau_j, \min(\tau_{j+1}, T_A) \right)$ the j -th epoch, and say epoch j occurs whenever $\tau_j < T_A$. We then let n_e denote the number of epochs which occur in a given renewal cycle.

Specification step. Since we are interested in bounding $\mathbb{E} \left[\int_0^{T_A} [N(t) - R] dt \right]$, we let our starting stopping time be $T_0 = 0$, our ending stopping time be $P = T_A$, our intervening stopping times be $T_j = \tau_j$, the process of interest $Y_t = N(t) - R$ and our counting variable be $F = n_e$. Let the quantity

$$p_{\text{rise}}^{(j)} \triangleq \Pr \left(\max_{t \in [\tau_j, \min(\tau_{j+1}, T_A)]} N(t) \geq R + C_3 \sqrt{\mu\beta R} \mid n_e \geq j \right) \quad (6)$$

be the probability that the total number of jobs $N(t)$ exceeds $R + C_3 \sqrt{\mu\beta R}$ during epoch j .

Required claims. From here, we can apply Lemma 4 after showing the following claims:

CLAIM 4 (Upper Bound on the Probability of Another Epoch). *Recall that the total number of epochs $n_e \triangleq \max \{j \in \mathbb{Z}^+ : \tau_j < T_A\}$. Then, taking $C_4 = 0.98$, we have $\Pr(n_e \geq j + 1 \mid n_e \geq j) \leq 1 - C_4 p_{\text{rise}}^{(j)}$.*

CLAIM 5 (Upper Bound on the Integral Over an Epoch). *Let $\tau_j \triangleq \min \{t \geq 0 : N(t) \leq R - j\}$, $T_A \triangleq \min \{t \geq 0 : Z(t) = R + 1\}$, and let $n_e \triangleq \max \{i \in \mathbb{Z}^+ : \tau_i < T_A\}$. Then,*

$$\mathbb{E} \left[\int_{\tau_j}^{\min(\tau_{j+1}, T_A)} [N(t) - R] dt \mid n_e \geq j \right] \leq B_1 \sqrt{\mu\beta R} \cdot \mathbb{E}[\min(\tau_{j+1}, T_A) - \tau_j \mid n_e \geq j] + C_2 \beta^2 \mu j p_{\text{rise}}^{(j)},$$

where $B_1 = 3.6$ and $C_2 = \frac{1}{2 \cdot 0.98} > 0.511$.

6.3.2. Proof of Lemma 1 assuming Claims 4 and 5. Before going further, we show how to complete the proof of Lemma 1, assuming the two prior claims. Applying Lemma 4, we find that

$$\begin{aligned} \mathbb{E} \left[\int_0^{T_A} [N(t) - R] dt \right] &\leq B_1 \sqrt{\mu\beta R} \cdot \mathbb{E}[T_A] + C_2 \beta^2 \mu \sum_{j=1}^R j p_{\text{rise}}^{(j)} \prod_{i=1}^{j-1} \left(1 - C_4 p_{\text{rise}}^{(i)} \right) \\ &\leq B_1 \sqrt{\mu\beta R} \cdot \mathbb{E}[T_A] + \frac{C_2}{C_4} \beta^2 \mu \left[\sum_{j=1}^R j C_4 p_{\text{rise}}^{(j)} \prod_{i=1}^{j-1} \left(1 - C_4 p_{\text{rise}}^{(i)} \right) \right] \\ &= B_1 \sqrt{\mu\beta R} \cdot \mathbb{E}[T_A] + \frac{C_2}{C_4} \beta^2 \mu \left[\sum_{j=1}^R \prod_{i=1}^j \left(1 - C_4 p_{\text{rise}}^{(i)} \right) \right], \end{aligned}$$

where we have used the ‘‘expectation as a sum of tails’’ trick. We now apply the following claim:

CLAIM 6 (Bound on the Probability of an Up-crossing $p_{\text{rise}}^{(j)}$). *Let $p_{\text{rise}}^{(j)}$ be the probability that the total number of jobs $N(t)$ exceeds $R + C_3 \sqrt{\mu\beta R}$ during epoch j defined in (6). Then, for any epoch $j \geq A_5 \sqrt{R}$, we have $p_{\text{rise}}^{(j)} \geq 0.99 \frac{A_5}{\sqrt{R}}$.*

Continuation: Proof of Lemma 1 assuming Claims 4, 5, and 6. We defer the proof of Claim 6 to Section G.3. Applying the claim's result, we find that

$$\sum_{j=1}^R \prod_{i=1}^j \left(1 - C_4 p_{\text{rise}}^{(j)}\right) \leq \sum_{j=1}^R \left(1 - \frac{0.99C_4 A_5}{\sqrt{R}}\right)^{[j - A_5 \sqrt{R}]^+} \leq \sum_{j=1}^{\infty} \left(1 - \frac{0.99C_4 A_5}{\sqrt{R}}\right)^{[j - A_5 \sqrt{R}]^+}. \quad (7)$$

Bounding this as a Geometric sum, we obtain

$$(7) = A_5 \sqrt{R} + \sum_{j=0}^{\infty} \left(1 - \frac{0.99C_4 A_5}{\sqrt{R}}\right)^j = A_5 \sqrt{R} + \frac{1}{0.99C_4 A_5} \sqrt{R}.$$

Returning to our original inequality, we obtain that

$$\mathbb{E} \left[\int_0^{T_A} [N(t) - R] dt \right] \leq B_1 \sqrt{\mu \beta R} \cdot \mathbb{E}[T_A] + \frac{C_2}{C_4} \left(A_5 + \frac{1}{0.99C_4 A_5} \right) \beta^2 \mu \sqrt{R}.$$

Noting that $A_5 = 1$ and taking $B_2 \triangleq 1.04 > \frac{C_2}{C_4} \left(A_5 + \frac{1}{0.99C_4 A_5} \right)$, we finish the proof of Lemma 1. \square

6.3.3. Proof of Claim 4, Upper Bound on Probability of Another Epoch.

Rewriting the claim. And so, assuming the preconditions of Lemma 4 (Claims 4 and 5) as well as the helper claim 6, we have proven Lemma 1. We thus begin proving Claim 4. We begin by rewriting the probability of another epoch occurring as

$$\Pr(n_e \geq j + 1 | n_e \geq j) = 1 - \Pr(n_e = j | n_e \geq j) = 1 - \Pr(T_A < \tau_{j+1} | n_e \geq j).$$

It thus suffices to show a bound on the probability that the accumulation phase ends in epoch j :

$$\Pr(T_A < \tau_{j+1} | n_e \geq j) \geq C_4 p_{\text{rise}}^{(j)}. \quad (8)$$

Lower bound based on up-crossing and down-crossing times. To show (8), we analyze a particular sequence of events which results in the accumulation phase ending in the current epoch, i.e. $T_A < \tau_{j+1}$. Specifically, we define the up-crossing time $u = \min \{t > \tau_j : N(t) \geq R + C_3 \sqrt{\mu \beta R}\}$ and the down-crossing time $d = \min \{t > u : N(t) \leq R\}$. We consider the event where (1) the up-crossing occurs during the accumulation phase ($u < T_A$) and (2) the accumulation phase ends before the next down-crossing occurs ($d > T_A$). Symbolically, we have (at the end, recalling that $p_{\text{rise}}^{(j)}$ is the probability of an up-crossing occurs)

$$(8) \geq \Pr(u < T_A < d | n_e \geq j) = \Pr(d > T_A | u < T_A) \Pr(u < T_A) = \Pr(d > T_A | u < T_A) p_{\text{rise}}^{(j)}.$$

Development of conditional probability. To bound the conditional probability $\Pr(d > T_A | u < T_A)$, we condition on the filtration at time u , then make a coupling argument. To begin, note that, if the number of jobs $N(t)$ does not fall to R before the $(R + 1)$ -th server finishes setting up, then the accumulation time T_A occurs exactly when the $(R + 1)$ -th server finishes, i.e. the accumulation time $T_A = u + Y_{R+1}(u)$.




Figure 8 A depiction of the up-crossings and down-crossings defined in Section 6.3.4. In this example, we see that the number of up-crossings in epoch 3 is $n_e^{(3)} = 2$ and that, in this case, epoch 3 ends when epoch 4 begins (i.e. at time τ_4).

Furthermore, the number of busy servers $Z(t) \leq R$ at any time during the accumulation phase $t < T_A$. Applying a basic coupling argument (Claim 1), we have a lower bound on $N(t)$ in the coupled process

$$\tilde{N}(t) \triangleq N(u) + \Pi_A((u, t]) - \mathcal{D}[R]((u, t]),$$

for any time $t \in [u, T_A]$. Let $\tilde{d} \triangleq \min \{t > u : \tilde{N}(t) \geq R\}$ be the analogous down-crossing time in the coupled system. Since the coupled $\tilde{N}(t)$ is a lower bound, the coupled down-crossing time $\tilde{d} \leq d$. Thus,

$$\Pr(d > T_A | \mathcal{F}_u, u < T_A) = \Pr(d > u + Y_{R+1}(u) | \mathcal{F}_u, u < T_A) \geq \Pr(\tilde{d} > u + Y_{R+1}(u) | \mathcal{F}_u, u < T_A). \quad (9)$$

Analyzing the coupled probability. Continuing, the probability that $\{\tilde{d} \geq \ell\}$ is decreasing in ℓ . Thus,

$$(9) \geq \Pr(\tilde{d} > u + \beta | \mathcal{F}_u, u < T_A) = \Pr(\tilde{d} - u > \beta) \geq 1 - 2\Phi\left(-\frac{C_3}{\sqrt{2}}\right) - \frac{2}{3\sqrt{\mu\beta R}} \geq 0.98,$$

where in the final inequalities we have applied both the down-crossing probability bound of Claim 3 and our assumptions. Taking $C_4 \triangleq 0.98$, we have the inter-epoch probability bound of Claim 4. \square

6.3.4. Proof of Claim 5, Upper Bound on the Integral Over an Epoch. We now prove Claim 5, the upper bound on the time integral over an epoch. We do this via another application of Lemma 4 —first specifying the intervening times, then completing the proof, then proving that the preconditions hold.

Definition of up-crossings and down-crossings. Let the 0-th down-crossing time in epoch j occur at time τ_j , i.e. let $d_0^{(j)} \triangleq \tau_j$. Next, define the first up-crossing in epoch j as the first time during epoch j that the total number of jobs $N(t)$ exceeds $R + C_3\sqrt{\mu\beta R}$, i.e.

$$u_1^{(j)} \triangleq \min \left\{ t > \tau_j : N(t) \geq R + C_3\sqrt{\mu\beta R} \right\}.$$

From here, define i -th down-crossing in epoch j and the $i + 1$ -th up-crossing in epoch j as

$$d_i^{(j)} \triangleq \min \left\{ t \geq u_i^{(j)} : N(t) \leq R \right\}$$

and

$$u_{i+1}^{(j)} \triangleq \min \left\{ t \geq d_i^{(j)} : N(t) \geq R + C_3 \sqrt{\mu\beta R} \right\},$$

respectively; we visualize these definitions in Figure 8. We say the i -th up-crossing occurs if $u_i^{(j)} < \min(T_A, \tau_{j+1})$ and let $n_u \triangleq \max \left\{ i \geq 0 : u_i^{(j)} < \min(\tau_{j+1}, T_A) \right\}$ be the random number of up-crossings which occur in epoch j . We call the interval $\left[d_i^{(j)}, \min(u_i^{(j)}, \min(T_A, \tau_{j+1})) \right)$ the i -th rise, and the interval $\left[u_i^{(j)}, \min(d_i^{(j)}, \min(T_A, \tau_{j+1})) \right)$ the i -th fall. Note that, if the i -th up-crossing occurs, then, by definition, $d_i < \tau_{j+1}$; this means that the i -th fall can always be written as $\left[u_i, \min(T_A, d_i) \right)$. For readability, we fix our epoch of interest and freely omit the superscript j on our up-crossings and down-crossings.

Specification step. With up-crossings and down-crossings defined, we are now ready to specify our application of the IST Lemma, Lemma 4. We define our starting time as $T_0 = \tau_j = d_0$, our ending time as $P = \min(T_A, \tau_{j+1})$, our intervening sequence as $(u_i)_{i=1}^{\infty}$, and our counting variable as $F = n_u$.

Required Claims. From here, in order to apply Lemma 4, we must show the following three claims:

CLAIM 7 (Upper Bound on Integral Until First Up-crossing). *Taking $B_1 = 3.6$, the integral until u_1 is*

$$\mathbb{E} \left[\int_{d_0}^{\min(u_1, \min(T_A, \tau_{j+1}))} [N(t) - R] dt \middle| n_u \geq j \right] \leq B_1 \sqrt{\mu\beta R} \cdot \mathbb{E} [\min(u_1, \min(T_A, \tau_{j+1})) - \tau_j | n_u \geq j].$$

CLAIM 8 (Upper Bound on Integral Between Up-crossings). *The integral between up-crossings u_i is*

$$\mathbb{E} \left[\int_{u_i}^{\min(u_{i+1}, \min(T_A, \tau_{j+1}))} [N(t) - R] dt \middle| n_u \geq i \right] \leq B_1 \sqrt{\mu\beta R} \cdot \mathbb{E} [\min(u_{i+1}, \min(T_A, \tau_{j+1})) - u_i | n_u \geq i] + \frac{1}{2} \beta^2 \mu j.$$

CLAIM 9 (Upper Bound on Probability of Another Up-crossing). *Recall that $p_{\text{rise}}^{(j)}$ is the probability that the number of jobs ($N(t) \geq C_3 \sqrt{\mu\beta R}$) at some point during epoch j , given that epoch j occurs. Then, $\Pr(n_u > 0) = p_{\text{rise}}^{(j)}$, and, for all counts $i \geq 1$ and $p_2 = 0.98$, we have $\Pr(n_u \geq i + 1 | n_u \geq i) \leq 0.02 = 1 - p_2$.*

6.3.5. Proof of Claim 5, assuming Claims 7, 8, and 9. Once again, before we move on to proving these claims, we show that they indeed suffice to prove Claim 5. By Lemma 4, taking $C_2 \triangleq \frac{0.5}{p_2}$,

$$\begin{aligned} \mathbb{E} \left[\int_{\tau_j}^{\min(T_A, \tau_{j+1})} [N(t) - R] dt \middle| n_e \geq j \right] &\leq B_1 \sqrt{\mu\beta R} \cdot \mathbb{E} [\min(T_A, \tau_{j+1}) - \tau_j | n_e \geq j] \\ &\quad + p_{\text{rise}}^{(j)} 0.5 \beta^2 \mu j \sum_{i=1}^{\infty} (1 - p_2)^{i-1} \\ &= B_1 \sqrt{\mu\beta R} \cdot \mathbb{E} [\min(T_A, \tau_{j+1}) - \tau_j | n_e \geq j] + p_{\text{rise}}^{(j)} \frac{0.5}{p_2} \beta^2 \mu j. \quad \square \end{aligned}$$

Proofs of Claims 7, 8, and 9. All that remains to be proven are our three aforementioned claims.

Proof of Claim 7: Upper Bound on Integral until First Up-crossing. Proving this claim is quite simple. In fact, we now prove a far more general claim, that the integral from a down-crossing to the next up-crossing

$$\int_{d_i}^{\min(u_i, \min(T_A, \tau_{j+1}))} [N(t) - R] dt \leq C_3 \sqrt{\mu\beta R} \cdot [\min(T_A, \tau_{j+1}) - d_i]. \quad (10)$$

To see this, note that, at any point between a down-crossing and up-crossing, the total number of jobs $N(t)$ must be strictly less than $R + C_3 \sqrt{\mu\beta R}$. Apply this to the 0-th down-crossing and we have the claim. \square

Proof of Claim 8: Upper Bound on Integral Between Up-crossings. This proof is a bit more involved. We separate the interval $\left[u_i, \min(u_{i+1}, \min(T_A, \tau_{j+1})) \right)$ into the i -th fall and the i -th rise, as discussed previously. For the rising portion, we can simply apply the simple bound from (10). For the falling portion, we apply the integral coupling claim, Claim 2. In particular, note that $Z(t) \geq R - j$ until time τ_{j+1} and that the interval $[u_i, \min(d_i, T_A))$ is equivalent to the interval $[u_i, \min(d_i, u_i + Y_{R+1}(u_i))]$. Applying Claim 2,

$$\begin{aligned} \mathbb{E} \left[\int_{u_i}^{\min(d_i, T_A)} [N(t) - R] dt \middle| S(u_i) \right] &\leq \frac{1}{2} \beta^2 \mu j + [N(u_i) - R] \cdot Y_{R+1}(u_i) \\ &= \frac{1}{2} \beta^2 \mu j + C_3 \sqrt{\mu\beta R} \cdot Y_{R+1}(u_i). \end{aligned}$$

By our analysis in the proof of Claim 4, we note that the remaining setup time $Y_{R+1}(u_i) \leq \min(d_i, T_A)$ with probability at least $p_2 \leq \min_{\mathcal{F}_{u_i}} \Pr(d_i < T_A | \mathcal{F}_{u_i}, n_u \geq i)$. By Markov's inequality,

$$Y_{R+1}(u_i) \leq \frac{1}{p_2} \mathbb{E}[\min(d_i, T_A) - u_i | S(u_i)].$$

Combining our bounds on the rises and falls and taking $B_1 = \frac{C_3}{p_2}$, we have Claim 8. \square

Proof of Claim 9: Upper Bound on Probability of Another Up-Crossing. We now proceed to our final claim, concerning the up-crossing probabilities $\Pr(n_u > 0)$ and $\Pr(n_u \geq i + 1 | n_u \geq i)$. To begin, we first note that, since the first up-crossing occurs precisely at the moment that $N(t)$ exceeds $R + C_3 \sqrt{\mu\beta R}$ during epoch j , one has $p_{\text{rise}}^{(j)} = \Pr(u_1 < \min(T_A, \tau_{j+1})) = \Pr(n_u > 0)$.

To prove the second part of the claim, first observe that an $(i + 1)$ -th up-crossing can only occur if an i -th down-crossing occurs, i.e. $\Pr(n_u \geq i + 1 | n_u \geq i) \leq \Pr(d_i < T_A | n_u \geq i)$.

To bound this conditional probability, we can apply a previous result. Recall the proof of the inter-epoch probability bound (Claim 4). In (8), we have already argued a bound on the conditional probability that the *first* down-crossing occurs, *in a state-independent manner*. The bound derived there thus *also applies here*:

$$\Pr(d_i < T_A | n_u \geq i) = 1 - \Pr(d_i > T_A | n_u \geq i) \leq 1 - C_4.$$

Taking $p_2 \triangleq C_4 = 0.98$, we have bounded the probability of another up-crossing (Claim 8). \square

6.4. Proof of Lemma 2, Upper Bound on Integral Over Draining Period.

To prove this lemma, we again make use of Lemma 4. We proceed through the usual two-step process, first defining the stopping time sequence we will analyze over, then proving the preconditions of the lemma.

Definition of the upward visits $v_i^{(up)}$ and downward visits $v_i^{(down)}$. Recall that the draining phase begins at time T_A . Let $M_B \triangleq \min(k - R, \sqrt{R})$ be a specially-set analysis threshold. Let the stopping time $v_1^{(down)} \triangleq \min\{t \geq T_A : N(t) < R + M_B\}$ be the first time the number of jobs $N(t)$ drops below $R + M_B$, and recursively define

$$v_i^{(up)} \triangleq \min\left\{t \geq v_i^{(down)} : N(t) \geq R + M_B\right\}$$

and

$$v_{i+1}^{(down)} \triangleq \min\left\{t \geq v_i^{(up)} : N(t) < R + M_B\right\}.$$

6.4.1. Application of Lemma 4. Applying Lemma 4, we take our initial stopping time to be the accumulation time T_A , our final stopping time to be the end of the renewal cycle X , our intervening stopping times to be the downward visits $v_i^{(down)}$, and our counting index to be n_b .

Required claims. To apply Lemma 4, we need to show is the usual three claims: a bound on the initial integral, a bound on the continuing integral, and a bound on the probability.

CLAIM 10 (Upper Bound on Integral Until First Downward Visit). Let $g(x, y, z) \triangleq x \frac{1}{2\mu z} + y \left[\frac{R}{\mu z^2} + \frac{3}{2\mu z} \right]$. Then, one can bound the integral immediately after time T_A with

$$\begin{aligned} \mathbb{E} \left[\int_{T_A}^{v_1^{(down)}} [N(t) - R] dt \right] &\leq \left[\beta + \frac{1}{\mu} \right] \left[3\mu\beta\sqrt{R} + \max\left(\sqrt{R}, \frac{\rho}{1-\rho}\right) \right] + \frac{2}{\mu} \ln \left(3 \frac{\mu\beta}{\sqrt{R}} \max\left(\sqrt{R}, \frac{\rho}{1-\rho}\right) \right) \\ &\quad + g\left(9(\mu\beta)^2 R, 3\mu\beta\sqrt{R}, k(1-\rho)\right) + \beta \frac{\rho}{1-\rho}. \end{aligned}$$

CLAIM 11 (Upper Bound on Integral Between Downward Visits). One can bound the integral between consecutive downward visits by

$$\mathbb{E} \left[\int_{v_i^{(down)}}^{v_{i+1}^{(down)}} [N(t) - R] dt \middle| \mathcal{F}_{v_i^{(down)}} \right] \leq \frac{1}{\mu M_B} \left[\max\left(\sqrt{R}, \frac{\rho}{1-\rho}\right) + 14 + \mu\beta\sqrt{R} + 2b_1\sqrt{\mu\beta R} + b_2\sqrt{R} \right]$$

CLAIM 12 (Upper Bound on Probability of Another Downward Visit). One can bound the probability of another downward visit occurring by

$$\Pr(n_b \geq i + 1 | n_b \geq i) \leq \frac{1}{M_B}. \quad (11)$$

6.4.2. Proof of Lemma 2 assuming Claims 10, 11, and 12. Simplifying the first bound further,

$$\begin{aligned} \mathbb{E} \left[\int_{T_A}^{v_1^{(down)}} [N(t) - R] dt \right] &\leq \left[\beta + \frac{3}{\mu} \right] \left[2.9\mu\beta\sqrt{R} + \max\left(\sqrt{R}, \frac{\rho}{1-\rho}\right) \right] \\ &\quad + g\left(9(\mu\beta)^2 R, 3\mu\beta\sqrt{R}, k(1-\rho)\right) + \beta \frac{\rho}{1-\rho} \\ &\leq 1.03\beta \left[2.91\mu\beta\sqrt{R} + \frac{\rho}{1-\rho} \right] + g\left(9(\mu\beta)^2 R, 3\mu\beta\sqrt{R}, k(1-\rho)\right) + \beta \frac{\rho}{1-\rho} \\ &\leq 3\mu\beta^2\sqrt{R} + 2.03\beta \frac{\rho}{1-\rho} + g\left(9(\mu\beta)^2 R, 3\mu\beta\sqrt{R}, k(1-\rho)\right), \end{aligned}$$

where we have bounded $\ln(x)/x < 0.1$ for values $x > 100$, noted that $\max(x, y) \leq x + y$, and done some upwards rounding. Simplifying the “continuing integral” term, we have

$$\begin{aligned} \mathbb{E} \left[\int_{v_1^{(\text{down})}}^X [N(t) - R] dt \right] &\leq \frac{1}{\text{Claim 12}} \cdot [\text{Claim 11}] \\ &= \frac{1}{\mu} \left[\max \left(\sqrt{R}, \frac{\rho}{1-\rho} \right) + 14 + \mu\beta\sqrt{R} + 2b_1\sqrt{\mu\beta R} + b_2\sqrt{R} \right] \\ &\leq \frac{1}{\mu} \left[\mu\beta\sqrt{R} \left(1 + \frac{2b_1}{\sqrt{\mu\beta}} + \frac{b_2}{\beta} + \frac{14}{\mu\beta\sqrt{R}} \right) + \frac{\rho}{1-\rho} \right] \\ &\leq \frac{1}{\mu} \left[1.6\mu\beta\sqrt{R} + \frac{\rho}{1-\rho} \right]. \end{aligned}$$

Combining these two pieces, we obtain, as desired,

$$\mathbb{E} \left[\int_{T_A}^X [N(t) - R] dt \right] \leq 3.01\mu\beta^2\sqrt{R} + 2.04\beta\frac{\rho}{1-\rho} + g \left(9(\mu\beta)^2R, 3\mu\beta\sqrt{R}, k(1-\rho) \right). \quad \square$$

Precursor: The “Wait-Busy” idea. As such, to complete our proof it suffices to show Claims 10, 11, and 12. To prove these claims, we make heavy use of the following idea.

CLAIM 13 (Wait Busy Claim). *Let τ be some stopping time, let the number of jobs $N(\tau) = R + h$, and define $\text{ns}(h) \triangleq \min\{h, k(1-\rho)\}$. Let the down-crossing $d_{\text{gen}} \triangleq \min\{t > 0 : N(\tau + t) = R + h - 1\}$. If $Z(\tau) \geq R$, then*

$$\mathbb{E} \left[\int_{\tau}^{\tau + d_{\text{gen}}} [N(t) - (R + h - 1)] dt \middle| \mathcal{F}_{\tau} \right] \leq Y_{R+\text{ns}(h)}(\tau) + g \left(1 + 2\mu R \mathbb{E} \left[\min \left(Y_{R+\text{ns}(h)}(\tau), d_{\text{gen}} \right) \right], 1, \mu \text{ns}(h) \right), \quad (12)$$

where the function $g(x, y, z) \triangleq x\frac{1}{2\mu z} + y \left[\frac{R}{\mu z^2} + \frac{3}{2\mu z} \right]$.

Furthermore,

$$\begin{aligned} \mathbb{E} \left[\int_{T_A}^{v_1^{(\text{down})}} [N(t) - R] dt \middle| \mathcal{F}_{T_A} \right] &\leq \left[\beta + \frac{1}{\mu} \right] \left[\mathbb{E} [N(T_A) - R] + \frac{R}{M_B} \right] + \frac{2}{\mu} \ln \left(\frac{\mathbb{E} [N(T_A) - R]}{M_B} \right) \\ &\quad + g \left(\mathbb{E} \left[[N(T_A) - R]^2 \right] + 2\mathbb{E} [N(T_A) - R], \mathbb{E} [N(T_A) - R], k(1-\rho) \right). \end{aligned} \quad (13)$$

Intuition. We defer the proof of Claim 13 until Section 6.4.2. For now, we give some brief intuition for how the bound is derived and how we use it in our proof. Essentially, we can consider performing the following procedure at time τ : First, watch the system for β time. If the number of jobs ever dips below $R + h$ during this watching period, we can end our integral immediately. If the number of jobs $N(t)$ never dips below $R + h$ during this watching period, then we know for sure that we have at least $\min(R + h, k)$ servers on at time $\tau + Y_{\min(R+h, k)}(\tau)$, since we have continually had at least $R + h$ servers either busy or setting up during that period. Moreover, since we only turn off servers when there isn't work for them to do, those servers will *stay* on until the number of jobs $N(t)$ dips below $R + h$; in other words, they will stay on until time d_{gen} . The proof of the claim follows along essentially the same lines, formalizing things and performing computations using coupling and martingales.

6.4.3. Proof of Claim 10: Bound on Integral Until First Visit. We return to proving our claims. The proof of this claim is simple; it is essentially rolled into the proof of Claim 13. From here, it suffices to apply the following claim, substituting in and simplifying constants:

CLAIM 14 (Upper Bound on $\mathbb{E}[N(T_A)]$). *Recall that $T_A \triangleq \min\{t > 0 : Z(t) = R + 1\}$. Then,*

$$\mathbb{E}[N(T_A) - R] \leq F_1 \mu \beta \sqrt{R} \left(1 + \frac{F_2}{\sqrt{\mu \beta}}\right) \leq 2.9 \mu \beta \sqrt{R}$$

and

$$\mathbb{E}\left[(N(T_A) - R)^2\right] \leq F_1^2 (\mu \beta)^2 R \left(1 + \frac{F_2}{\sqrt{\mu \beta}}\right)^2 + 2 \mu \beta R \leq 8.4 (\mu \beta)^2 R + 2 \mu \beta R$$

where $F_1 = 2.12$ and $F_2 = 3.645$.

6.4.4. Proof of Claim 11: Bound on Integral Between Visits. To prove Claim 11, we break the integral into two parts: from the down-crossing $v_i^{(\text{down})}$ to the up-crossing $v_i^{(\text{up})}$, and vice-versa.

First part: from $v_i^{(\text{down})}$ to $v_i^{(\text{up})}$. To bound the integral from the down-crossing to the next up-crossing, we first make the simple observation that

$$\int_{v_i^{(\text{down})}}^{\min(v_i^{(\text{up})}, X)} [N(t) - R] dt \leq \left[\min(v_i^{(\text{up})}, X) - v_i^{(\text{down})}\right] \cdot M_B, \quad (14)$$

since $v_i^{(\text{up})}$ is the next time $N(t) \geq R + M_B$. To bound $\mathbb{E}\left[\min(v_i^{(\text{up})}, X) - v_i^{(\text{down})} \mid \mathcal{F}_{v_i^{(\text{down})}}\right]$, we couple the system to an infinite-server $M/M/\infty$ queue at time $v_i^{(\text{down})}$, and note that the coupled up-crossing time

$$\tilde{T}_{(R+M_B-1) \rightarrow (R+M_B)} + v_i^{(\text{down})} \geq v_i^{(\text{up})} \geq \min(v_i^{(\text{up})}, X).$$

Since $M_B \leq \sqrt{R}$, from standard results on the $M/M/\infty$ (reproduced in Section F.4), we have that

$$M_B \mathbb{E}\left[\min(v_i^{(\text{up})}, X) - v_i^{(\text{down})} \mid \mathcal{F}_{v_i^{(\text{down})}}\right] \leq M_B \frac{b_2}{\mu \sqrt{R}} = \frac{1}{\mu M_B} b_2 \frac{M_B^2}{\sqrt{R}} \leq \frac{1}{\mu M_B} b_2 \sqrt{R}.$$

Second part: from $v_i^{(\text{up})}$ to $v_{i+1}^{(\text{down})}$. From $v_i^{(\text{up})}$ onwards, we use the “wait-busy” bound. Applying the “Wait Busy” Claim (Claim 13) with $h = M_B$, we obtain

$$\begin{aligned} \mathbb{E}\left[\int_{v_i^{(\text{up})}}^{v_{i+1}^{(\text{down})}} [N(t) - R] dt \mid \mathcal{F}_{v_i^{(\text{up})}}\right] &\leq \beta + g \left(1 + 2\mu R \mathbb{E}\left[\min\left(\beta, v_{i+1}^{(\text{down})} - v_i^{(\text{up})} \mid \mathcal{F}_{v_i^{(\text{up})}}\right)\right], 1, M_B\right) \\ &\quad + M_B \mathbb{E}\left[\min\left(\beta, v_{i+1}^{(\text{down})} - v_i^{(\text{up})} \mid \mathcal{F}_{v_i^{(\text{up})}}\right)\right]. \end{aligned}$$

To bound the conditional expectation of $\mathbb{E}\left[\min\left(\beta, v_{i+1}^{(\text{down})} - v_i^{(\text{up})} \mid \mathcal{F}_{v_i^{(\text{up})}}\right)\right]$, we make our usual coupling argument. Define a coupled system $M/M/1$ queue with departure rate μR , and let \tilde{d}_{gen} be the length of its

busy period. It suffices to bound $\mathbb{E} \left[\min \left(d_{\text{gen}}^{\sim}, \beta \right) \right]$, for the coupled relative down-crossing time d_{gen}^{\sim} . From standard results on simple random walks (Claim 21), we have $\mathbb{E} \left[\min \left(d_{\text{gen}}^{\sim}, \beta \right) \right] \leq \frac{b_1 \sqrt{\mu \beta}}{\mu \sqrt{R}} + \frac{6}{\mu R}$, giving

$$\begin{aligned} \mathbb{E} \left[\int_{v_i^{(\text{up})}}^{v_{i+1}^{(\text{down})}} [N(t) - R] dt \middle| \mathcal{F}_{v_i^{(\text{up})}} \right] &\leq \frac{1}{\mu M_B} \left[\frac{R}{M_B} + 2 \right] + \beta + \left[\frac{b_1 \sqrt{\mu \beta}}{\mu \sqrt{R}} + \frac{6}{\mu R} \right] \left(M_B + \frac{R}{M_B} \right) \\ &= \frac{1}{\mu M_B} \left[\frac{R}{M_B} + 2 + \mu \beta M_B + (M_B^2 + R) \left[\frac{b_1 \sqrt{\mu \beta}}{\sqrt{R}} + \frac{6}{R} \right] \right] \\ &\leq \frac{1}{\mu M_B} \left[\max \left(\sqrt{R}, \frac{\rho}{1 - \rho} \right) + 14 + \mu \beta \sqrt{R} + 2b_1 \sqrt{\mu \beta R} \right], \end{aligned}$$

where in the last line we used $M_B \leq \sqrt{R}$; combining these two parts gives Claim 11. \square

6.4.5. Proof of Claim 12, Upper Bound on the Probability of Another Visit. To see (11), we first note that, if there is another upcrossing, then there must be another downcrossing. As such, it suffices to upper bound $\Pr \left(v_i^{(\text{up})} < X \middle| v_i^{(\text{down})} \right)$. To do this, we note that the number of busy servers $Z(t) \geq R$. From Claim 1, it thus suffices to bound the corresponding probability in the coupled system with exactly R busy servers. But this is simply the probability that a simple random walk started at $W(0) = M_B - 1$ hits $W(t) = M_B$ before it hits $W(t) = 0$. Classically, this probability is $\frac{1}{M_B}$; this proves the claim, and thus Lemma 2. \square

6.5. Proof of Lemma 3: Lower Bound on the Cycle Length.

Preliminaries. The proof of this lemma is much simpler than the others. Before describing our strategy, we first state some preliminaries. Recall the definition of the start of the j -th epoch $\tau_j \triangleq \min \{ t \geq 0 : N(t) \leq R - j \}$, that we call the period $\left[\tau_j, \min(\tau_{j+1}, T_A) \right)$ the j -th epoch, and we say epoch j occurs if $\tau_j < T_A$. Now, say epoch j is *long* if it lasts longer than a setup time β ; note that such an epoch must exist, since servers can only turn on during long epochs, and a server must turn on before the accumulation phase ends at time T_A . Let $L \triangleq \min \{ j \in \{0, 1, 2, \dots, R\} : \min(\tau_{j+1}, T_A) - \tau_j > \beta \}$ be the index of the *first* long epoch. Note that, although the random time τ_L is *not* a stopping time (we do not know how long an epoch will last when the epoch starts), the first moment we can identify epoch L , the random time $\tau_L + \beta$, is a stopping time. Moreover, we know that $\tau_L + \beta < T_A$. From here, one sees that $\mathbb{E}[X] = \mathbb{E}[\tau_L + \beta] + \mathbb{E}[X - (\tau_L + \beta)] \geq \beta + \mathbb{E}[X - (\tau_L + \beta)]$. To complete the proof, it suffices to show

$$\mathbb{E}[X - (\tau_L + \beta)] \geq \frac{L_1 \mu \beta \sqrt{R}}{\mu k (1 - \rho)}. \quad (15)$$

6.5.1. Proof of (15): Lower Bound on the Remaining Cycle Length. To show (15), we first show we can bound an analogous quantity in a coupled process, then appeal to standard results on the M/M/1 queue.

Defining the coupled process $\tilde{N}(t)$. To define the coupled process, note that the number of busy servers $Z(t) \leq k$. It follows from Claim 1 that, for any time $t \geq \tau_L + \beta$, the coupled process

$$\tilde{N}(t) \triangleq N(\tau_L + \beta) + A((\tau_L + \beta, t]) - \mathcal{D}[k]((\tau_L + \beta, t])$$

satisfies $\tilde{N}(t) \leq N(t)$.

Using the coupled process to bound $\mathbb{E}[X - \tau_L + \beta]$. We now use this process to bound $\mathbb{E}[X - \tau_L + \beta]$. Recall that the end of the renewal cycle $X \triangleq \min\{t > 0 : Z(t^-) = R + 1, Z(t) = R\}$ occurs when the $(R + 1)$ -th server turns off. It is useful to view X in a different way: since the accumulation time T_A is the moment when the $(R + 1)$ -th server turns on, we also know that the end of the renewal cycle $X = \min\{t > T_A : N(t) \leq R\}$ is the first moment after time T_A that the number of jobs $N(t) \leq R$. Furthermore, since the time $\tau_L + \beta$ happens *before* any server could possibly turn on, the time $\tau_L + \beta < T_A$. Denoting the end of the coupled renewal as \tilde{X} as the first moment the coupled process $\tilde{N}(t) \leq R$, we have

$$\tilde{X} \triangleq \min\{t > \tau_L + \beta : \tilde{N}(t) \leq R\} \leq \min\{t > \tau_L + \beta : N(t) \leq R\} \leq \min\{t > T_A : N(t) \leq R\} = X.$$

Bounding the end of coupled renewal $\mathbb{E}[X - \tau_L + \beta]$. To bound the quantity $\mathbb{E}[X - \tau_L + \beta]$, we condition on the filtration at time $\tau_L + \beta$ and use standard results on the M/M/1 busy period. Note that, since the departure rate of the coupled system is fixed at μk , the period $[\tilde{X} - \tau_L + \beta]$ is precisely the length of an M/M/1 busy period with 1) arrival rate $k\lambda$, 2) departure rate $k\mu$, and 3) started by $[N(\tau_L + \beta) - R]^+$ jobs. It follows that $\mathbb{E}[X - (\tau_L + \beta) | \mathcal{F}_{\tau_L + \beta}] \geq \frac{[N(\tau_L + \beta) - R]^+}{\mu k(1 - \rho)}$. Taking expectations, applying Jensen's and results from the lower bound ((21) and (18)), we obtain, proving (15), Lemma 3, and Theorem 1 simultaneously,

$$\mathbb{E}[X - (\tau_L + \beta)] \geq \frac{[\mathbb{E}[N(\tau_L + \beta) - R]]^+}{\mu k(1 - \rho)} \geq \frac{\mu \beta \mathbb{E}[L]}{\mu k(1 - \rho)} \geq \frac{L_1 \mu \beta \sqrt{R}}{\mu k(1 - \rho)}. \quad \square$$

7. Conclusion and Future Work

Summary. In this work, we provided a comprehensive characterization of the average waiting time within the M/M/k/Setup-Deterministic. In particular, we proved the first multiplicatively-tight upper and lower bounds, respectively, on the average waiting time. By combining our upper and lower bounds, we also obtained an approximation which we demonstrate to be highly accurate. These results significantly advance our understanding of how setup times affect queueing systems beyond the prevalent Exponential setup model, providing new insights into capacity provisioning in systems with setup times.

7.1. Future Work.

Tail of Response Time. Our work leaves open many possible directions for future research. One such possibility is studying other aspects of system performance in the M/M/k/Setup-Deterministic. For context, in this work, we only analyze the average waiting time of a job in the system. However, in many applications,

a customer's quality of service is determined by more advanced metrics like the *tail* of their waiting time, e.g. the fraction of jobs which are served within 100 ms of their arrival. While our analysis lends itself most easily to studying the average waiting time, it's conceivable that, with the right selection of stopping time and martingales, our results could be extended to higher moments and possibly even to a direct analysis of the tail.

Other natural policies. Another possible expansion on this work is studying other natural policies. For example, consider the following variation on the policy studied in this paper: instead of turning a server off instantly when there is no work for that server to do, we allow the server to sit idle for, say, 5 seconds before turning off. One would imagine that such a policy would have a lower average waiting time and a higher power consumption than this paper's setup policy. Moreover, one would anticipate that the tradeoff between the waiting time and the power consumption could be tuned by changing the amount of allowed idle time γ ; at $\gamma = 0$, we recover the policy studied in this paper, and as $\gamma \rightarrow \infty$, we recover the "Always On" policy. This class of policies, often referred to as the DelayedOff class of policies, have been studied extensively in the Exponential setup model due to their tunability, ease of implementation, and excellent empirical performance Gandhi et al. (2012), Pender and Phung-Duc (2016), Gandhi et al. (2010a), Gandhi and Harchol-Balter (2011); it would be a great contribution to extend our knowledge of these policies to settings with Deterministic setup times.

Tradeoff between waiting time and power. Another important avenue of study lies in characterizing the fundamental tradeoff between average waiting time and average power consumption. One point on the waiting-v.s.-power Pareto curve is clear: to obtain the minimum possible waiting time, one should keep all k servers on all the time. However, all the other points on the optimal tradeoff curve are unknown; furthermore, even if we knew a given (waiting time, power consumption) pair was feasible, it would be entirely unclear how we should go about achieving it. For example, it's conceivable that the DelayedOff class of policies mentioned previously could perform near-optimally as one varies the "idle time" parameter γ . A more principled understanding of the optimal tradeoff between waiting time and power consumption could allow us to directly derive a policy which achieves this optimal tradeoff, and could grant us greater insight into how to design more performant setup systems in the practical setting.

More general setup distributions. Lastly, we note that our results might be generalized to setup distributions beyond Deterministic. A promising step would be the establishment of closed form bounds for the *Exponential* setup time model; note that no closed-form result exists for the Exponential model, excluding the infinite server result of Gandhi et al. (2010b). With the establishment of an Exponential result, it seems likely that our arguments could be used to obtain a result for a Hyperexponential setup system, or even bootstrapped into a full phase-type result. Another way to generalize our setup model is to study a system where setup times take on two possible values (a short setup or a long setup); whether a setup is long or

short (and how the setup time of a server changes over time) could be used to model various real-world scenarios, e.g. where variable data loading times or intermittent network congestion. By generalizing our setup model in these ways, we could greatly expand the practical utility of our results and gain even deeper insight into what fundamentally governs the performance of systems with setup times.

References

- Allahverdi A, Soroush H (2008) The significance of reducing setup times/setup costs. *European journal of operational research* 187(3):978–984.
- Artalejo JR, Economou A, Lopez-Herrero MJ (2005) Analysis of a Multiserver Queue with Setup Times. *Queueing Syst.* 51(1):53–76.
- Behroozi M, Shirolkar S, Barik T, Parnin C (2020) Debugging hiring: What went right and what went wrong in the technical interview process. *Proceedings of the ACM/IEEE 42nd International Conference on Software Engineering: Software Engineering in Society*, 71–80.
- Bischof W (2001) Analysis of M/G/1-Queues with Setup Times and Vacations under Six Different Service Disciplines. *Queueing Syst.* 39(4):265–301.
- Choudhury G (1998) On a batch arrival Poisson queue with a random setup time and vacation period. *Comp. & Oper. Res.* 25(12):1013–1026.
- Cook JD (2024) Details for error bound on normal approximation to the Poisson distribution. https://www.johndcook.com/blog/berry_esseen_poisson/, accessed: 01/16/2024.
- Gandhi A, Doroudi S, Harchol-Balter M, Scheller-Wolf A (2013) Exact analysis of the M/M/k/setup class of Markov chains via Recursive Renewal Reward. *Queueing Syst.*, 153–166.
- Gandhi A, Gupta V, Harchol-Balter M, Kozuch M (2010a) Optimality analysis of energy-performance trade-off for server farm management. *Proc. Int. Symp. Computer Performance, Modeling, Measurements and Evaluation (IFIP Performance)* (Namur, Belgium).
- Gandhi A, Harchol-Balter M (2011) How Data Center Size Impacts the Effectiveness of Dynamic Power Management. *Proc. Ann. Allerton Conf. Communication, Control and Computing*, 1164–1169 (Urbana-Champaign, IL).
- Gandhi A, Harchol-Balter M (2013a) M/g/k with staggered setup. *Operations Research Letters* 41(4):317–320.
- Gandhi A, Harchol-Balter M (2013b) M/G/k with staggered setup. *Oper. Res. Lett.* 41(4):317–320.
- Gandhi A, Harchol-Balter M, Adan I (2010b) Server farms with setup costs. *Performance Evaluation* 67(11):1123–1138.
- Gandhi A, Harchol-Balter M, Kozuch M (2011) The case for sleep states in servers. *SOSP Workshop on Power-Aware Computing and Systems (HotPower)*, 1–5 (Cascais, Portugal).
- Gandhi A, Harchol-Balter M, Raghunathan R, Kozuch MA (2012) AutoScale: Dynamic, Robust Capacity Management for Multi-Tier Data Centers. *ACM Trans. Comput. Syst.* 30(4):1–26.

- Halfin S, Whitt W (1981) Heavy-traffic limits for queues with many exponential servers. *Oper. Res.* 29(3):567–588.
- Hao J, Jiang T, Wang W, Kim IK (2021) An empirical analysis of vm startup times in public iaas clouds. *2021 IEEE 14th International Conference on Cloud Computing (CLOUD)*, 398–403 (IEEE).
- He QM, Jewkes E (1995) Flow time in the MAP/G/1 queue with customer batching and setup times. *Stochastic Models* 11(4):691–711.
- Hong Y, Scully Z (2023) Performance of the gittins policy in the $g/g/1$ and $g/g/k$, with and without setup times. *ACM SIGMETRICS Performance Evaluation Review* 51(2):33–35.
- Hyytiä E, Down D, Lassila P, Aalto S (2018) Dynamic Control of Running Servers. *Int. Conf. Measurement, Modelling and Evaluation of Comput. Systems*, 127–141 (Erlangen, Germany: Springer).
- Kara A (2017) *Energy Consumption in Data Centers with Deterministic Setup Times*. Master's thesis, Middle East Technical University.
- Karatzas I, Shreve S (2012) *Brownian motion and stochastic calculus*, volume 113 (Springer Science & Business Media).
- Kingman JF (2009) The first Erlang century—and the next. *Queueing systems* 63(1):3.
- Kleinrock L (1975) *Queueing Systems* (John Wiley & Son).
- Liu X, Ying L (2022) Universal scaling of distributed queues under load balancing in the super-Halfin-Whitt regime. *IEEE/ACM Trans. Netw.* 30(1):190–201.
- Maheshwari S, Raychaudhuri D, Seskar I, Bronzino F (2018) Scalability and performance evaluation of edge cloud systems for latency constrained applications. *2018 IEEE/ACM Symposium on Edge Computing (SEC)*, 286–299 (IEEE).
- Mao M, Humphrey M (2012) A Performance Study on the VM Startup Time in the Cloud. *IEEE Int. Conf. Cloud Computing (CLOUD)*, 423–430 (Honolulu, HI).
- Mogul JC, Kompella RR (2015) Inferring the network latency requirements of cloud tenants. *15th Workshop on Hot Topics in Operating Systems (HotOS XV)*.
- Mukherjee D, Dhara S, Borst SC, van Leeuwen JS (2017) Optimal Service Elasticity in Large-Scale Distributed Systems. *Proc. ACM SIGMETRICS Int. Conf. Measurement and Modeling of Computer Systems* 1:1–28.
- Mukherjee D, Stolyar A (2019) Join Idle Queue with Service Elasticity: Large-Scale Asymptotics of a Nonmonotone System. *Stoch. Syst.* 9(4):338–358.
- Pender J, Phung-Duc T (2016) A law of large numbers for $m/m/c$ /delayoff-setup queues with nonstationary arrivals. *Int. Conf. on Analytical and Stochastic Modeling Techniques and Applications*, 253–268 (Cardiff, UK: Springer).
- Phung-Duc T (2017) Exact solutions for $M/M/c$ /setup queues. *Telecommun. Syst.* 64(2):309–324.
- Rzadca K, Findeisen P, Swiderski J, Zych P, Broniek P, Kusmieriek J, Nowak P, Strack B, Witusowski P, Hand S, et al. (2020) Autopilot: workload autoscaling at Google. *Proc. European Conf. Computer Systems (EuroSys)*, 1–16 (Heraklion, Crete, Greece).

Shevtsova I (2011) On the absolute constants in the Berry-Esseen type inequalities for identically distributed summands.

Tuttas C (2013) Travel nurses and 21st century nurse staffing. *Nurse Leader* 11(4):43–46.

Welch PD (1964) On a Generalized M/G/1 Queuing Process in Which the First Customer of Each Busy Period Receives Exceptional Service. *Oper. Res.* 12(5):736–752.

Williams JK, Harchol-Balter M, Wang W (2022) The M/M/k with Deterministic Setup Times. *Proc. ACM SIGMETRICS Int. Conf. Measurement and Modeling of Computer Systems* 6(3), URL <http://dx.doi.org/10.1145/3570617>.

Appendix A: Proof of Lemma 4, the Intervening Stopping Time (IST) Lemma.

LEMMA 4 (Intervening Stopping Time Lemma). *Given a starting stopping time T_0 , an ending stopping time P , and a collection of intervening stopping times $(T_i : i \in \mathbb{Z}^+)$, define the random variable F to be such that $T_F \leq P < T_{F+1}$. Now, given some time-varying random variable $Y_t \geq 0$ which is a function of the underlying Markov state of the system $\mathcal{S}(t)$, suppose that:*

1. $\mathbb{E} \left[\int_{T_0}^{\min(T_1, P)} Y_t dt \middle| \mathcal{F}_{T_0} \right] \leq G_0(\mathcal{S}(T_0))$,
2. $\mathbb{E} \left[\int_{T_i}^{\min(T_{i+1}, P)} Y_t dt \middle| \mathcal{F}_{T_i}, F \geq i \right] \leq G_i + B \cdot \mathbb{E} [\min(T_{i+1}, P) - T_i | \mathcal{F}_{T_i}, F \geq i] + C [T_i - T_{i-1}]$,
3. and $\Pr(F \geq i | \mathcal{F}_{T_i}, F \geq i-1) \leq 1 - p_i$,

where G_0 is also some function of the system state, and the G_i 's, the p_i 's, C , and B are all constants (possibly depending on system parameters). Then one can bound the time integral of Y_t from T_0 to P by

$$\mathbb{E} \left[\int_{T_0}^P Y_t dt \right] \leq \mathbb{E}[G_0(\mathcal{S}(T_0))] + \Pr(F > 0) \sum_{j=1}^{\infty} G_j \prod_{i=2}^j (1 - p_i) + (B + C) \cdot \mathbb{E}[P - T_0].$$

A.1. Proof of Lemma 4.

We begin with a manipulation of the integral, finding

$$\begin{aligned} \int_{T_0}^P Y_t dt &= \int_{T_0}^{\min(T_1, P)} Y_t dt + \sum_{i=1}^{\infty} \int_{\min(T_i, P)}^{\min(T_{i+1}, P)} Y_t dt \\ &= \int_{T_0}^{\min(T_1, P)} Y_t dt + \sum_{i=1}^{\infty} \mathbf{1}_{T_i < P} \int_{T_i}^{\min(T_{i+1}, P)} Y_t dt. \end{aligned}$$

Applying linearity of expectation and the tower property, we find that

$$\begin{aligned} &\mathbb{E} \left[\int_{T_0}^P Y_t dt \right] \\ &= \mathbb{E} \left[\mathbb{E} \left[\int_{T_0}^{\min(T_1, P)} Y_t dt \middle| \mathcal{F}_{T_0} \right] + \sum_{i=1}^{\infty} \mathbb{E} \left[\mathbb{E} \left[\mathbf{1}_{T_i < P} \int_{T_i}^{\min(T_{i+1}, P)} Y_t dt \middle| \mathcal{F}_{T_i} \right] \right] \right] \\ &= \mathbb{E} \left[\mathbb{E} \left[\int_{T_0}^{\min(T_1, P)} Y_t dt \middle| \mathcal{F}_{T_0} \right] + \sum_{i=1}^{\infty} \mathbb{E} \left[\mathbf{1}_{T_i < P} \mathbb{E} \left[\int_{T_i}^{\min(T_{i+1}, P)} Y_t dt \middle| \mathcal{F}_{T_i} \right] \right] \right]. \end{aligned}$$

Noting that the event $\{T_i < P\} = \{F \geq i\}$, we have

$$\begin{aligned}
&= \mathbb{E} \left[\mathbb{E} \left[\int_{T_0}^{\min(T_1, P)} Y_t dt \middle| \mathcal{F}_{T_0} \right] \right] + \sum_{i=1}^{\infty} \mathbb{E} \left[\mathbf{1}_{F \geq i} \mathbb{E} \left[\int_{T_i}^{\min(T_{i+1}, P)} Y_t dt \middle| \mathcal{F}_{T_i} \right] \right] \\
&\leq \mathbb{E} [G_0(S(T_0))] + \sum_{i=1}^{\infty} \mathbb{E} [\mathbf{1}_{F \geq i} (G_i + B \cdot \mathbb{E} [\min(T_{i+1}, P) - T_i | \mathcal{S}(T_i), F \geq i] + C [T_i - T_{i-1}])] \\
&= \mathbb{E} [G_0(S(T_0))] + B \cdot \mathbb{E} [P - T_0] + C \mathbb{E} [T_F - T_0] + \sum_{i=1}^{\infty} G_i \Pr(F \geq i) \\
&\leq \mathbb{E} [G_0(S(T_0))] + B \cdot \mathbb{E} [P - T_0] + C \mathbb{E} [P - T_0] + \sum_{i=1}^{\infty} G_i \Pr(F \geq i).
\end{aligned}$$

Applying our final assumption to bound $\Pr(F \geq i)$,

$$\begin{aligned}
\Pr(F \geq i) &= \Pr(F > 0) \prod_{j=2}^i \Pr(F \geq j | F \geq j-1) \\
&= \Pr(F > 0) \prod_{j=2}^i \mathbb{E} [\Pr(F \geq j | F \geq j-1, \mathcal{F}_{T_{j-1}})] \\
&\leq \Pr(F > 0) \prod_{j=2}^i \mathbb{E} [1 - p_j] \\
&= \Pr(F > 0) \prod_{j=2}^i (1 - p_j).
\end{aligned}$$

Applying this final result, we find

$$\mathbb{E} \left[\int_{T_0}^P Y_t dt \right] \leq \mathbb{E} [G_0(S(T_0))] + \Pr(F > 0) \sum_{j=1}^{\infty} G_j \prod_{i=2}^j (1 - p_i) + (B + C) \cdot \mathbb{E} [P - T_0],$$

as desired. \square

Appendix B: Proof of the Improved Lower Bound, Theorem 2.

B.1. The New Lower Bound: Proof Outline.

Basic Structure. We prove Theorem 2 via the MIST method. As noted in Section 5, we begin by applying the Renewal-Reward theorem to the queue length $Q(t)$, defining our renewal points as those points in time where the $(R + 1)$ -th server turns off. Defining time 0 to be one of these points, and defining the cycle time $X \triangleq \min\{t > 0 : Z(t^-) = R + 1, Z(t) = R\}$ as the next point, this gives

$$\mathbb{E} [Q(\infty)] = \frac{\mathbb{E} \left[\int_0^X Q(t) dt \right]}{\mathbb{E} [X]}.$$

To obtain our lower bound, it suffices to lower bound the numerator and upper bound the denominator of this fraction, i.e. lower bound $\mathbb{E} \left[\int_0^X Q(t) dt \right]$ and upper bound $\mathbb{E} [X]$. The time integral lower bound is handled by Lemma 5, which we state at the end of this section. The cycle length upper bound is split into two separate lemmas: Lemma 6 upper bounds the length of the cycle's "first part" and Lemma 7 bounds the length of its "second part".

Decomposition into phases. However, before we state or prove these lemmas, we first discuss the decomposition of the renewal cycle $[0, X)$ into two parts; one might think of this as a “miniature” application of the MIST method. We begin by noting that the end of the renewal cycle is moment when the $(R + 1)$ -th server turns off. Since the $(R + 1)$ -th server is off at the start of a renewal period, we can break the renewal cycle into two phases based on whether the $(R + 1)$ -th server has turned on yet. Formally, we define the accumulation time $T_A \triangleq \min \{t > 0 : Z(t) = R + 1\}$ as the first moment that the $(R + 1)$ -th server turns on. From here, we can focus separately on the accumulation phase, from time 0 to time T_A , and the draining phase, from time T_A to time X .

With this decomposition, we can now state our main lemmas. Their proofs follow in sequence afterwards.

LEMMA 5 (Lower bound on Cycle Integral). Define busy period integral $I^{\text{busy}}(x, z)$ as

$$I^{\text{busy}}(x, z) \triangleq \frac{x}{\mu z} \left[\frac{x+1}{2} + \frac{1}{1 - \frac{k\lambda}{k\lambda + \mu z}} \right] = \frac{x}{\mu z} \left[\frac{x+1}{2} + \frac{R}{z} \right].$$

For the time integral of the queue length $Q(t)$ over a renewal cycle, we have

$$\mathbb{E} \left[\int_0^X Q(t) dt \right] \geq \frac{1}{2} \mu \beta^2 L_1 \sqrt{R} + I^{\text{busy}} \left(\left[\mu \beta L_1 \sqrt{R} - (k - R) \right]^+, k - R \right).$$

LEMMA 6 (Upper bound on Accumulation Phase Length). Recall that

$$T_A \triangleq \min \{t > 0 : Z(t) \geq R + 1\}$$

is the amount of time until the $(R + 1)$ -th server turns on. Then we can bound the expectation $\mathbb{E}[T_A]$ by

$$\mathbb{E}[T_A] \leq e^{\frac{1}{24R\beta}} \left(\sqrt{1 + \frac{1}{2R\beta}} \right) \left[1 + \frac{e^{\frac{1}{12R}}}{\sqrt{2\mu\beta}} \right] \beta \leq 1.08 * \beta.$$

LEMMA 7 (Upper bound on Draining Phase Length). Recall that the accumulation time

$$T_A \triangleq \min \{t > 0 : Z(t) \geq R + 1\}$$

is the amount of time until the $(R + 1)$ -th server turns on and the cycle time X is the moment when it turns off. Then, one can bound $\mathbb{E}[X - T_A]$ by

$$\mathbb{E}[X - T_A] \leq \beta + \frac{1}{\mu} \frac{F_1 \beta \sqrt{R}}{k - R} + \frac{1}{\mu} \frac{3}{2} \ln(\beta) + \frac{1}{\mu} \ln(F_1 D_1) + \frac{2}{\mu} + \left[D_2 + \frac{D_3}{\sqrt{R}} \right] \max \left(\frac{1}{D_1 \sqrt{\mu\beta}}, \frac{1}{\sqrt{R}} \right),$$

where F_1 , D_1 , D_2 , and D_3 are constants not depending on system parameters.

B.2. Proof of Lemma 5: Lower Bound on Cycle Integral.

B.2.1. Lemma 5 Proof Outline

Basic Strategy. First, we split the first phase $[0, T_A)$ into epochs, where epoch i begins when the number of busy servers $Z(t)$ first drops to $R - i$, and an epoch ends either when the next epoch starts or when the first phase ends. Our goal will be to analyze a specific “significant” epoch. In particular, we say that an epoch is long if it lasts for longer than a setup time β . Because the accumulation phase ends when the $(R + 1)$ -th server turns on, at least one epoch must be long. We use L to denote the index of the *first* long epoch. From here, we argue via a martingale/coupling argument that the expected time integral over the first β time in epoch L is at least $\frac{1}{2}\beta^2\mathbb{E}[L]$. To bound the integral afterwards, we couple the behavior of the total number of jobs $N(t)$ to the queue length in an M/M/1 queue with arrival rate $k\lambda$ and departure rate $k\mu$.

Formalization. Define the stopping time $\tau_i \triangleq \min\{t \geq 0 : N(t) \leq R - i\}$ as the beginning of epoch i . We say that the epoch *occurs* is $\tau_i < T_A$, and define the end of epoch i as $\gamma_i \triangleq \min(\tau_{i+1}, T_A)$ the moment when either epoch $i + 1$ begins or when the first phase ends. If epoch i occurs, we say it is long if $\gamma_i - \tau_i \geq \beta$. Let $L \triangleq \min\{i \in \mathbb{N} : \gamma_i - \tau_i \geq \beta\}$ be the index of the first long epoch. It suffices to show two claims; we state and prove them in sequence.

B.3. Lower Bound on Integral until $\tau_L + \beta$.

We show the following claim.

CLAIM 15. *Let L be the index of the first long epoch. Then,*

$$\mathbb{E} \left[\int_0^{\tau_L + \beta} Q(t) dt \right] \geq \frac{1}{2} \mu \beta^2 L_1 \sqrt{R}, \quad (16)$$

where L_1 is some absolute constant.

Claim 15 Proof Strategy. First, we show that the initial integral is bounded by

$$\mathbb{E} \left[\int_0^{\tau_L + \beta} Q(t) dt \right] \geq \frac{1}{2} \mu \beta^2 \mathbb{E}[L]. \quad (17)$$

Afterwards, we give a bound on $\mathbb{E}[L]$, showing that

$$\mathbb{E}[L] \geq L_1 \sqrt{R}. \quad (18)$$

Proof of (17), Bound in terms of $\mathbb{E}[L]$.

To show (17), we first condition on whether $L \geq i$, giving

$$\begin{aligned} \mathbb{E} \left[\int_0^{\tau_L + \beta} Q(t) dt \right] &= \sum_{i=0}^{\infty} \mathbb{E} \left[\int_{\tau_i}^{\min(\tau_i + \beta, \tau_{i+1})} Q(t) dt \mathbf{1}_{L \geq i} \right] \\ &= \sum_{i=0}^{\infty} \mathbb{E} \left[\int_{\tau_i}^{\min(\tau_i + \beta, \tau_{i+1})} Q(t) dt \middle| \mathcal{F}_{\tau_i} \right] \Pr(L \geq i). \end{aligned}$$

To further develop this conditional expectation, we note that during the interval $[\tau_i, \min(\tau_i + \beta, \tau_{i+1}))$, the system must have exactly $Z(t) = R - i$ busy servers running, meaning that $Q(t) = N(t) - (R - i)$. Defining a coupled process $\tilde{Q}(t)$ as

$$\tilde{Q}(t) = A(\tau_i, t) - \mathcal{D}[R - i](\tau_i, t),$$

we see that $Q(t)$ and $\tilde{Q}(t)$ coincide during the interval in question. Moreover, one can redefine the stopping time $\gamma = \tau_{i+1}$ as $\min\{t > \tau_i : \tilde{Q}(t) = -1\}$. Noting that $Q(\min(\gamma, t)) = -1$ for any time $t > \gamma$, we find that

$$\begin{aligned} \int_{\tau_i}^{\min(\tau_i + \beta, \tau_{i+1})} Q(t) dt &= \int_{\tau_i}^{\min(\tau_i + \beta, \tau_{i+1})} \tilde{Q}(t) dt \\ &= \int_{\tau_i}^{\min(\tau_i + \beta, \tau_{i+1})} \tilde{Q}(\min(t, \tau_{i+1})) dt + \int_{\min(\tau_i + \beta, \tau_{i+1})}^{\tau_i + \beta} (\tilde{Q}(\min(t, \tau_{i+1})) + 1) dt \\ &= \int_{\tau_i}^{\tau_i + \beta} \tilde{Q}(\min(t, \tau_{i+1})) dt + [\beta - \min(\beta, \tau_{i+1} - \tau_i)] \\ &\geq \int_{\tau_i}^{\tau_i + \beta} \tilde{Q}(\min(t, \tau_{i+1})) dt. \end{aligned}$$

Taking the conditional expectation at time τ_i , we find

$$\mathbb{E} \left[\int_{\tau_i}^{\tau_i + \beta} \tilde{Q}(\min(t, \tau_{i+1})) dt \middle| \mathcal{F}_{\tau_i} \right] = \int_{\tau_i}^{\tau_i + \beta} \mathbb{E} \left[\tilde{Q}(\min(t, \tau_{i+1})) \middle| \mathcal{F}_{\tau_i} \right] dt.$$

Noting that $V_L(t) = \tilde{Q}(t) - \mu i [t - \tau_i]$ is a martingale, and that $\min(t, \tau_{i+1})$ is an almost-surely bounded stopping time, we have that

$$\begin{aligned} \tilde{Q}(\tau_i) &= V_L(\tau_i) = 0 \\ &= \mathbb{E}[V_L(\min(t, \tau_{i+1})) | \mathcal{F}_{\tau_i}] \\ &= \mathbb{E} \left[\tilde{Q}(\min(t, \tau_{i+1})) \middle| \mathcal{F}_{\tau_i} \right] - \mu i \mathbb{E}[\min(t, \tau_{i+1}) | \mathcal{F}_{\tau_i}]. \end{aligned}$$

Since

$$\mathbb{E}[\min(t, \tau_{i+1}) | \mathcal{F}_{\tau_i}] \geq t \cdot \Pr(\tau_{i+1} - \tau_i \geq t) \geq t \cdot \Pr(\tau_{i+1} - \tau_i \geq \beta) = t \Pr(L = i | L \geq i),$$

we have

$$\begin{aligned} \Pr(L \geq i) \mathbb{E} \left[\int_{\tau_i}^{\min(\tau_i + \beta, \tau_{i+1})} Q(t) dt \middle| L \geq i \right] &\geq \Pr(L \geq i) \mathbb{E} \left[\int_{\tau_i}^{\tau_i + \beta} \tilde{Q}(\min(t, \tau_{i+1})) dt \middle| L \geq i \right] \\ &\geq \int_{\tau_i}^{\tau_i + \beta} \mu i t \Pr(L = i) \\ &= \mu \frac{\beta^2}{2} i \Pr(L = i) \end{aligned}$$

Summing across all i , we obtain (17).

B.3.1. Proof Sketch for (18), bound on $\mathbb{E}[L]$. We defer the full proof of this to Section G.6, and for now give a proof sketch.

We prove (18) by first showing that

$$\Pr(L > j | L \geq j) \geq \left(1 - \frac{j}{R}\right) \left(1 - \frac{b_1}{\sqrt{\mu\beta R}}\right),$$

where $b_1 = \frac{2}{\sqrt{\pi}}$. Next, we show that this implies that, for any $\delta \in (0, 1)$ and any $j < \delta R$,

$$\Pr(L > j) \geq \left(1 - \frac{b_1}{\sqrt{\mu\beta R}}\right)^{j+1} e^{-\frac{j(j+1)}{2R} \frac{1}{1-\delta}}.$$

From here, we use the sum of tails formula $\mathbb{E}[L] = \sum_{j=0}^{\infty} \Pr(L > j)$ to show

$$\mathbb{E}[L] \geq \left(1 - \frac{b_1}{\sqrt{\mu\beta R}}\right) \left(\left[\sqrt{\frac{\pi}{2}}(1-\delta) - \frac{1.15(1-\delta)}{\sqrt{\mu\beta}} \right] \sqrt{R} - \frac{1}{2} - \frac{2(1-\delta)}{\delta} e^{-R \frac{\delta^2}{1-\delta}} \right).$$

Choosing $\delta = \frac{2}{\sqrt{R}}$ then noting that $\mu\beta \geq 100$ and $R \geq 100$ gives the result.

B.4. Lower Bound on Integral after $\tau_L + \beta$.

To finish our lower bound on the integral, we now show the following claim.

CLAIM 16. *Let L be the index of the first long epoch. Then,*

$$\mathbb{E} \left[\int_{\tau_L + \beta}^X Q(t) dt \right] \geq I^{\text{busy}} \left(\left[\mu\beta L_1 \sqrt{R} - (k - R) \right]^+, k - R \right) \quad (19)$$

where L_1 is some absolute constant.

Claim 16: Proof Strategy. First, we show that the remaining integral is bounded by

$$\mathbb{E} \left[\int_{\tau_L + \beta}^X Q(t) dt \right] \geq I^{\text{busy}} \left(\left[\mathbb{E}[N(\tau_L + \beta)] - k \right]^+, k - R \right). \quad (20)$$

Then, we use martingales again to show that

$$\mathbb{E}[N(\tau_L + \beta)] \geq R + \mu\beta \mathbb{E}[L]. \quad (21)$$

Applying (18), our bound on $\mathbb{E}[L]$, we obtain the result.

B.4.1. Proof of (20), Bound in terms of $\mathbb{E}[N(\cdot)]$.

To prove (20), we make a simple coupling argument. Let $\eta_k \triangleq \min \{t \geq \tau_L + \beta : N(t) \leq k\}$. Since the draining phase starts at $T_A \geq \tau_L + \beta$ and the end of the cycle $X = \min \{t \geq T_A : N(t) \leq R\}$, we know that $X \geq \eta_k$. Moreover, we know the number of busy servers $Z(t) \leq k$; it follows by Claim 1 that we can define $\tilde{N}(t)$ as

$$\tilde{N}(t) \triangleq N(\tau_L + \beta) + A(\tau_L + \beta, t) - \mathcal{D}[k]((\tau_L + \beta, t))$$

and have $\tilde{N}(t) \leq N(t)$ for any $t > \tau_L + \beta$. Even further, we can defined a coupled hitting time $\tilde{\eta}_k \triangleq \min \left\{ t > \tau_L + \beta : \tilde{N}(t) \leq k \right\}$ which must happen before η_k . In other words,

$$\begin{aligned} \int_{\tau_L + \beta}^X Q(t) dt &\geq \int_{\tau_L + \beta}^{\eta_k} Q(t) dt \\ &\geq \int_{\tau_L + \beta}^{\eta_k} [N(t) - k] dt \\ &\geq \int_{\tau_L + \beta}^{\tilde{\eta}_k} [\tilde{N}(t) - k] dt. \end{aligned}$$

This final term is just the time integral of the number of jobs in system over a M/M/1 busy period started by $[N(\tau_L + \beta) - k]^+$ jobs, where jobs arrive at rate $k\lambda$ and depart at rate $k\mu$. Accordingly, we have

$$\begin{aligned} \mathbb{E} \left[\int_{\tau_L + \beta}^X Q(t) dt \right] &\geq \mathbb{E} \left[I^{\text{busy}} \left([N(\tau_L + \beta) - k]^+, k - R \right) \right] \\ &\geq I^{\text{busy}} \left(\mathbb{E} [N(\tau_L + \beta) - k]^+, k - R \right) \\ &\geq I^{\text{busy}} \left(\mathbb{E} \left[[N(\tau_L + \beta) - R - (k - R)]^+ \right], k - R \right), \end{aligned}$$

where in the last two lines we have applied Jensen's inequality. \square

B.4.2. Proof of (21), Bound on $\mathbb{E}[N(\cdot)]$. To bound $\mathbb{E}[N(\tau_L + \beta)]$, we condition on the value of L , then make a martingale argument.

$$\begin{aligned} \mathbb{E} [N(\tau_L + \beta)] &= \sum_{i=0}^R \mathbb{E} [N(\tau_i + \beta) \mathbf{1}_{L=i}] \\ &= \sum_{i=0}^R \mathbb{E} [N(\tau_i + \beta) \mathbf{1}_{L=i, L \geq i}] \\ &\geq \sum_{i=0}^R \Pr(L \geq i) \mathbb{E} [N(\tau_i + \beta) \mathbf{1}_{L=i} | L \geq i]. \end{aligned}$$

Continuing with this conditional expectation,

$$\begin{aligned} &\mathbb{E} [N(\tau_i + \beta) \mathbf{1}_{L=i} | \mathcal{F}_{\tau_i}] \\ &= \mathbb{E} [N(\tau_i + \beta) \mathbf{1}_{\tau_i + \beta < \tau_{i+1}} | \mathcal{F}_{\tau_i}] \\ &= \mathbb{E} [[N(\tau_i + \beta) - (R - (i + 1))] \mathbf{1}_{\tau_i + \beta < \tau_{i+1}} | \mathcal{F}_{\tau_i}] + (R - i - 1) \Pr(\tau_i + \beta < \tau_{i+1} | \mathcal{F}_{\tau_i}) \\ &= \mathbb{E} [N(\min(\tau_i + \beta, \tau_{i+1})) - (R - (i + 1)) | \mathcal{F}_{\tau_i}] + (R - i - 1) \Pr(\tau_i + \beta < \tau_{i+1} | \mathcal{F}_{\tau_i}) \\ &= 1 + \mu i \mathbb{E} [\min(\beta, \tau_{i+1} - \tau_i)] + (R - i - 1) \Pr(\tau_i + \beta < \tau_{i+1}) \\ &\geq 1 + \mu i \beta \Pr(\tau_{i+1} - \tau_i \geq \beta) + (R - i - 1) \Pr(\tau_i + \beta < \tau_{i+1}) \\ &= 1 + \mu i \beta \Pr(L = i | L \geq i) + (R - i - 1) \Pr(L = i | L \geq i). \end{aligned} \tag{22}$$

Summing across i , we find

$$\begin{aligned}\mathbb{E}[N(\tau_L + \beta)] &= \sum_{i=0}^R \Pr(L \geq i) \quad ((22)) \\ &= (1 + \mathbb{E}[L]) + (\mu\beta\mathbb{E}[L]) + (R - \mathbb{E}[L] - 1) \\ &= R + \mu\beta\mathbb{E}[L],\end{aligned}$$

as desired. □

Combining Claims 15 and 16, we obtain a lower bound on $\mathbb{E}\left[\int_0^X Q(t)dt\right]$, proving Lemma 5.

B.5. Proof of Lemma 6: Upper Bound on the Accumulation Time $\mathbb{E}[T_A]$.

Defining a coupling. To prove Lemma 6, we first note that, during the accumulation phase, we have two bounds on the number of busy servers $Z(t)$: it must be less than the total number of jobs $N(t)$ and it must be less than R ; the former because every busy server must be working on a job, and the latter because otherwise the accumulation phase would be over. Thus, we can define a coupled $M/M/R$ system for which the number of jobs $\tilde{N}(t)$ in the coupled system is always at least the number of jobs $N(t)$ in the original system.

How we use the coupling. To use this coupled process to bound $\mathbb{E}[T_A]$, recall that the accumulation point T_A is the first time the $(R+1)$ -th server turns on. Accordingly, one can also think of this as the first time that there has been at least $R+1$ jobs in the system for β time. Thus, if we define a coupled accumulation point $\tilde{T}_A \triangleq \min\left\{t \geq \beta : \min_{s \in [t-\beta, t]} \tilde{N}(s) \geq R+1\right\}$, then we know $\tilde{T}_A \geq T_A$. In other words, it suffices to bound $\mathbb{E}[\tilde{T}_A]$.

General Strategy. We bound $\mathbb{E}[\tilde{T}_A]$ using the MIST method of Lemma 4. As such, we define a few stopping times, then list the preconditions/claims that we will satisfy to complete the proof of Lemma 6.

Definition of γ and α . Let the initial cycle-downcrossing occur at $\alpha_0 \triangleq 0$ and iteratively define the upcrossings γ and downcrossings α as

$$\gamma_i \triangleq \min\left\{t \geq \alpha_i : \tilde{N}(t) \geq R+1\right\}$$

and

$$\alpha_{i+1} \triangleq \min\left\{t \geq \gamma_i : \tilde{N}(t) \geq R+1\right\}.$$

Application of Lemma 4, the IST Lemma. Applying Lemma 4 using $0 = \alpha_0$ as our starting point, the coupled accumulation point \tilde{T}_A as our ending point, our test function as $Y_t = 1$, and the cycle-upcrossings (α_i) as our intervening stopping times, we now must prove that

$$\mathbb{E}[\gamma_i - \alpha_i | n_\alpha \geq i] \leq \frac{1}{\mu} e^{\frac{1}{12R}} \sqrt{1 + \frac{1}{R} \frac{\sqrt{2\pi}}{\sqrt{R}}} \leq \frac{c_3}{\mu\sqrt{R}}, \quad (23)$$

$$\mathbb{E} \left[\min \left(\tilde{T}_A, \alpha_{i+1} \right) - \gamma_i \mid n_\alpha \geq i \right] \leq b_1 \sqrt{\frac{\beta}{\mu R}} + \frac{6}{\mu R}, \quad (24)$$

and

$$\Pr(n_\alpha \geq i+1 \mid n_\alpha \geq i) \leq 1 - \frac{b_1}{\sqrt{2}} e^{-\frac{1}{3(\mu^2 R \beta - 1)}} \frac{1}{\sqrt{\mu^2 R \beta + 2}} \leq 1 - \frac{b_1 c_4}{\sqrt{\mu R \beta}}, \quad (25)$$

where $b_1 \triangleq \sqrt{\frac{2}{\pi}}$, $c_3 = 1.001\sqrt{2\pi}$, and $c_4 = 0.499$.

Completion of Proof, assuming (23), (24), and (25). Applying Lemma 4, one finds that

$$\begin{aligned} \mathbb{E} [\tilde{T}_A] &= \sum_{i=0}^{\infty} \mathbb{E} \left[\min \left(\tilde{T}_A, \alpha_{i+1} \right) - \alpha_i \mid n_\alpha \geq i \right] \Pr(n_\alpha \geq i) \\ &\leq \left[\frac{c_3}{\mu\sqrt{R}} + \frac{b_1\sqrt{\beta}}{\sqrt{\mu R}} + \frac{6}{\mu R} \right] \frac{\sqrt{\mu R \beta}}{b_1 c_4} \\ &= \frac{1}{\mu} \left[\frac{c_3}{b_1 c_4} \sqrt{\mu \beta} + \frac{1}{c_4} \beta + \frac{6}{b_1 c_4} \sqrt{\frac{\mu \beta}{R}} \right]. \end{aligned}$$

B.5.1. Proof of (23): Upper bound on initial up-crossing time. To prove (23), we note that, since our coupled system is an $M/M/R$, the expected time $\mathbb{E}[\gamma_i - \alpha_i \mid n_\alpha \geq i]$ is simply the expected passage time from state R to $(R+1)$ in an $M/M/R$ (and equivalently an $M/M/R/(R+1)$, an $M/M/R$ which can contain only $R+1$ jobs. Solving, one finds that

$$\begin{aligned} \mathbb{E} [T_{R \rightarrow (R+1)}] &\leq \mathbb{E} [T_{(R+1) \rightarrow (R+1)}] \\ &= \frac{1}{\mu(R+1)} \frac{1}{\pi_{R+1}} \\ &= \frac{1}{\mu(R+1)} \frac{\sum_{i=0}^{R+1} \frac{R^i}{i!}}{\frac{R^{R+1}}{(R+1)!}} \\ &\leq \frac{1}{\mu(R+1)} e^R \frac{(R+1)!}{R^{R+1}} \\ &\leq \frac{1}{\mu(R+1)} e^R \frac{e^{\frac{1}{12(R+1)}} \sqrt{2\pi(R+1)} (R+1)^{R+1} e^{-(R+1)}}{R^{R+1}} \quad (26) \end{aligned}$$

$$\begin{aligned} &= e^{\frac{1}{12(R+1)}} \frac{1}{\mu} \sqrt{2\pi} \frac{\sqrt{R+1}}{R} \left(1 + \frac{1}{R}\right)^R e^{-1} \\ &\leq \frac{1}{\mu} e^{\frac{1}{12(R+1)}} \sqrt{1 + \frac{1}{R}} \frac{\sqrt{2\pi}}{\sqrt{R}} \quad (27) \\ &\leq \frac{1}{\mu} 1.006 \frac{\sqrt{2\pi}}{\sqrt{R}} \\ &\triangleq \frac{c_3}{\mu\sqrt{R}}, \end{aligned}$$

where in (26) we have applied Stirling's approximation and in (27) we have used that $R \geq 100$.

B.5.2. Proof of (24): Bound on time between up-crossings. To bound the expected time $\mathbb{E} \left[\min \left(\tilde{T}_A, \alpha_{i+1} \right) - \gamma_i \mid n_\alpha \geq i \right]$, we first note that, if $\gamma_i + \beta \leq \alpha_{i+1}$, then $\tilde{T}_A = \gamma_i + \beta$. Likewise, if $\gamma_i + \beta > \alpha_{i+1}$, then $\tilde{T}_A > \alpha_{i+1}$. It follows that, given that $n_\alpha \geq i$, the time $\min \left(\tilde{T}_A, \alpha_{i+1} \right) = \min \left(\beta + \gamma_i, \alpha_{i+1} \right)$. Thus, we have that

$$\mathbb{E} \left[\min \left(\tilde{T}_A, \alpha_{i+1} \right) - \gamma_i \mid n_\alpha \geq i \right] = \mathbb{E} \left[\min \left(\beta, \alpha_{i+1} - \gamma_i \right) \mid n_\alpha \geq i \right] = \int_0^\beta \Pr \left(\alpha_{i+1} - \gamma_i > s \mid n_\alpha \geq i \right) ds.$$

We continue by bounding this tail probability. To begin, note that, while $\tilde{N}(t)$ stays above $R + 1$, the dynamics of \tilde{N} are precisely that of a critically-loaded $M/M/1$ queue with arrival rate and departure rate equal to $k\lambda$. The tail probability we are interested in bounding is precisely the probability that a busy period (started with 1 job) in such a system lasts longer than s time. Applying Claim 20, one finds that, for any $t \geq \frac{3}{\mu 2R}$,

$$\Pr \left(\alpha_{i+1} - \gamma_i > s \mid n_\alpha \geq i \right) \leq b_1 \left(\frac{1}{\sqrt{\mu 2R s}} + \frac{b_2}{(\mu 2R s)^{3/2}} \right).$$

Integrating, we find that

$$\begin{aligned} \int_0^\beta \Pr \left(\alpha_{i+1} - \gamma_i > s \mid n_\alpha \geq i \right) ds &\leq \frac{3}{\mu 2R} + \frac{b_1}{\sqrt{2}} \int_{\frac{3}{\mu 2R}}^\beta \frac{1}{\sqrt{\mu 2R s}} + \frac{b_2}{(\mu 2R s)^{3/2}} ds \\ &\leq \frac{3}{\mu 2R} + \frac{b_1}{\sqrt{2}} \left[\sqrt{\frac{2\beta}{\mu R}} + b_2 \sqrt{\frac{2}{3}} \frac{1}{\mu R} \right] \\ &= b_1 \sqrt{\frac{\beta}{\mu R}} + \left(\frac{3}{2} + (b_1 + 2.5) \sqrt{\frac{2}{3}} \right) \frac{1}{\mu R} \\ &\leq \frac{2}{\sqrt{\pi}} \sqrt{\frac{\beta}{\mu R}} + \frac{6}{\mu R}. \end{aligned}$$

□

B.5.3. Proof of (25): Bound on probability of another γ up-crossing. To prove (25), it suffices to note that, upon conditioning on the filtration at γ_i , the probability $\Pr \left(n_\alpha \geq i + 1 \mid n_\alpha \geq i \right)$ is simply the probability that a busy period in a critically-loaded $M/M/1$, with arrival and departure rate equal to μR , ends before β time has passed. Applying Claim 20, one finds that this is

$$\Pr \left(n_\alpha \geq i + 1 \mid n_\alpha \geq i \right) \geq 1 - \frac{b_1}{\sqrt{2}} e^{-\frac{1}{3(\mu 2R\beta - 1)}} \frac{1}{\sqrt{\mu 2R\beta + 2}}.$$

□

B.6. Proof of Lemma 7: Upper Bound on the Remaining Cycle Time $\mathbb{E}[X - T_A]$.

We now prove the upper bound on $\mathbb{E}[X - T_A]$. We make use of the “wait-busy” idea from Section 6.4.2 as well as our main tool, Lemma 4. As such, we begin by defining some stopping times.

Definition of $v_i^{(\text{down})}$ and $v_i^{(\text{up})}$. Recall that the draining phase begins at time T_A . Let $M_L \triangleq \min\left(k - R, \max\left(\frac{\sqrt{R}}{D_1\sqrt{\beta}}, 1\right)\right)$ be a specially-set analysis threshold. Let the stopping time $v_1^{(\text{down})} \triangleq \min\{t \geq T_A : N(t) < R + M_L\}$ be the first time the number of jobs $N(t)$ drops below $R + M_L$, and recursively define

$$v_i^{(\text{up})} \triangleq \min\left\{t \geq v_i^{(\text{down})} : N(t) \geq R + M_L\right\}$$

and

$$v_{i+1}^{(\text{down})} \triangleq \min\left\{t \geq v_i^{(\text{up})} : N(t) < R + M_L\right\}.$$

Specification Step. Now, we apply Lemma 4 using the accumulation point T_A as our initial point, the cycle end X as our ending point, the constant function $Y_t = 1$ as our test function, and the draining-downcrossing points $(v_i^{(\text{down})})$ as our intervening points; we use n_ζ to count the number of intervening points. To complete the proof, we must show that the following claims:

$$\mathbb{E}\left[v_1^{(\text{down})} - T_A\right] \leq \beta + \frac{1}{\mu} \frac{F_1\beta\sqrt{R}}{k-R} + \frac{1}{\mu} \frac{3}{2} \ln(\beta) + \frac{1}{\mu} \ln(F_1 D_1), \quad (28)$$

$$\mathbb{E}\left[\min\left(X, v_{i+1}^{(\text{down})}\right) - v_i^{(\text{down})} \mid n_\zeta \geq i\right] \leq \frac{D_2}{\mu\sqrt{R}} + \frac{D_3}{\mu R} + \frac{2}{\mu M_L} \quad (29)$$

$$\Pr(n_\zeta \geq i+1 \mid n_\zeta \geq i) \leq \frac{1}{M_L}. \quad (30)$$

Completion of Proof assuming (28), (29), and (30). Before proving the claims, we now prove the lemma. It suffices to give a bound on $\mathbb{E}\left[X - v_1^{(\text{down})}\right]$; applying Lemma 4 gives

$$\begin{aligned} \mathbb{E}\left[X - v_1^{(\text{down})}\right] &\leq M_L \left[\frac{D_2}{\mu\sqrt{R}} + \frac{D_3}{\mu R} + \frac{2}{\mu M_L} \right] \\ &= \frac{2}{\mu} + \left[D_2 + \frac{D_3}{\sqrt{R}} \right] \max\left(\frac{1}{D_1\sqrt{\mu\beta}}, \frac{1}{\sqrt{R}}\right). \end{aligned}$$

B.6.1. Proof of (28): Upper bound on time until first downward visit. To bound $\mathbb{E}\left[v_1^{(\text{down})} - T_A\right]$, we make a coupling argument then apply basic results on M/M/1 busy periods. Moreover, instead of proving (28) directly, we first show a more general claim.

CLAIM 17. For $M_L \leq j \leq N(T_A) - R$, define η_j as the first time after T_A that $N(t) \leq R + j$. Note that this means that $\eta_{N(T_A)} = T_A$ and $\eta_{M_L} = v_1^{(\text{down})}$. Then we have the following bound:

$$\mathbb{E}\left[\eta_{M_L} - \eta_j \mid \mathcal{F}_{\eta_j}\right] \leq Y_{R+j}(\eta_j) + \frac{1}{\mu} \sum_{i=M_L}^j \frac{1}{\min(i, k-R)}.$$

Afterwards, we complete the proof by noting that $[N(T_A) - k]^+ \leq [N(T_A) - R]$, taking expectations, applying Jensen's inequality to the minimum function and the $\ln(\cdot)$ (which is concave), using the bound on $\mathbb{E}[N(T_A) - R]$ from Claim 14, then letting $h = M_L$.

Proof of Claim 17. We prove Claim 17 by induction. In the base case, suppose that $j = M_L + 1$. Note that at time η_{M_L+1} , the numbers of jobs $N(\eta_{M_L+1} = R + M_L + 1$ and the remaining time until the $(R + M_L + 1)$ -th server turns on is $Y_{R+M_L+1}(\eta_{M_L+1})$. As such, we can simply wait until either that server turns on, in which case we can analyze the system as an M/M/1 busy period with departure rate $\mu \min(R + M_L + 1, k)$, or the number of jobs $N(t)$ drops below $R + M_L + 1$ on its own. In other words, (using j here to save space)

$$\mathbb{E}[\eta_{j-1} - \eta_j | \mathcal{F}_{\eta_j}] \leq Y_{R+j}(\eta_j) + \frac{\mathbb{E}\left[\left[N(\eta_j + Y_{R+j}(\eta_j)) - (R + (j-1))\right] \mathbf{1}_{\eta_{j-1} > \eta_j + Y_{R+j}(\eta_j)} \middle| \mathcal{F}_{\eta_j}\right]}{\mu \min(j, k - R)}.$$

Now, we reframe the expectation as an expectation up to a stopping time. We note that, if $\eta_{j-1} > \eta_j + Y_{R+j}(\eta_j)$, then we have that

$$N(\eta_j + Y_{R+j}(\eta_j)) = N(\min(\eta_j + Y_{R+j}(\eta_j), \eta_{j-1})).$$

Likewise, if $\eta_{j-1} \leq \eta_j + Y_{R+j}(\eta_j)$, then

$$R + j - 1 = N(\eta_{j-1}) = N(\min(\eta_j + Y_{R+j}(\eta_j), \eta_{j-1})).$$

Using this and applying a simple coupling argument, one sees that

$$\begin{aligned} & \mathbb{E}\left[\left[N(\eta_j + Y_{R+j}(\eta_j)) - (R + (j-1))\right] \mathbf{1}_{\eta_{j-1} > \eta_j + Y_{R+j}(\eta_j)} \middle| \mathcal{F}_{\eta_j}\right] \\ &= \mathbb{E}\left[N(\min(\eta_j + Y_{R+j}(\eta_j), \eta_{j-1})) - (R + j - 1) \middle| \mathcal{F}_{\eta_j}\right] \\ &\leq N(\eta_j) - (R + j - 1) = 1. \end{aligned}$$

Thus, we find that

$$\mathbb{E}[\eta_{j-1} - \eta_j | \mathcal{F}_{\eta_j}] \leq Y_{R+j}(\eta_j) + \frac{1}{\mu \min(j, k - R)}.$$

Inductive case. The inductive case proceeds in much the same way, except now, if $N(t)$ does drop below $R + j$ “early”, then we can factor in the time that has elapsed in the value of $Y_{R+j}(\eta_j)$. In particular, note that, since the $(R + j)$ -th server would have already turned on,

$$\mathbb{E}[\eta_{M_L} - \eta_j | \mathcal{F}_{\eta_j}] \mathbf{1}_{\eta_j \geq \eta_{j+1} + Y_{R+j+1}(\eta_{j+1})} \leq \frac{1}{\mu} \sum_{i=M_L}^j \frac{1}{\mu \min(i, k - R)} \mathbf{1}_{\eta_j \geq \eta_{j+1} + Y_{R+j+1}(j+1)}.$$

It follows that

$$\mathbb{E}[\eta_{M_L} - \eta_j | \mathcal{F}_{\eta_j}] \leq Y_{R+j}(\eta_j) \mathbf{1}_{\eta_j < \eta_{j+1} + Y_{R+j+1}(\eta_{j+1})} + \frac{1}{\mu} \sum_{i=M_L}^j \frac{1}{\mu \min(i, k - R)}.$$

Now, we note that

$$\begin{aligned} Y_{R+j}(\eta_j) \mathbf{1}_{\eta_j < \eta_{j+1} + Y_{R+1+j}(\eta_{j+1})} &= [Y_{R+j}(\eta_j) + \eta_j - \eta_j] \mathbf{1}_{\eta_j < \eta_{j+1} + Y_{R+1+j}(\eta_{j+1})} \\ &= [Y_{R+j}(\eta_{j+1}) + \eta_{j+1} - \eta_j] \mathbf{1}_{\eta_j < \eta_{j+1} + Y_{R+1+j}(\eta_{j+1})} \\ &\leq [Y_{R+j+1}(\eta_{j+1}) + \eta_{j+1} - \eta_j] \mathbf{1}_{\eta_j < \eta_{j+1} + Y_{R+1+j}(\eta_{j+1})} \\ &= [Y_{R+j+1}(\eta_{j+1}) + \eta_{j+1} - \eta_j]^+, \end{aligned}$$

so that we find

$$\mathbb{E} [\eta_{M_L} - \eta_j | \mathcal{F}_{\eta_j}] \leq [Y_{R+j+1}(\eta_{j+1}) + \eta_{j+1} - \eta_j]^+ + \frac{1}{\mu} \sum_{i=M_L}^j \frac{1}{\mu \min(i, k-R)}.$$

Finally, we note that

$$\mathbb{E} [\eta_j - \eta_{j+1} | \mathcal{F}_{\eta_{j+1}}] \leq \mathbb{E} [\min(\eta_j - \eta_{j+1}, Y_{R+j+1}(\eta_{j+1})) | \mathcal{F}_{\eta_{j+1}}] + \frac{1}{\mu \min(j+1, k-R)}.$$

Summing these final two expressions gives the inductive result, proving Claim 17.

Using Claim 17. Thus, we obtain that, using H_i to denote the i -th harmonic number,

$$\begin{aligned} \mathbb{E} [v_1^{(\text{down})} - T_A | \mathcal{F}_{T_A}] &\leq \beta + \frac{1}{\mu} \frac{[N(T_A) - k]^+}{k-R} + \frac{1}{\mu} [H_{\min(N(T_A)-R, k-R)} - H_{M_L}] \\ &\leq \beta + \frac{1}{\mu} \frac{[N(T_A) - R]^+}{k-R} + \frac{1}{\mu} \ln \left(\frac{\min(N(T_A) - R, k-R)}{M_L} \right). \end{aligned}$$

Taking expectations and applying Jensen's inequality twice, we find

$$\begin{aligned} \mathbb{E} [v_1^{(\text{down})} - T_A | \mathcal{F}_{T_A}] &\leq \beta + \frac{1}{\mu} \frac{F_1 \mu \beta \sqrt{R}}{k-R} + \frac{1}{\mu} \ln \left(\frac{F_1 \mu \beta \sqrt{R}}{M_L} \right) \\ &\leq \beta + \frac{1}{\mu} \frac{F_1 \mu \beta \sqrt{R}}{k-R} + \frac{1}{\mu} \ln \left(\frac{\min(F_1 \mu \beta \sqrt{R}, k-R)}{\min(\max(1, \frac{\sqrt{R}}{D_1 \sqrt{\beta}}), k-R)} \right) \\ &\leq \beta + \frac{1}{\mu} \frac{F_1 \mu \beta \sqrt{R}}{k-R} + \frac{1}{\mu} \ln(F_1 D_1 \beta^{3/2}) \\ &= \beta + \frac{1}{\mu} \frac{F_1 \mu \beta \sqrt{R}}{k-R} + \frac{1}{\mu} \frac{3}{2} \ln(\beta) + \frac{1}{\mu} \ln(F_1 D_1). \quad \square \end{aligned}$$

B.6.2. Proof of (29): Upper Bound on Time between Consecutive Downward Visits.

To bound the expectation $\mathbb{E} [\min(v_{i+1}^{(\text{down})}, X) - v_i^{(\text{down})} | \mathcal{F}_{v_i^{(\text{down})}}]$, we split the interval into two parts, $[\min(v_i^{(\text{up})}, X) - v_i^{(\text{down})}]$ and $[v_{i+1}^{(\text{down})} - v_i^{(\text{down})}]$.

To bound the expectation of the first quantity, it suffices to note that, if we couple the system to an $M/M/\infty$, the coupled number of jobs $\tilde{N}(t)$ will reach $R + M_L$ only after the original system. Using Claim 22 to bound this passage time, we thus know that

$$\begin{aligned} \mathbb{E} [\min(v_i^{(\text{up})}, X) - v_i^{(\text{down})} | \mathcal{F}_{v_i^{(\text{down})}}] &\leq \mathbb{E} [\min(T_{(R+M_L-1) \rightarrow (R+M_L)}^{M/M/\infty} + v_i^{(\text{down})}, X) - v_i^{(\text{down})} | \mathcal{F}_{v_i^{(\text{down})}}] \\ &\leq \mathbb{E} [T_{(R+M_L-1) \rightarrow (R+M_L)}^{M/M/\infty}] \\ &\leq \frac{D_2}{\sqrt{R}}. \end{aligned}$$

To bound the expectation of the second quantity, we provide two bounds. First, we again make use of the “wait-busy” idea; as we argued in the proof of (28),

$$\mathbb{E} [v_{i+1}^{(\text{down})} - v_i^{(\text{up})} | \mathcal{F}_{v_i^{(\text{up})}}] \leq \mathbb{E} [\min(v_{i+1}^{(\text{down})} - v_i^{(\text{up})}, \beta) | \mathcal{F}_{v_i^{(\text{up})}}] + \frac{1}{\mu M_L}.$$

From here, we note, by coupling to an M/M/1 with arrival rate and departure rate both equal to $k\lambda$, we can bound $\mathbb{E} \left[\min \left(v_{i+1}^{(\text{down})} - v_i^{(\text{up})}, \beta \right) \middle| v_i^{(\text{up})} < X \right]$ by the expected minimum between β and the length of a single-job busy period in that system. Applying Claim 21, we can complete the proof, finding that

$$\mathbb{E} \left[\min \left(v_{i+1}^{(\text{down})} - v_i^{(\text{up})}, \beta \right) \middle| \mathcal{F}_{v_i^{(\text{up})}} \right] \leq D_1 \frac{\sqrt{\beta}}{\sqrt{\mu R}} + \frac{6}{\mu R}.$$

For the second bound, we simply note that, during the draining phase, the number of busy servers $Z(t) \geq R + 1$. It follows from a simple coupling argument that

$$\mathbb{E} \left[\min \left(v_{i+1}^{(\text{down})} - v_i^{(\text{up})}, \beta \right) \middle| \mathcal{F}_{v_i^{(\text{up})}} \right] \leq \frac{1}{\mu}.$$

Combining the bounds pessimistically, we find that

$$\begin{aligned} \mathbb{E} \left[\min \left(v_{i+1}^{(\text{down})}, X \right) - v_i^{(\text{down})} \middle| \mathcal{F}_{v_i^{(\text{down})}} \right] &\leq \frac{D_2}{\mu\sqrt{R}} + \frac{D_3}{\mu R} + \frac{1}{\mu M_L} + \min \left(D_1 \frac{\sqrt{\beta}}{\sqrt{\mu R}}, \frac{1}{\mu} \right) \\ &\leq \frac{D_2}{\mu\sqrt{R}} + \frac{D_3}{\mu R} + \frac{2}{\mu M_L} \end{aligned}$$

B.6.3. Proof of (30): Upper Bound on Probability of Another Downward Visit. To bound the probability of an additional downcrossing, we again make a coupling argument. In particular, we couple again to the system which only has R servers busy, which gives an upper bound on the number of jobs in the system $N(t)$. If, in our coupled system, we reach $\tilde{N}(t) = R + M_L$ before we reach $\tilde{N}(t) = R$, then another upcrossing *must* have previously occurred in the original system, and thus another downcrossing must also occur. But, of course, we know classically that the probability that this happens is just $\frac{1}{M_L}$; this is precisely what is asserted by (30). \square

Appendix C: Proof of Multiplicative Tightness

We now show that the upper and lower bounds of Theorems 1 and 2, respectively, differ by at most a multiplicative factor.

THEOREM 3. *The bounds of Theorems 1 and 2 lie within a constant multiplicative factor of each other. In particular, using $=_c$ to denote equivalence modulo a multiplicative constant,*

$$\mathbb{E}[Q(\infty)] =_c \beta\sqrt{R} + \frac{1}{1-\rho}. \quad (31)$$

C.1. Proof for Lower Bound.

We prove Theorem 3 in two parts, showing equivalence for the lower bound, then for the upper bound. For the lower bound, we first discard all constants and a number of terms in the denominator, since $\beta > \frac{1}{\mu}$ by assumption. Doing so, we obtain

$$\mathbb{E}[Q(\infty)] \geq_c \frac{\mu\beta^2\sqrt{R} + \frac{[L_1\beta\sqrt{R} - k(1-\rho)]^+}{\mu k(1-\rho)} \left[[L_1\beta\sqrt{R} - k(1-\rho)]^+ + \frac{1}{1-\rho} \right]}{\beta + \frac{\beta\sqrt{R}}{\mu k(1-\rho)}}, \quad (32)$$

where \geq_c denotes that the inequality holds up to an (unspecified) constant factor.

C.1.1. Replacing the $[\cdot]^+$ term. Now, we show that the $[\beta\sqrt{R} - k(1 - \rho)]^+$ term can be replaced by the term $\mu\beta\sqrt{R}$, while losing only a constant factor; this turns out to be the difficult part. We approach this by casing on whether the positive term $\frac{1}{2}L_1\mu\beta\sqrt{R} \geq k(1 - \rho)$.

First case. If we have $\frac{1}{2}L_1\mu\beta\sqrt{R} \geq k(1 - \rho)$, then $[\mu L_1\beta\sqrt{R} - k(1 - \rho)]^+ \geq \frac{1}{2}L_1\mu\beta\sqrt{R} =_c \mu\beta\sqrt{R}$.

Second case. In the second case, assume that $\frac{1}{2}L_1\mu\beta\sqrt{R} < k(1 - \rho)$. In this case, even if we increase the value of the numerator by replacing the $[\cdot]^+$ term, the relevant term in the numerator becomes

$$\frac{\beta\sqrt{R}}{k(1 - \rho)} \left[\mu\beta\sqrt{R} + \frac{1}{1 - \rho} \right] = \mu\beta^2\sqrt{R} \cdot \frac{\sqrt{R}}{k(1 - \rho)} \cdot \left[1 + \frac{1}{\mu\beta\sqrt{R}} \right] =_c \mu\beta^2\sqrt{R} \cdot \frac{\sqrt{R}}{k(1 - \rho)} \leq_c \mu\beta^2\sqrt{R} \cdot \frac{1}{\beta}$$

where in the second equality we have used that $\sqrt{R} \geq 1$ and $\mu\beta \geq 1$, and in the final inequality we have used our case assumption. From here, it's clear that one can replace the term $[L_1\mu\beta\sqrt{R} - k(1 - \rho)]^+$ with the term $\mu\beta\sqrt{R}$ without altering the scaling behavior of numerator. In other words,

$$\mathbb{E}[Q(\infty)] \geq_c \frac{\mu\beta^2\sqrt{R} + \frac{\mu\beta\sqrt{R}}{\mu k(1 - \rho)} \left[\beta\sqrt{R} + \frac{1}{1 - \rho} \right]}{\beta + \frac{\beta\sqrt{R}}{\mu k(1 - \rho)}} =_c \mu\beta\sqrt{R} + \frac{\sqrt{R}}{k(1 - \rho) + \sqrt{R}} \frac{1}{1 - \rho}. \quad (33)$$

C.1.2. Bounding the final term. We now show equivalence for this final term, i.e. that

$$\mu\beta\sqrt{R} + \frac{\sqrt{R}}{k(1 - \rho) + \sqrt{R}} \frac{1}{1 - \rho} =_c \mu\beta\sqrt{R} + \frac{1}{1 - \rho}. \quad (34)$$

To do so, we bound the rightmost term in (33). Note that, since $\frac{\sqrt{R}}{k - R + \sqrt{R}} \leq \frac{R}{k - R + R} = \rho$, in order for this term to have an appreciable effect on the scaling, we must have that $\frac{\rho}{1 - \rho} \geq_c \mu\beta\sqrt{R}$, or, phrased more usefully, we must have $\sqrt{R} \geq_c \mu\beta k(1 - \rho)$. But even in this case, we can bound the factor in the rightmost term of (33) with $\frac{\sqrt{R}}{k(1 - \rho) + \sqrt{R}} \geq_c \frac{\beta k(1 - \rho)}{k(1 - \rho) + \mu\beta k(1 - \rho)} = \frac{\mu\beta}{1 + \mu\beta} =_c 1$; the multiplicative equivalence (34) follows.

C.2. Proof for Upper Bound.

C.2.1. Initial Steps. The proof for the upper bound follows along the same lines. First, note that the terms outside of the fraction are $\frac{R}{M} \leq \frac{R}{k(1 - \rho)} = \frac{\rho}{1 - \rho}$ and $\sqrt{\mu\beta R} \ll \mu\beta\sqrt{R}$. Discarding the lower order terms and constants, we obtain

$$\mathbb{E}[Q(\infty)] \leq_c \mu\beta\sqrt{R} + \frac{\rho}{1 - \rho} + \frac{\mu\beta^2\sqrt{R} + g\left(9(\mu\beta)^2R, 3\mu\beta\sqrt{R}, k(1 - \rho)\right)}{\beta + \frac{\mu\beta\sqrt{R}}{\mu k(1 - \rho)}}, \quad (35)$$

where one should recall that $M =_c \min(k(1 - \rho), \sqrt{\mu\beta R})$. For the terms in the fraction, the denominator of the upper bound is already up-to-constants-equivalent to the denominator of (32). It thus suffices to show that the numerator of the upper bound aligns with the numerator of (33). However, note that, by definition, the function

$$g\left(9(\mu\beta)^2R, 3(\mu\beta)\sqrt{R}, k(1 - \rho)\right) \triangleq \frac{3\mu\beta\sqrt{R}}{k(1 - \rho)} \left[\frac{3\mu\beta\sqrt{R} + 1}{2} + \frac{1}{1 - \rho} \right];$$

thus, the terms are clearly equivalent up to scaling. \square

Appendix D: The Approximation

In Section D, we present our approximation for the average waiting time in the M/M/k/Setup-Deterministic. We begin by discussing why we need such an approximation, then state the approximation, then give a short justification for its form.

D.1. Why we need an approximation

Despite our success in analyzing the M/M/k/Setup-Deterministic, our upper and lower bounds *alone* are not suitable for practical use in predicting the value of the average waiting time $\mathbb{E}[T_Q]$.

There are two reasons for this. First, although we can prove that our bounds are both within a constant factor of the true waiting time (Theorem 3), it's not *a priori* obvious whether the true value of $\mathbb{E}[T_Q]$ will get closer to one bound or the other as we vary the system parameters. Although the true value does not *seem* to ever get closer to a particular bound (and so we could conceivably just scale our lower bound to serve as a predictor), it would be better to have a more concrete theoretical justification for our prediction.

The second reason why our bounds are unsuitable for practical use is their complexity. Although both the upper and lower bounds are far more straightforward to compute than, for example, the average waiting time in the M/M/k/Setup-Exponential, both bounds incorporate a large number of terms and are thus somewhat difficult to reason about on the fly. As such, it would be better to have a predictor which incorporates only a few, easy-to-remember terms.

D.2. Approximation: Statement

To this end, we introduce the following approximation; the justification for the approximation follows. We evaluate the approximation in Figure 4; it is extremely accurate, across a variety of parameter settings.

Approximation 2 (Approximation to the average queue length.) *In the M/M/k/Setup-Deterministic, for offered loads $R \triangleq k\rho > 2$ and taking $C_{\text{apx}} \triangleq \sqrt{\frac{\pi}{2}}$,*

$$\mathbb{E}[Q(\infty)] \approx Q_{\text{apx}} \triangleq \frac{\frac{1}{2}\beta^2 C_{\text{apx}} \sqrt{R} + \frac{\beta C_{\text{apx}} \sqrt{R}}{\mu k(1-\rho)} \left[\frac{\beta C_{\text{apx}} \sqrt{R} + 1}{2} + \frac{1}{1-\rho} \right]}{\beta + \frac{\beta C_{\text{apx}} \sqrt{R}}{\mu k(1-\rho)}}. \quad (36)$$

D.3. Justification

We arrive at this bound via a straightforward combination of the bounds from Theorems 1 and 2, along with a few modifications. We follow our renewal-reward analysis, separately approximating the expected time integral over our renewal cycle and the expected length of that renewal cycle, the numerator and denominator of 36, respectively.

D.3.1. Justification of Numerator. We first approximate the numerator of our expression, the expected time integral over our chosen renewal cycle. We begin by recalling the lower bound on the time integral, Lemma 5, which states

$$\mathbb{E} \left[\int_0^X Q(t) dt \right] \geq L_1 \beta^2 \sqrt{R} + I^{\text{busy}} \left(\left[L_1 \beta \sqrt{R} - (k - R) \right]^+, k - R \right),$$

where

$$I^{\text{busy}}(x, z) \triangleq \frac{x}{\mu z} \left[\frac{x+1}{2} + \frac{1}{1 - \frac{k\lambda}{k\lambda + \mu z}} \right]$$

represents the time integral of the queue length a certain M/M/1 queue over a busy period started by x jobs.

To obtain the appropriate constant C_{apx} , we next note that, although our theorem states L_1 as an absolute constant, as the setup time β and the offered load R grow, the best possible constant will become $C_{\text{apx}} = \sqrt{\frac{\pi}{2}}$. Under the hood, this convergence stems from the fact that

$$\sum_{j=1}^R \prod_{i=1}^j \left(1 - \frac{j}{R} \right) \approx \int_0^\infty e^{-\frac{j^2}{2R}} dj = \frac{1}{2} \sqrt{2\pi R};$$

see the proof of Lemma 5 for more details.

To complete the approximation, it suffices to remove the subtraction of $(k - R)$ in the busy period term, which we anticipate is an artifact of our analysis. Removing it, we obtain the desired approximation

$$\mathbb{E} \left[\int_0^X Q(t) dt \right] \approx \frac{1}{2} \beta^2 C_{\text{apx}} \sqrt{R} + \frac{\beta C_{\text{apx}} \sqrt{R}}{\mu k (1 - \rho)} \left[\frac{\beta C_{\text{apx}} \sqrt{R} + 1}{2} + \frac{1}{1 - \rho} \right]. \quad (37)$$

D.3.2. Justification of Denominator. We next approximate the denominator of our expression, the expected length of our chosen renewal cycle. To do so, we again make use of the lower bound on the expected cycle length $\mathbb{E}[X]$ from Lemma 3, which states

$$\mathbb{E}[X] \geq \beta + \frac{L_1 \beta \sqrt{R}}{\mu k (1 - \rho)}.$$

By making the same convergence argument for L_1 , i.e. that $L_1 \rightarrow C_{\text{apx}}$ for large setup times β and large offered loads R , we obtain the denominator, completing both parts of our bound.

Appendix E: Coupling Claims.

E.1. Proof of Claim 2, the Coupling Integral Bound.

CLAIM 2 (Coupling Integral Bound). *Let τ be some stopping time and d_{gen} be the next down-crossing as described in Section 6.2.2. Suppose that, at time τ , we have a lower bound on the number of busy servers over a period, i.e. we know that the number of busy servers $Z(t) \geq R - j$, for all $t \in [\tau, \tau + \min(\ell, d_{\text{gen}})]$ and for some non-negative j . Then we have the following bound on the integral over this time period:*

$$\mathbb{E} \left[\int_{\tau}^{\tau + \min(d_{\text{gen}}, \ell)} [N(t) - h] dt \middle| \mathcal{F}_{\tau} \right] \leq \ell \cdot [N(\tau) - h]^+ + \frac{1}{2} \mu j \ell^2.$$

We prove this claim in three parts. First, we construct a coupled process $\tilde{N}(t) \geq N(t)$ on the interval of interest. Then, we give an upper bound on $\mathbb{E} \left[\int_{\tau}^{\tau + \min(\ell, d_{\text{gen}})} \tilde{N}(t) dt \middle| \mathcal{F}_{\tau} \right]$. Define $\tilde{N}(t)$ as

$$\tilde{N}(t) \triangleq N(\tau) + A(\tau, t) - \mathcal{D}[R - j](\tau, t).$$

Then, by Claim 1, we have that

$$\tilde{N}(t) \geq N(t).$$

on the interval of interest. To develop the integral, we first move the minimum from the bounds of integration into the integrand. In particular, we note that the quantity $N(d_{\text{gen}}) - h = 0$, and thus, for any $t > \tau + d_{\text{gen}}$, the quantity $N(\min(\tau + d_{\text{gen}}, t)) - h = 0$. On the other hand, for any $t < \tau + d_{\text{gen}}$, the quantity $N(\min(\tau + d_{\text{gen}}, t)) = N(t)$. It follows that

$$\begin{aligned} \int_{\tau}^{\tau + \min(\ell, d_{\text{gen}})} [N(t) - h] dt &= \int_{\tau}^{\tau + \min(\ell, d_{\text{gen}})} [N(\min(t, \tau + d_{\text{gen}})) - h] dt \\ &= \int_{\tau}^{\tau + \min(\ell, d_{\text{gen}})} [N(\min(t, \tau + d_{\text{gen}})) - h] dt + \int_{\tau + \min(\ell, d_{\text{gen}})}^{\tau + \ell} [N(\min(t, \tau + d_{\text{gen}})) - h] dt \\ &= \int_{\tau}^{\tau + \ell} [N(\min(t, \tau + d_{\text{gen}})) - h] dt \\ &\leq \int_{\tau}^{\tau + \ell} [\tilde{N}(\min(t, \tau + d_{\text{gen}})) - h] dt. \end{aligned}$$

Defining $\tilde{d}_{\text{gen}} \triangleq \min \{ t > 0 : \tilde{N}(\tau + t) \leq h \}$, since $\tilde{N}(t) \geq N(t)$, we know both that $\tilde{d}_{\text{gen}} \geq d_{\text{gen}}$ and that, for any $t \in [\tau + d_{\text{gen}}, \tau + \tilde{d}_{\text{gen}}]$,

$$\tilde{N}(t) - h \geq 0.$$

Moreover, the process $V(t)$ defined as

$$V(t) \triangleq \tilde{N}(t) - \mu j t$$

is a martingale. Thus, we have

$$\int_{\tau}^{\tau + \ell} [\tilde{N}(\min(t, \tau + d_{\text{gen}})) - h] dt \leq \int_{\tau}^{\tau + \ell} [\tilde{N}(\min(t, \tau + \tilde{d}_{\text{gen}})) - h] dt.$$

Taking the expectation, we find that

$$\begin{aligned} \mathbb{E} \left[\int_{\tau}^{\tau + \ell} [\tilde{N}(\min(t, \tau + \tilde{d}_{\text{gen}})) - h] dt \middle| \mathcal{F}_{\tau} \right] &= \int_{\tau}^{\tau + \ell} \mathbb{E} [\tilde{N}(\min(t, \tau + \tilde{d}_{\text{gen}})) - h \middle| \mathcal{F}_{\tau}] dt \\ &= \int_{\tau}^{\tau + \ell} \mathbb{E} [V(\min(t, \tau + \tilde{d}_{\text{gen}})) + \mu j (\min(\tau + \tilde{d}_{\text{gen}}, t)) - h \middle| \mathcal{F}_{\tau}] dt \\ &= \int_{\tau}^{\tau + \ell} \mathbb{E} [V(\tau) + \mu j (\min(\tau + \tilde{d}_{\text{gen}}, t)) - h \middle| \mathcal{F}_{\tau}] dt \end{aligned} \tag{38}$$

$$\begin{aligned}
&\leq \int_{\tau}^{\tau+\ell} \mathbb{E}[V(\tau) + \mu j t - h | \mathcal{F}_{\tau}] dt \\
&= \int_{\tau}^{\tau+\ell} \mathbb{E}[\tilde{N}(\tau) - \mu j \tau + \mu j t - h | \mathcal{F}_{\tau}] dt \\
&= [\tilde{N}(\tau) - h] \ell + \frac{1}{2} \mu j \ell^2,
\end{aligned}$$

where (38) is an application of Doob's Optimal Stopping Theorem.

E.2. Proof of Claim 3, the Coupling Probability Bound.

We now prove Claim 3, restated here for the reader's convenience.

CLAIM 3 (Coupling Probability Bound). *Let τ be some stopping time and d_{gen} be the next down-crossing as described in Section 6.2.2. We consider two cases.*

*In the first case, suppose that we have a **lower** bound on the number of busy servers $Z(t)$ over some length ℓ interval starting at time τ , i.e. the busy servers $Z(t) \geq R - j$, for all $t \in [\tau, \tau + \min(\ell, d_{\text{gen}})]$ and for some non-negative j . Then, we can bound the threshold-crossing probability by*

$$\Pr(d_{\text{gen}} < \ell | \mathcal{F}_{\tau}) \geq 2\Phi\left(-\left[\frac{N(\tau) - h + \mu j \ell}{\sqrt{\ell(2k\lambda - \mu j)}}\right]\right) - \frac{2}{3\sqrt{\ell(2k\lambda - \mu j)}}.$$

In particular, if $N(\tau) - h = c_1 \sqrt{\mu\beta R}$, then the probability $\Pr(d_{\text{gen}} < \ell | \mathcal{F}_{\tau}) \geq 2\Phi\left(-\frac{c_1}{\sqrt{2}}\right) - \frac{1}{100}$.

*In the second case, suppose that we instead have the **upper** bound on $Z(t) \leq R$ during this interval instead. Then,*

$$\Pr(d_{\text{gen}} < \ell | \mathcal{F}_{\tau}) \leq 2\Phi\left(-\left[\frac{N(\tau) - h}{\sqrt{2\ell k\lambda}}\right]\right) - \frac{2}{3\sqrt{2\ell k\lambda}}.$$

As before, if $N(\tau) - h = c_2 \sqrt{\mu\beta R}$, then the probability $\Pr(d_{\text{gen}} < \ell | \mathcal{F}_{\tau}) \leq 2\Phi\left(-\frac{c_2}{\sqrt{2}}\right) + \frac{1}{100}$.

We prove this result in three parts. First, we use Claim 1 to construct a process $\tilde{N}(t) \geq N(t)$ on the interval of interest. Afterwards, we analyze the down-crossing probability of this coupled process. In particular, we use a reflection argument to show that

$$\Pr(d_{\text{gen}} < \ell) \geq 2\Pr\left(\tilde{N}(\tau + \ell) \leq h\right),$$

then use a Berry-Esseen bound to bound this final probability. In what follows, we focus on the lower-bound; the upper bound follows in precisely the same way.

To construct our coupled process, we note that, by assumption, the number of busy servers $Z(t) \geq R - j$ for any $t \in [\tau, \tau + \min(\ell, d_{\text{gen}})]$. Thus, by Claim 1, the process $\tilde{N}(t)$ defined as

$$\tilde{N}(t) \triangleq N(\tau) + A(\tau, \tau + t) + \mathcal{D}[R - j](\tau, \tau + t)$$

is an upper bound for $N(t + \tau)$, i.e.

$$\tilde{N}(t) \geq N(\tau + t)$$

for any $t \in [0, \min(\ell, d_{\text{gen}})]$. By definition, we have that

$$\begin{aligned} \Pr(d_{\text{gen}} < \ell) &= \Pr\left(\inf_{t \in [0, \ell]} N(\tau + t) \leq h\right) \\ &\geq \Pr\left(\inf_{t \in [0, \ell]} \tilde{N}(t) \leq h\right). \end{aligned}$$

From a reflection argument, since \tilde{N} is upwards-biased,

$$\begin{aligned} \Pr\left(\inf_{t \in [0, \ell]} \tilde{N}(t) \leq h\right) &= \Pr\left(\inf_{t \in [0, \ell]} \tilde{N}(t) \leq h, \tilde{N}(\ell) < h\right) + \Pr\left(\inf_{t \in [0, \ell]} \tilde{N}(t) \leq h, \tilde{N}(\ell) \geq h\right) \\ &\geq 2 \Pr\left(\inf_{t \in [0, \ell]} \tilde{N}(t) \leq h, \tilde{N}(\ell) < h\right) \\ &= 2 \Pr\left(\tilde{N}(\ell) < h\right). \end{aligned}$$

Let $\sigma \triangleq \sqrt{\ell(2k\lambda - \mu j)}$. We now apply Now, assume that, for any x ,

$$\left| \Pr\left(\tilde{N}(\ell) < \tilde{N}(0) + \mu j \ell + x \sigma\right) - \Phi(x) \right| \leq \frac{0.3328}{\sigma}, \quad (39)$$

we have

$$\begin{aligned} \Pr\left(\tilde{N}(\ell) < h\right) &= \Pr\left(\tilde{N}(\ell) < \tilde{N}(0) + \mu j \ell + \frac{h - \mu j \ell - \tilde{N}(0)}{\sigma} \cdot \sigma\right) \\ &\geq \Phi\left(\frac{h - \mu j \ell - \tilde{N}(0)}{\sigma}\right) - \frac{1}{3\sigma} \\ &= \Phi\left(-\frac{[N(\tau) - h + \mu j \ell]}{\sigma}\right) - \frac{1}{3\sigma}. \end{aligned}$$

Putting this all together, we find

$$\Pr(d_{\text{gen}} < \ell | \mathcal{F}_\tau) \geq 2\Phi\left(-\frac{[N(\tau) - h + \mu j \ell]}{\sigma}\right) - \frac{2}{3\sigma}.$$

From here, then, it suffices to show (39). To begin, note that, if we choose some arbitrarily large n and define

$$X_i \triangleq \Pi'_i\left(\frac{k\lambda\ell}{n}\right) - \Pi''_i\left(\frac{\mu(R-j)\ell}{n}\right) - \frac{\mu j \ell}{n},$$

where each $\Pi(y)$ is an independent Poisson random variable with mean y , then

$$\tilde{N}(\ell) =_d \sum_{i=1}^n X_i + \mu j \ell + \tilde{N}(0).$$

To compute the moments of X_i , note that one can define centered Poisson random variables $A_i = \Pi\left(\frac{k\lambda\ell}{n}\right) - \frac{k\lambda\ell}{n}$ and $B_i = \Pi\left(\frac{\mu(R-j)\ell}{n}\right) - \frac{\mu(R-j)\ell}{n}$, and then take $X_i = A_i - B_i$. Doing this, one finds that

$$\mathbb{E}[X_i^2] = \mathbb{E}[(A_i - B_i)^2] = \frac{k\lambda\ell}{n} + \frac{\mu(R-j)\ell}{n} = \frac{\mu(2R-j)\ell}{n}$$

and, using the triangle inequality, that

$$\mathbb{E} \left[|X_i|^3 \right] = \mathbb{E} \left[|A_i - B_i|^3 \right] \leq \mathbb{E} \left[|A_i|^3 \right] + \mathbb{E} \left[|B_i|^3 \right] = \frac{\mu(2R-j)\ell}{n} + o\left(\frac{1}{n^2}\right).$$

We now apply the main result of Shevtsova (2011). Let $\sigma_n \triangleq \sqrt{\mathbb{E}[X_i^2]} = \sqrt{\frac{\mu(2R-j)\ell}{n}} = \frac{\sigma}{\sqrt{n}}$ and note that $\rho_n = \mathbb{E} \left[|X_i|^3 \right] < \sigma_n + o\left(\frac{1}{n^2}\right)$ (from Cook (2024)). Then, noting that $\rho_n \geq 1.286\sigma_n^3$ for sufficiently large n , we have

$$\max_x \left| \Pr \left(\frac{\sum X_i}{\sqrt{n}\sigma_n} < x \right) - \Phi(x) \right| \leq \frac{0.3328\rho_n + 0.429\sigma_n^3}{\sigma_n^3\sqrt{n}} = \frac{0.3328}{\sqrt{\mu(2R-j)\ell}} + o\left(\frac{1}{n}\right).$$

Now noting that

$$\frac{\sum_{i=1}^n X_i}{\sqrt{n}\sigma} = \frac{\tilde{N}(\ell) - \tilde{N}(0) - \mu j \ell}{\sigma}$$

and taking $n \rightarrow \infty$, we have our result. \square

E.3. Proof of Claim 18: Bound on Expected Value After Coupling.

CLAIM 18 (Bound on Expected Value after Coupling.) *Let τ be some stopping time and d_{gen} be the next down-crossing as described in Section 6.2.2. Suppose that we have a **lower** bound on the number of busy servers $Z(t)$ over some length ℓ interval starting at time τ , i.e. the busy servers $Z(t) \geq R - j$, for all $t \in [\tau, \tau + \min(\ell, d_{\text{gen}})]$ and for some non-negative j . Then, bounding the first moment,*

$$\mathbb{E} \left[[N(\tau + \ell) - h] \mathbf{1}_{d_{\text{gen}} > \ell} \middle| \mathcal{F}_\tau \right] \leq [N(\tau) - h] + \mu j \ell, \quad (40)$$

and, bounding the second moment,

$$\mathbb{E} \left[[N(\tau + \ell) - h] \mathbf{1}_{d_{\text{gen}} \geq \ell} \right] \leq [N(\tau) - h + \mu j \ell]^2 + 2\mu R \ell. \quad (41)$$

E.3.1. Proof. The proof is essentially an application of Doob's Optional Stopping Theorem to an appropriately selected martingale. To begin, we define a coupled process $\tilde{N}(t)$ with

$$\tilde{N}(t - \tau) \triangleq N(\tau) + A[\tau, t] - \mathcal{D}[R - j](\tau, t);$$

by Claim 1, we know that $\tilde{N}(t - \tau) \geq N(t)$ for any $t \in [\tau, \tau + \min(d_{\text{gen}}, \ell)]$, and that the coupled hitting time $\tilde{d}_{\text{gen}} \triangleq \min \left\{ t > 0 : \tilde{N}(t) \leq h \right\}$ can not be smaller than the original hitting time d_{gen} . It follows that

$$N(\tau + \ell) \mathbf{1}_{d_{\text{gen}} > \ell} \leq \tilde{N}(\ell) \mathbf{1}_{\tilde{d}_{\text{gen}} > \ell}.$$

Thus, we bound coupled versions of (40) and (41).

Construction of martingales. We now construct our martingales and set up the language of optional stopping. Note that, for any process $\tilde{N}(t)$ with independent, stationary increments, both functions V_1 and V_2 , defined as

$$V_1(t) \triangleq [\tilde{N}(t) - h] - \mathbb{E} [\tilde{N}(t) - \tilde{N}(0)]$$

and

$$\begin{aligned} V_2(t) &\triangleq [\tilde{N}(t) - h - \mathbb{E} [\tilde{N}(t) - \tilde{N}(0)]]^2 - \mathbb{E} \left[[\tilde{N}(t) - h - \mathbb{E} [\tilde{N}(t) - \tilde{N}(0)]]^2 \right] \\ &= \left(\tilde{N}(t) - h - \mu j t \right)^2 - \mu (2R - j) t \end{aligned}$$

are martingales Karatzas and Shreve (2012). Moreover, one has that

$$\begin{aligned} [\tilde{N}(\ell) - h] \mathbf{1}_{d_{\text{gen}}^{\tilde{}} > \ell} &= [\tilde{N}(\min(d_{\text{gen}}^{\tilde{}}, \ell)) - h] \mathbf{1}_{d_{\text{gen}}^{\tilde{}} > \ell} \\ &= [\tilde{N}(\min(d_{\text{gen}}^{\tilde{}}, \ell)) - h] \mathbf{1}_{d_{\text{gen}}^{\tilde{}} > \ell} + [\tilde{N}(\min(d_{\text{gen}}^{\tilde{}}, \ell)) - h] \mathbf{1}_{\ell \leq d_{\text{gen}}^{\tilde{}}} \\ &= [\tilde{N}(\min(d_{\text{gen}}^{\tilde{}}, \ell)) - h]. \end{aligned}$$

Proof of (40). Combining these facts allows us to prove our desired result. Applying Doob's Optional Stopping Theorem along with our previous deductions, we obtain

$$\begin{aligned} \mathbb{E} [[N(\tau + \ell) - h] \mathbf{1}_{d_{\text{gen}} > \ell} | \mathcal{F}_\tau] &\leq \mathbb{E} \left[[\tilde{N}(\ell) - h] \mathbf{1}_{d_{\text{gen}}^{\tilde{}} > \ell} \right] \\ &= \mathbb{E} \left[\tilde{N}(\min(d_{\text{gen}}^{\tilde{}}, \ell)) - h \right] \\ &= \mathbb{E} \left[V_1(\min(d_{\text{gen}}^{\tilde{}}, \ell)) \right] + \mu j \mathbb{E} \left[\min(d_{\text{gen}}^{\tilde{}}, \ell) \right] \\ &= \mathbb{E} [V_1(0)] + \mu j \mathbb{E} \left[\min(d_{\text{gen}}^{\tilde{}}, \ell) \right] \\ &= [\tilde{N}(0) - h] + \mu j \mathbb{E} \left[\min(d_{\text{gen}}^{\tilde{}}, \ell) \right] \\ &\leq [\tilde{N}(0) - h] + \mu j \ell \\ &= [N(\tau) - h] + \mu j \ell. \end{aligned}$$

Proof of (41). To do the same for the squared martingale $V_2(t)$, we must first note, via some algebra, that

$$\left(\tilde{N}(t) - h \right)^2 = V_2(t) + \left(\tilde{N}(t) - h \right) \mu j t - \mu j^2 t^2 + \mu (2R - j) t.$$

Now, applying the same deductions we made previously,

$$\begin{aligned} &\mathbb{E} [[N(\tau + \ell) - h] \mathbf{1}_{d_{\text{gen}} > \ell} | \mathcal{F}_\tau] \\ &\leq \mathbb{E} \left[[\tilde{N}(\ell) - h]^2 \mathbf{1}_{d_{\text{gen}}^{\tilde{}} > \ell} \right] \\ &= \mathbb{E} \left[\left(\tilde{N}(\min(d_{\text{gen}}^{\tilde{}}, \ell)) - h \right)^2 \right] \end{aligned}$$

$$\begin{aligned}
&= \mathbb{E} \left[V_2 \left(\min \left(\tilde{d}_{\text{gen}}, \ell \right) \right) \right] + \mathbb{E} \left[\left(\tilde{N} \left(\min \left(\tilde{d}_{\text{gen}}, \ell \right) \right) - h \right) \mu j \min \left(\tilde{d}_{\text{gen}}, \ell \right) \right] - \mu j^2 \left(\min \left(\tilde{d}_{\text{gen}}, \ell \right) \right)^2 \\
&\quad + \mu (2R - j) \mathbb{E} \left[\min \left(\tilde{d}_{\text{gen}}, \ell \right) \right] \\
&\leq \mathbb{E} \left[V_2 \left(\min \left(\tilde{d}_{\text{gen}}, \ell \right) \right) \right] + \mathbb{E} \left[\left(\tilde{N} \left(\min \left(\tilde{d}_{\text{gen}}, \ell \right) \right) - h \right) \right] \mu j \ell + \mu (2R) \ell \\
&\leq \mathbb{E} [V_2(0)] + \left[\tilde{N}(0) - h + \mu j \ell \right] \mu j \ell + \mu (2R) \ell \\
&= \left[\tilde{N}(0) - h \right]^2 + \left[\tilde{N}(0) - h \right] \mu j \ell + (\mu j \ell)^2 + \mu (2R) \ell \\
&= \left[\tilde{N}(0) - h + \mu j \ell \right]^2 - \left[\tilde{N}(0) - h \right] \mu j \ell + \mu 2R \ell \\
&\leq \left[\tilde{N}(0) - h + \mu j \ell \right]^2 + 2\mu R \ell \\
&= [N(\tau) - h + \mu j \ell]^2 + 2\mu R \ell. \quad \square
\end{aligned}$$

Appendix F: Hitting Time Bounds.

F.1. Proof of Claim 19, Discrete-Time Hitting Time Tail Bound.

CLAIM 19 (Discrete-Time Hitting Time Tail Bound). *Suppose one has an upwards-biased discrete random walk $V(t)$ where in each step*

$$\Pr(V(t+1) = V(t) + 1 | \mathcal{F}_t) = p = 1 - q,$$

where $p \geq \frac{1}{2} \geq q$. Suppose that $V(0) = 1$ and let the hitting time $\gamma \triangleq \min \{t \in \mathcal{N} : V(t) = 0\}$ be the first timestep where the walk $V(t) = 0$. Then, for $n \geq 1$,

$$\Pr(\gamma \geq 2m + 1) \leq \frac{1}{\sqrt{\pi}} \frac{2q}{\sqrt{m}} \left(1 + \frac{1}{2(m+1)} \right).$$

Moreover, if $p = q = \frac{1}{2}$, then

$$\Pr(\gamma \geq 2m + 1) \geq \frac{1}{\sqrt{\pi}} e^{-\frac{1}{6m}} \frac{1}{\sqrt{m+1}}.$$

F.1.1. Proof We first note, as in Williams et al. (2022), that by a counting argument $\Pr(\gamma = 2\ell + 1) = q(qp)^\ell C_\ell$, where $C_\ell \triangleq \frac{1}{\ell+1} \frac{(2\ell)!}{\ell!\ell!}$ is the ℓ -th Catalan number; note that γ can not be even, since the number of downward steps must exceed the number of upward steps by exactly 1.

We proceed by bounding the Catalan numbers using Stirling's approximation. For $m = 0$, then $\Pr(\gamma \geq 1) = \Pr(\gamma \geq 2) = p$, i.e. the probability that the first step is an upward step. For $m \geq 1$, applying Stirling's approximation and simplifying gives

$$e^{-\frac{1}{6\ell}} \frac{1}{\sqrt{\pi\ell(\ell+1)}} q(4pq)^\ell \leq \Pr(\gamma = 2\ell + 1) \leq \frac{1}{\sqrt{\pi\ell(\ell+1)}} q(4pq)^\ell.$$

Lower bound. Since we are interested in the lower bound only when $q = p = \frac{1}{2}$, we obtain that

$$\begin{aligned} \Pr(\gamma \geq 2m + 1) &\geq \frac{1}{\sqrt{pi}} \frac{1}{2} \sum_{\ell=m}^{\infty} \frac{e^{-\frac{1}{6\ell}}}{\sqrt{\ell(\ell+1)}} \\ &\geq \frac{1}{\sqrt{\pi}} e^{-\frac{1}{6m}} \frac{1}{2} \sum_{\ell=m}^{\infty} \frac{1}{\sqrt{\ell(\ell+1)}} \\ &\geq \frac{1}{\sqrt{\pi}} e^{-\frac{1}{6m}} \frac{1}{2} \int_m^{\infty} \frac{1}{(\ell+1)^{3/2}} d\ell \\ &= \frac{1}{\sqrt{\pi}} e^{-\frac{1}{6m}} \frac{1}{\sqrt{m+1}}. \end{aligned}$$

Upper bound. Noting that $4pq \leq 1$, we have likewise that

$$\begin{aligned} \Pr(\gamma \geq 2m + 1) &\leq \frac{1}{\sqrt{\pi}} q \sum_{\ell=m}^{\infty} \frac{1}{\sqrt{\ell(\ell+1)}} \\ &\leq \frac{1}{\sqrt{\pi}} q \frac{1}{\sqrt{m(m+1)}} + \int_m^{\infty} \frac{1}{\ell^{3/2}} d\ell \\ &= \frac{1}{\sqrt{\pi}} q \frac{2}{\sqrt{m}} \left(1 + \frac{1}{2(m+1)} \right). \end{aligned}$$

F.2. Proof of Claim 20, Continuous-Time Hitting Time Tail Bound.

We further extend this discrete-time bound into a continuous-time bound.

CLAIM 20 (Continuous-Time Hitting Time Tail Bound). *Suppose one has an Poisson arrival process $Y_A(t)$ of rate $k\lambda$ and a Poisson departure process $Y_D(t)$ of rate $\mu(R - j)$, for some integer $j \geq 0$. Let the continuous random walk $X_c(t) = Y_A(t) - Y_D(t)$, with $X_c(0) = 1$, and define $\gamma_c \triangleq \min\{t > 0 : X_c(t) = 0\}$. Let $\nu = (2R - j)\mu t$. For any $\nu \geq 3$, we have*

$$\Pr(\gamma_c \geq t) \leq \frac{b_1}{\sqrt{2}} \left(\frac{1}{\sqrt{\nu}} + \frac{b_2}{\nu^{3/2}} \right)$$

where $b_1 = \sqrt{\frac{2}{\pi}}$ and $b_2 = 1 + \frac{2.5}{b_1\sqrt{2}}$.

Moreover, if $j = 0$, then

$$\Pr(\gamma_c \geq t) \geq \frac{b_1}{\sqrt{2}} e^{-\frac{1}{3(\nu-1)}} \frac{1}{\sqrt{\nu+2}}.$$

F.2.1. Proof of Upper Bound. To prove this claim, we first condition on the value of $Y_T = Y_A(t) + Y_D(t)$, the total number of Poisson events during the interval $[0, t]$, then relate that to the same question in a discrete-time random walk, a la Claim 19. Note that $Y_T \sim \text{Poisson}(\nu)$, and thus

$$\begin{aligned} \Pr(\gamma_c \geq t) &= \Pr(\gamma \geq Y_T) \\ &= \sum_{j=0}^{\infty} e^{-\nu} \frac{\nu^j}{j!} \Pr(\gamma \geq j) \\ &= e^{-\nu} + 2p\nu e^{-\nu} + \sum_{j=3}^{\infty} e^{-\nu} \frac{\nu^j}{j!} \Pr(\gamma \geq j + \mathbf{1}_{j \text{ is even}}) \\ &= e^{-\nu} + 2p\nu e^{-\nu} + \sum_{j=0}^{\infty} e^{-\nu} \frac{\nu^j}{j!} \Pr\left(\gamma \geq 2 \left(\frac{j + \mathbf{1}_{j \text{ is even}} - 1}{2} \right) + 1\right). \end{aligned}$$

Applying the discrete upper bound to the sum, we obtain

$$\begin{aligned}
& \sum_{j=3}^{\infty} e^{-\nu} \frac{\nu^j}{j!} \Pr \left(\gamma \geq 2 \left(\frac{j + \mathbf{1}_{j \text{ is even}} - 1}{2} \right) + 1 \right) \\
& \leq b_1 \sqrt{2q} \sum_{j=3}^{\infty} e^{-\nu} \frac{\nu^j}{j!} \frac{1}{\sqrt{j + \mathbf{1}_{j \text{ is even}} - 1}} \left(1 + \frac{1}{(j + \mathbf{1}_{j \text{ is even}} + 1)} \right) \\
& = b_1 \sqrt{2q} \frac{1}{\nu} \sum_{j=3}^{\infty} e^{-\nu} \frac{\nu^{(j+1)}}{(j+1)!} \frac{1}{\sqrt{j + \mathbf{1}_{j \text{ is even}} - 1}} \left(j + 1 + \frac{j + 1}{j + \mathbf{1}_{j \text{ is even}} + 1} \right) \\
& \leq b_1 \sqrt{2q} \frac{1}{\nu} \sum_{j=3}^{\infty} e^{-\nu} \frac{\nu^{(j+1)}}{(j+1)!} \frac{j + 2}{\sqrt{j + \mathbf{1}_{j \text{ is even}} - 1}} \\
& \leq b_1 \sqrt{2q} \frac{1}{\nu} \sum_{j=3}^{\infty} e^{-\nu} \frac{\nu^{(j+1)}}{(j+1)!} \frac{j + \mathbf{1}_{j \text{ is even}} - 1 + 3}{\sqrt{j + \mathbf{1}_{j \text{ is even}} - 1}} \\
& = b_1 \sqrt{2q} \frac{1}{\nu} \sum_{j=3}^{\infty} e^{-\nu} \frac{\nu^{(j+1)}}{(j+1)!} \left(\sqrt{j + \mathbf{1}_{j \text{ is even}} - 1} + \frac{3}{\sqrt{j + \mathbf{1}_{j \text{ is even}} - 1}} \right).
\end{aligned}$$

From here, we note that the function $f(x) = \sqrt{x} + \frac{3}{\sqrt{x}}$ is both increasing and concave for all $x \geq 3$. After increasing the argument and applying Jensen's inequality, we find that

$$\begin{aligned}
& \leq b_1 \sqrt{2q} \frac{1}{\nu} \sum_{j=3}^{\infty} e^{-\nu} \frac{\nu^{(j+1)}}{(j+1)!} \left(\sqrt{j+1} + \frac{3}{\sqrt{j+1}} \right) \\
& \leq b_1 \sqrt{2q} \frac{1}{\nu} \left(\sqrt{\nu} + \frac{3}{\sqrt{\nu}} \right),
\end{aligned}$$

where in the final line we have used that the function $f(x)$ is increasing in x for any $x \geq 3$, and that $\mathbb{E}[Y_T \mathbf{1}_{Y_T \geq 4}] \geq \nu - 3 \geq 3$. Thus, we have that

$$\begin{aligned}
\Pr(\gamma_c \geq t) & \leq (3\nu) e^{-\nu} + 2q \sqrt{\frac{2}{\pi}} \left(\frac{1}{\sqrt{\nu}} + \frac{1}{\nu^{3/2}} \right) \\
& \leq \frac{2.5}{\nu^{3/2}} + 2q \sqrt{\frac{2}{\pi}} \left(\frac{1}{\sqrt{\nu}} + \frac{1}{\nu^{3/2}} \right)
\end{aligned}$$

F.2.2. Proof of Lower Bound. We approach the initial stages of the proof in the precisely the same way, obtaining

$$\begin{aligned}
\Pr(\gamma_c \geq t) & = \Pr(\gamma \geq Y_T) \\
& = e^{-\nu} + 2p\nu e^{-\nu} + \sum_{j=3}^{\infty} e^{-\nu} \frac{\nu^j}{j!} \Pr \left(\gamma \geq 2 \left(\frac{j + \mathbf{1}_{j \text{ is even}} - 1}{2} \right) + 1 \right) \\
& \geq \sum_{j=3}^{\infty} e^{-\nu} \frac{\nu^j}{j!} b_1 e^{-\frac{1}{3(j+\mathbf{1}_{j \text{ is even}}-1)}} \frac{q\sqrt{2}}{\sqrt{(j + \mathbf{1}_{j \text{ is even}} + 1)}} \\
& \geq \sum_{j=3}^{\infty} e^{-\nu} \frac{\nu^j}{j!} b_1 e^{-\frac{1}{3(j-1)}} \frac{q\sqrt{2}}{\sqrt{(j+2)}}.
\end{aligned}$$

Applying Jensen's inequality, we obtain

$$\geq b_1 q \sqrt{2} e^{-\frac{1}{3(\nu-1)}} \frac{1}{\sqrt{\nu+2}}.$$

□

F.3. Proof of Claim 21, Bound on Expected Length of Stopped Random Walk.

CLAIM 21 (Bound on Expected Length of Stopped Random Walk). *Suppose we have a critically loaded M/M/1 queue with arrival rate and departure rate both equal to $k\lambda$, with offered load $R > 100$ and setup time $\beta > 100$. Suppose also that at time 0, a job arrives. Let τ be the length of the busy period which follows. Then, letting $b_1 \triangleq \sqrt{\frac{2}{\pi}}$,*

$$\mathbb{E}[\min(\beta, \tau)] \leq b_1 \frac{\sqrt{\beta}}{\sqrt{\mu R}} + \frac{6}{\mu R}.$$

F.3.1. Proof. From Claim 20, the continuous-time random walk hitting time bound, we have that

$$\Pr(\tau \geq t) \leq \frac{b_1}{\sqrt{2}} \left(\frac{1}{\sqrt{\nu}} + \frac{b_2}{\nu^{3/2}} \right), \quad (42)$$

where $\nu = 2\mu R t$ and we require that $\nu \geq 3$. By integrating this bound (using a bound of 1 wherever this bound doesn't apply), we obtain

$$\begin{aligned} \mathbb{E}[\min(\beta, \tau)] &= \int_0^\beta \Pr(\tau > t) dt \leq \frac{3}{2\mu R} + \int_{\frac{3}{2\mu R}}^\beta \frac{b_1}{\sqrt{2}} \left(\frac{1}{\sqrt{2\mu R t}} + \frac{b_2}{(2\mu R)^{3/2}} \right) dt \\ &\leq \frac{3}{2\mu R} + b_1 \frac{\sqrt{\beta}}{\sqrt{\mu R}} + \frac{b_1 b_2}{4\mu R} \left[2\sqrt{\frac{2}{3}} \right] \leq b_1 \frac{\sqrt{\beta}}{\sqrt{\mu R}} + \frac{6}{\mu R}. \quad \square \end{aligned}$$

F.4. Proof of Claim 22, Bound on the Expected Hitting Time in the M/M/ ∞ .

CLAIM 22 (M/M/ ∞ Passage Time Bound). *Given an M/M/ ∞ queue, let $T_{x \rightarrow y}$ denote the random amount of time taken to go from state x to state y . Suppose this system has an arrival rate of $k\lambda$ and a per-server departure rate of μ . Let $R \triangleq k\frac{\lambda}{\mu}$. Then, for any h such that $1 \leq h \leq \sqrt{R}$,*

$$\mathbb{E}[T_{(R+h-1) \rightarrow (R+h)}] \leq \frac{\sqrt{2\pi}}{\mu\sqrt{R}} \left(1 + \frac{h}{R} \right)^{h-\frac{1}{2}} e^{\frac{1}{12R}} \leq D_2 \frac{\sqrt{\pi}}{\mu\sqrt{R}}.$$

F.4.1. Proof. The proof here is quite simple. First, we note that the passage time in the M/M/ ∞ from state $(R+h-1)$ to state $(R+h)$ is exactly the passage time from those states in the M/M/ $(R+h)/(R+h)$. This new system has a nice product form, so that

$$\begin{aligned} \mathbb{E}[T_{(R+h-1) \rightarrow (R+h)}] &\leq \mathbb{E}[T_{(R+h) \rightarrow (R+h)}] \\ &= \frac{1}{\mu(R+h)} \frac{1}{\pi_{R+h}} \\ &= \frac{1}{\mu(R+h)} \frac{\sum_{i=0}^{R+h} \frac{R^i}{i!}}{\frac{R^{R+h}}{(R+h)!}} \end{aligned}$$

$$\begin{aligned}
&\leq \frac{1}{\mu(R+h)} e^R \frac{(R+h)!}{R^{R+h}} \\
&\leq \frac{1}{\mu(R+h)} e^R \frac{e^{\frac{1}{12(R+h)}} \sqrt{2\pi(R+h)} (R+h)^{R+h} e^{-(R+h)}}{R^{R+h}} \\
&\leq e^{\frac{1}{12R}} \frac{1}{\mu} \frac{1}{\mu\sqrt{R+h}} \sqrt{2\pi} \left(1 + \frac{h}{R}\right)^{R+h} e^{-h} \\
&\leq \frac{\sqrt{2\pi}}{\mu\sqrt{R}} \left(1 + \frac{h}{R}\right)^{h-\frac{1}{2}} e^{\frac{1}{12R}} \\
&\leq \frac{1}{\mu} \frac{\sqrt{2\pi}}{\sqrt{R}} e^{\frac{h^2}{R}} e^{\frac{1}{12R}} \\
&\leq \frac{7}{\mu\sqrt{R}},
\end{aligned}$$

where we have made extensive use of Stirling's approximation and the bound $(1+x) \leq e^x$.

Appendix G: Miscellaneous Claims

G.1. Proof of Claim 23, the Busy Period Integral Bound.

CLAIM 23 (Busy Period Integral Bound). *Suppose that, at time τ , we can guarantee that $N(\tau) \geq Z(\tau) \geq R+j$. Let $\eta_i \triangleq \min\{t > 0 : N(t) \leq R+i\}$, for $i \in \{j, j+1, \dots, [N(\tau) - R]\}$. Then,*

$$\mathbb{E} \left[\int_{\tau}^{\eta_j} [N(t) - R] dt \middle| \mathcal{F}_{\tau} \right] \leq (N(\tau) - (R+j)) \left[\frac{3}{2\mu j} + \frac{1}{\mu} + \frac{R}{\mu j^2} \right] + \frac{(N(\tau) - (R+j))^2}{2\mu j} \triangleq I^{\text{busy}}([N(\tau) - R], j).$$

We prove this claim via an appeal to conventional M/M/1 busy period analysis. In particular, we first note that

$$\int_{\tau}^{\eta_j} [N(t) - R] dt = \sum_{i=j+1}^{N(\tau)-R} \int_{\eta_i}^{\eta_{i-1}} [N(t) - R] dt,$$

meaning we need only bound the integrals between the η_i 's. To bound that process, we define a coupled process $\tilde{N}(t)$ and bound the integrals over that process.

To do so, note that, until time η_j , the number of busy servers $Z(t) \geq R+j$. By Claim 1, we can define, for each index i , the i -th coupled process $\tilde{N}_i(t)$ as

$$\tilde{N}_i(t) = N(\eta_{i+1}) + A(\eta_{i+1}, t) - \mathcal{D}[R+j](\eta_{i+1}, t),$$

and have $\tilde{N}(t) \geq N(t)$ on the interval $[\eta_{i+1}, \eta_i]$. Furthermore, we can extend our integral of interest from the interval $[\eta_{i+1}, \eta_i]$ to the interval $[\eta_{i+1}, \tilde{\eta}_i]$, where $\tilde{\eta}_i \triangleq \min\{t > 0 : N(t) \leq R+i\}$. Now, we note that

$$\mathbb{E} \left[\int_{\eta_{i+1}}^{\tilde{\eta}_i} [\tilde{N}_i(t) - R] dt \middle| \mathcal{F}_{\tau} \right] = \mathbb{E} \left[\int_{\eta_{i+1}}^{\tilde{\eta}_i} [\tilde{N}_i(t) - (R+i)] dt \middle| \mathcal{F}_{\tau} \right] + i \mathbb{E}[\eta_{i+1} - \tilde{\eta}_i | \mathcal{F}_{\tau}].$$

The first term on the right is simply the expected time integral of the number of jobs in an M/M/1 queue over a busy period, with arrival rate $k\lambda$ and departure rate $\mu(R+j)$. The second term is simply the quantity

i multiplied by the expected length of that M/M/1 busy period. Let $\rho_j = \frac{k\lambda}{\mu(R+j)}$. Then, from standard results on the M/M/1 busy period,

$$\mathbb{E} \left[\int_{\eta_{i+1}}^{\tilde{\eta}_i} [\tilde{N}_i(t) - (R+i)] dt \middle| \mathcal{F}_\tau \right] = \frac{1}{\mu j} \left[\frac{1}{1-\rho_j} \right] = \frac{1}{\mu j} \left[\frac{R}{j} + 1 \right] = \frac{1}{\mu j} + \frac{R}{\mu j^2}.$$

Summing over all values of i , we obtain

$$\begin{aligned} & \mathbb{E} \left[\int_{\tau}^{\eta_j} [N(t) - R] dt \middle| \mathcal{F}_\tau \right] \\ & \leq \sum_{i=j+1}^{N(\tau)-R} \mathbb{E} \left[\int_{\eta_i}^{\tilde{\eta}_{i-1}} [\tilde{N}_i(t) - R] dt \middle| \mathcal{F}_\tau \right] \\ & = \sum_{i=j+1}^{N(\tau)-R} \left[\frac{1}{\mu j} + \frac{R}{\mu j^2} \right] + i \frac{1}{\mu j} \\ & = (N(\tau) - (R+j)) \left[\frac{1}{\mu j} + \frac{R}{\mu j^2} \right] + (N(\tau) - (R+j)) \frac{1}{\mu} + \frac{1}{\mu j} \left[\frac{(N(\tau) - (R+j))(N(\tau) - (R+j) + 1)}{2} \right] \\ & = (N(\tau) - (R+j)) \left[\frac{3}{2\mu j} + \frac{1}{\mu} + \frac{R}{\mu j^2} \right] + \frac{(N(\tau) - (R+j))^2}{2\mu j}, \end{aligned}$$

as desired. \square

G.2. Proof of Claim 13, the Wait Busy Claim.

CLAIM 13 (Wait Busy Claim). *Let τ be some stopping time, let the number of jobs $N(\tau) = R + h$, and define $\text{ns}(h) \triangleq \min\{h, k(1-\rho)\}$. Let the down-crossing $d_{\text{gen}} \triangleq \min\{t > 0 : N(\tau+t) = R + h - 1\}$. If $Z(\tau) \geq R$, then*

$$\mathbb{E} \left[\int_{\tau}^{\tau+d_{\text{gen}}} [N(t) - (R+h-1)] dt \middle| \mathcal{F}_\tau \right] \leq Y_{R+\text{ns}(h)}(\tau) + g(1 + 2\mu R \mathbb{E}[\min(Y_{R+\text{ns}(h)}(\tau), d_{\text{gen}})], 1, \mu \text{ns}(h)), \quad (12)$$

where the function $g(x, y, z) \triangleq x \frac{1}{2\mu z} + y \left[\frac{R}{\mu z^2} + \frac{3}{2\mu z} \right]$.

Furthermore,

$$\begin{aligned} \mathbb{E} \left[\int_{T_A}^{v_1^{(\text{down})}} [N(t) - R] dt \middle| \mathcal{F}_{T_A} \right] & \leq \left[\beta + \frac{1}{\mu} \right] \left[\mathbb{E}[N(T_A) - R] + \frac{R}{M_B} \right] + \frac{2}{\mu} \ln \left(\frac{\mathbb{E}[N(T_A) - R]}{M_B} \right) \\ & \quad + g \left(\mathbb{E} \left[[N(T_A) - R]^2 \right] + 2\mathbb{E}[N(T_A) - R], \mathbb{E}[N(T_A) - R], k(1-\rho) \right). \end{aligned} \quad (13)$$

G.2.1. Proof of (12). We prove the two parts of Claim 13 separately; we first show (12) by applying coupling, martingales, and busy period analysis. First, note that, if the down-crossing at $\tau + d_{\text{gen}}$ does not occur by time $Y_{R+\text{ns}(h)}(\tau)$, then the system must have at least $(R + \text{ns}(h))$ servers at its disposal afterwards. (Note that the $\text{ns}(\cdot)$ function here is just to account for the case where you have more jobs than servers.) Accordingly, we split our analysis into two parts.

First portion. For the first portion, since the number of busy servers $Z(t) \geq R$, by coupling our system to a critically-loaded M/M/1 using Claim 2, we have

$$\mathbb{E} \left[\int_{\tau}^{\tau + \min(d_{\text{gen}}, Y_{R+\text{ns}(h)}(\tau))} [N(t) - (R + h - 1)] dt \middle| \mathcal{F}_{\tau} \right] \leq Y_{R+\text{ns}(h)}(\tau).$$

Second portion. For the second portion, since, at that point the number of busy servers $Z(t) \geq R + \text{ns}(h)$, we can apply a stronger bound. In particular, at time $(\tau + Y_{R+\text{ns}(h)}(\tau))$, we can couple to an accordingly-stronger M/M/1 with the same number of jobs. From basic busy period analysis and Claim 1, this tells us that, letting the adjusted number of jobs $N_{\text{adj}}(t) \triangleq N(\tau + t) - (R + h - 1)$ and the important remaining setup time $Y_{\text{imp}} \triangleq Y_{R+\text{ns}(h)}(\tau)$ as a shorthand,

$$\mathbb{E} \left[\mathbf{1}_{Y_{\text{imp}} < d_{\text{gen}}} \int_{Y_{\text{imp}}}^{d_{\text{gen}}} [N_{\text{adj}}(t)] dt \middle| \mathcal{F}_{\tau + Y_{\text{imp}}} \right] \leq \mathbf{1}_{Y_{\text{imp}} < d_{\text{gen}}} g \left(N_{\text{adj}}(Y_{\text{imp}})^2, N_{\text{adj}}(Y_{\text{imp}}), \text{ns}(h) \right),$$

since the function $g(x^2, x, z)$ describes the integral of the number of jobs over a busy period started by x jobs in an M/M/1 with arrival rate $k\lambda$ and departure rate $\mu(R + z)$. Note that $\mathbf{1}_{Y_{\text{imp}} < d_{\text{gen}}} [N(\tau + Y_{\text{imp}}) - (R + h - 1)] = [N(\tau + \min(d_{\text{gen}}, Y_{\text{imp}})) - (R + h - 1)]$, since $N(\tau + d_{\text{gen}}) = R + h - 1$ by definition. Coupling our system to the critically-loaded M/M/1 \tilde{N} system as before, then taking expectations and applying Doob's Optional Stopping Theorem, we obtain, as desired,

$$\mathbb{E} \left[\mathbf{1}_{Y_{R+\text{ns}(h)}(\tau) < d_{\text{gen}}} \int_{Y_{R+\text{ns}(h)}(\tau)}^{d_{\text{gen}}} N_{\text{adj}}(t) dt \middle| \mathcal{F}_{\tau + Y_{R+\text{ns}(h)}(\tau)} \right] \leq g \left(1 + 2\mu R \mathbb{E} [\min(Y_{R+\text{ns}(h)}(\tau), d_{\text{gen}})], 1, \text{ns}(h) \right).$$

Combining these two terms, we obtain (12).

G.2.2. Proof of (13). We now apply (12) to prove (13).

Decomposition in terms of η_i 's. We first fix the filtration/state at the accumulation time T_A , then define the hitting times $\eta_i \triangleq \min \{t > T_A : N(t) \leq R + i\}$ for a set number of jobs $i \in [M_B, N(T_A)]$; we accordingly omit the filtration at time T_A in our expectations. Note that the down-crossing $v_1^{(\text{down})} = \eta_{M_B}$, by this definition. From there, it's clear that

$$\int_{T_A}^{v_1^{(\text{down})}} [N(t) - R] dt = \sum_{i=M_B}^{N(T_A)-R-1} \int_{\eta_{i+1}}^{\eta_i} [N(t) - R] dt.$$

For each of these terms, we can separate $[N(t) - R] = [N(t) - (R + i)] + i$ and apply (12) to find

$$\mathbb{E} \left[\int_{\eta_{i+1}}^{\eta_i} [N(t) - R] dt \right] \leq \mathbb{E} [Y_{R+\text{ns}(h)}(\eta_{i+1})] + i \mathbb{E} [\min(Y_{R+\text{ns}(i+1)}(\eta_{i+1}), \eta_i - \eta_{i+1})] \quad (43)$$

$$+ g \left([1 + 2\mu R \mathbb{E} [\min(Y_{R+\text{ns}(i+1)}(\eta_{i+1}), (\eta_i - \eta_{i+1}))]], 1, \mu \text{ns}(i + 1) \right) \quad (44)$$

$$+ i \cdot \frac{1}{\mu \text{ns}(i + 1)}. \quad (45)$$

We analyze each of these terms separately.

Bound on (43), the remaining setup time portion. From here, it suffices to note that the sum

$$\sum_{i=1}^{N(T_A)-R} i \min(Y_{R+ns(i)}(\tau), \eta_i - \eta_{i+1}) + \sum_{i=1}^{N(T_A)-R} \mathbb{E}[Y_{R+ns(i+1)}(\eta_{i+1})] \leq \beta [N(T_A) - R]; \quad (46)$$

actually, the statement is true without expectations. To see this, we make an interchange of summation argument. First note that, if the $(R+i)$ -th server becomes busy, then, by the monotonicity of server states, all servers of index smaller than $(R+i)$ must also be busy; in other words, the remaining setup time $Y_{R+ns(i+1)}(\eta_{i+1}) = 0$ for all $i < s$. To use this, we let s be the largest index for which $Y_{R+ns(s)}(\eta_s) < (\eta_{s-1} - \eta_s)$. Note also that η_i is the first time after T_A that the number of jobs $N(t) \leq R+i$ (and so must have been continuously decreasing from time T_A); it follows that the remaining setup time $Y_{R+ns(i)}(\eta_i) \leq [\beta - (\eta_i - T_A)]^+$. Breaking things down further,

$$\sum_{i=s}^{N(T_A)-R} i \min(Y_{R+ns(i+1)}(\eta_{i+1}), (\eta_i - \eta_{i+1})) = \sum_{i=s}^{N(T_A)-R} i(\eta_i - \eta_{i+1}) + sY_{R+ns(s)}(\eta_s).$$

From here, by an interchange of summation argument,

$$\sum_{i=s}^{N(T_A)-R} i(\eta_i - \eta_{i+1}) = \sum_{i=s}^{N(T_A)-R} \sum_{j=1}^i (\eta_i - \eta_{i+1}) = \sum_{j=s}^{N(T_A)-R} \sum_{i=j}^{N(T_A)-R} (\eta_i - \eta_{i+1}) = \sum_{j=1}^{N(T_A)-R} \eta_s - T_A,$$

where, by definition, the (relative to T_A) hitting time $(\eta_s - T_A) \leq \beta - Y_{R+ns(s)}(\eta_s)$; using this, (46) follows.

Bound on (44), the busy period integral portion. Applying similar reason to the sum of the first terms in g , and using the independence of g in its first and second arguments (i.e. that $g(x, y, z) = f_1(x, z) + f_2(y, z)$ for two functions f_1 and f_2 linear in their first argument),

$$\sum_{i=M_B}^{N(T_A)-R} (44) \leq g(2\mu R\beta, 0, M_B) + g\left([N(T_A) - k]^+, [N(T_A) - k]^+, k(1-\rho)\right) + \sum_{i=M_B}^{\min(N(T_A)-R-1, k(1-\rho))} g(1, 1, i+1).$$

This last term above can be bounded by replacing it with an integral, which gives

$$\begin{aligned} \sum_{i=M_B}^{\min(N(T_A)-R, k(1-\rho))} g(1, 1, i) &\leq \int_{M_B}^{\min(N(T_A)-R, k(1-\rho))} \left(\frac{2}{\mu i} + \frac{R}{\mu i^2} \right) di \\ &\leq \frac{2}{\mu} \ln \left(\frac{\min(N(T_A) - R, k(1-\rho))}{M_B} \right) + \frac{R}{\mu} \left[\frac{1}{M_B} \right] \\ &\leq \frac{2}{\mu} \ln \left(\frac{N(T_A) - R}{M_B} \right) + \frac{R}{\mu M_B}. \end{aligned}$$

Bound on (45), the busy period length portion. Using the definition of the function g , we also find that $\sum_{i=M_B}^{N(T_A)-R} \frac{i}{\mu ns(i)} \leq g\left(\left([N(T_A) - k]^+\right)^2 + [N(T_A) - k]^+, 0, k(1-\rho)\right) + \frac{1}{\mu} \min(k(1-\rho), N(T_A) - R, \cdot)$

Combining the terms to bound (13), the integral from T_A to $v_1^{(\text{down})}$. Combining terms, noting that $N(T_A) - k \geq N(T_A) - R \geq 0$, and applying Jensen's inequality to the $\ln(\cdot)$ term, we find that

$$\begin{aligned}
 (13) &\leq \beta \mathbb{E}[N(T_A) - R] + g(2\mu R\beta, 0, M_B) + \mathbb{E}\left[g\left([N(T_A) - k]^+, [N(T_A) - k]^+, k(1 - \rho)\right)\right] \\
 &\quad + \mathbb{E}\left[\frac{2}{\mu} \ln\left(\frac{N(T_A) - R}{M_B}\right)\right] + \frac{1}{\mu} \frac{R}{M_B} \\
 &\quad + \mathbb{E}\left[g\left(\left([N(T_A) - k]^+\right)^2 + [N(T_A) - k]^+, 0, k(1 - \rho)\right)\right] + \frac{1}{\mu} \mathbb{E}[\min(k(1 - \rho), N(T_A) - R)] \\
 &\leq \left[\beta + \frac{1}{\mu}\right] \left[\mathbb{E}[N(T_A) - R] + \frac{R}{M_B}\right] + \frac{2}{\mu} \ln\left(\frac{\mathbb{E}[N(T_A) - R]}{M_B}\right) \\
 &\quad + g\left(\mathbb{E}[N(T_A) - R]^2 + 2\mathbb{E}[N(T_A) - R], \mathbb{E}[N(T_A) - R], k(1 - \rho)\right). \quad \square
 \end{aligned}$$

G.3. Proof of Claim 6

We now prove Claim 6, restated here.

CLAIM 6 (Bound on the Probability of an Up-crossing $p_{\text{rise}}^{(j)}$). Let $p_{\text{rise}}^{(j)}$ be the probability that the total number of jobs $N(t)$ exceeds $R + C_3\sqrt{\mu\beta R}$ during epoch j defined in (6). Then, for any epoch $j \geq A_5\sqrt{R}$, we have $p_{\text{rise}}^{(j)} \geq 0.99\frac{A_5}{\sqrt{R}}$.

We show a more general claim: that, for $j \geq A_5\sqrt{R}$,

$$p_{\text{rise}}^{(j)} \geq 0.99\frac{j}{R}. \quad (47)$$

G.3.1. Proof of (47): Lower Bound on $p_{\text{rise}}^{(j)}$. We begin with a simple probability manipulation:

$$\begin{aligned}
 p_{\text{rise}}^{(j)} &\triangleq \Pr\left(N(t) \geq C_3\sqrt{\mu\beta R} \text{ at some point during epoch } j \mid \mathcal{F}_{\tau_j}\right) \\
 &\geq \Pr\left(N(t) \geq C_3\sqrt{\mu\beta R} \text{ during the interval } [\tau_j, \min(\tau_j + \beta, \tau_{j+1})] \mid \mathcal{F}_{\tau_j}\right).
 \end{aligned}$$

From here, we make with a useful observation: since there are no server *in setup* at the beginning of an epoch (as we have just turned off a server), no servers can complete setup in the first β time of an epoch. Thus, the number of busy servers $Z(t) \leq R - j$ during this time, and, by Claim 1, the coupled process

$$\tilde{N}(t) \triangleq N(\tau_j) + A(\tau_j, t) - \mathcal{D}[R - j](\tau_j, t)$$

must be a lower bound on $N(t)$, during the interval $[\tau_j, \tau_j + \beta]$. Moreover, the number of busy servers $Z(t)$ can not be smaller than $R - j$ until the beginning of epoch $j + 1$ either. Thus, we find that the behavior of $N(t)$ corresponds *exactly* with the behavior of $\tilde{N}(t)$ during the interval $[\tau_j, \min(\tau_{j+1}, \tau_j + \beta)]$.

We now use this coupled process to analyze our original probability. Define the up-crossing time τ_{up} as

$$\tau_{\text{up}} \triangleq \min\left\{t > 0 : \tilde{N}(\tau_j + t) \geq R + C_3\sqrt{\mu\beta R}\right\}.$$

Likewise, define the down-crossing time τ_{down} as

$$\tau_{\text{down}} \triangleq \min \left\{ t > 0 : \tilde{N}(\tau_j + t) \leq R - (j + 1) \right\}.$$

It follows that

$$\begin{aligned} & \Pr \left(\text{reach } N(t) \leq R - (j + 1) \text{ during the interval } [\tau_j, \min(\tau_j + \beta, \tau_{j+1})] \middle| \mathcal{F}_{\tau_j} \right) \\ &= \Pr \left(\text{reach } \tilde{N}(t - \tau_j) \leq R - (j + 1) \text{ during the interval } [\tau_j, \min(\tau_j + \beta, \tau_{j+1})] \middle| \mathcal{F}_{\tau_j} \right) \\ &= \Pr(\tau_{\text{up}} \leq \beta, \tau_{\text{up}} < \tau_{\text{down}}) \\ &= \Pr(\tau_{\text{up}} \leq \beta) - \Pr(\tau_{\text{up}} \leq \beta, \tau_{\text{up}} \geq \tau_{\text{down}}) \\ &= \Pr(\tau_{\text{up}} \leq \beta) - \Pr(\tau_{\text{up}} \leq \beta | \tau_{\text{up}} \geq \tau_{\text{down}}) \Pr(\tau_{\text{up}} \geq \tau_{\text{down}}). \end{aligned}$$

We now observe that

$$\Pr(\tau_{\text{up}} \leq \beta | \tau_{\text{up}} \geq \tau_{\text{down}}) \leq \Pr(\tau_{\text{up}} \leq \beta), \quad (48)$$

since the process has farther to go, less time to do so, and the process's behavior is translation-invariant (this last point is why we needed to analyze the coupled process instead).

Continuing from where we left off, we find that

$$\begin{aligned} p_{\text{rise}}^{(j)} &= \Pr(\tau_{\text{up}} \leq \beta) - \Pr(\tau_{\text{up}} \leq \beta | \tau_{\text{up}} \geq \tau_{\text{down}}) \Pr(\tau_{\text{up}} \geq \tau_{\text{down}}) \\ &\geq \Pr(\tau_{\text{up}} \leq \beta) - \Pr(\tau_{\text{up}} \leq \beta) \Pr(\tau_{\text{up}} \geq \tau_{\text{down}}) \\ &= \Pr(\tau_{\text{down}} > \tau_{\text{up}}) \Pr(\tau_{\text{up}} \leq \beta) \\ &\geq \Pr(\tau_{\text{down}} > \infty) \Pr(\tau_{\text{up}} \leq \beta) \\ &= \frac{j}{R} \Pr(\tau_{\text{up}} \leq \beta), \end{aligned}$$

where the last equality is a classical result on upwards-biased discrete random walks (one can think of \tilde{N} as a discrete random walk driven by a Poisson process of rate $(k\lambda + \mu(R - j))$, where the probability that \tilde{N} increases at a Poisson event is $\frac{k\lambda}{k\lambda + \mu(R - j)} = \frac{R}{2R - j}$).

From here, it suffices to lower bound $\Pr(\tau_{\text{up}} \leq \beta)$. To begin, note

$$\begin{aligned} \Pr(\tau_{\text{up}} \leq \beta) &= \Pr \left(\sup_{t \in [0, \beta]} \tilde{N}(t) \geq R + C_3 \sqrt{\mu\beta R} \right) \\ &\geq \Pr \left(\tilde{N}(\beta) \geq R + C_3 \sqrt{\mu\beta R} \right) \\ &= \Pr \left(A(\tau_j, \tau_j + \beta) - \mathcal{D}[R - j](\tau_j, \tau_j + \beta) \geq j + C_3 \sqrt{\mu\beta R} \right). \end{aligned}$$

Noting that the number of arrivals $A(\tau_j, \tau_j + \beta)$ and the number of departures $\mathcal{D}[R - j]$ ($[\tau_j, \tau_j + \beta]$) are independent Poisson r.v.'s, we can apply the Berry-Esseen bound of Claim 24 to find

$$\begin{aligned} &= 1 - \Phi\left(\frac{\mu\beta j - j - C_3\sqrt{\mu\beta R}}{\sqrt{\mu\beta(2R - j)}}\right) - \frac{1}{3\sqrt{\mu\beta(2R - j)}} \\ &\geq 1 - \Phi\left(\frac{0.99\mu\beta j - C_3\sqrt{\mu\beta R}}{\sqrt{2\mu\beta R}}\right) - \frac{1}{3\sqrt{\mu\beta R}} \\ &= 1 - \Phi\left(-0.99\frac{j}{\sqrt{R}}\sqrt{\mu\beta} + \frac{C_3}{\sqrt{2}}\right) - \frac{1}{3\sqrt{\mu\beta R}} \\ &\geq 1 - \Phi\left(-9.9A_5 + \frac{C_3}{\sqrt{2}}\right) - \frac{1}{300}. \end{aligned}$$

To complete the proof, we set the constant A_5 such that the final probability is ≥ 0.99 . In particular, we need

$$\Phi\left(-9.9A_5 + \frac{C_3}{\sqrt{2}}\right) \leq \frac{2}{300},$$

which is achieved when $A_5 > \frac{C_3}{9.9\sqrt{2}} + 0.25$; choosing $A_5 = 1$ gives the result. \square

G.4. Proof of Claim 24.

CLAIM 24 (Berry-Esseen bound for the Skellam distribution). *Given two independent random variables $Y_1 \sim \text{Poisson}(\mu_1)$ and $Y_2 \sim \text{Poisson}(\mu_2)$, as well as a constant C with $\mu_1 > \mu_2 + C$, one has*

$$\Pr(Y_1 - Y_2 \geq C) \geq 1 - \Phi\left(-\left\lceil \frac{\mu_1 - \mu_2 - C}{\mu_1 + \mu_2} \right\rceil\right) - \frac{1}{3\sqrt{\mu_1 + \mu_2}}.$$

G.5. Proof.

This follows directly from the Poisson Berry-Esseen bound of Cook (2024), applied twice; first approximating Y_1 then approximating Y_2 . \square

G.6. Proof of (18): Lower Bound on $\mathbb{E}[L]$, Expected Value of First Long Epoch Index.

We prove this result by first showing that

$$\Pr(L > j | L \geq j) \geq \left(1 - \frac{j}{R}\right) \left(1 - \frac{b_1}{\sqrt{\mu\beta R}}\right), \quad (49)$$

where $b_1 = \frac{2}{\sqrt{\pi}}$. Next, we show that this implies that for any $\delta \in (0, 1)$ and any $j < \delta R$,

$$\Pr(L > j) \geq \left(1 - \frac{b_1}{\sqrt{\mu\beta R}}\right)^{j+1} e^{-\frac{j(j+1)}{2R} \frac{1}{1-\delta}}. \quad (50)$$

From here, we use the sum of tails formula $\mathbb{E}[L] = \sum_{j=0}^{\infty} \Pr(L > j)$ to show

$$\mathbb{E}[L] \geq \left(1 - \frac{b_1}{\sqrt{\mu\beta R}}\right) \left(\left[\sqrt{\frac{\pi}{2}}(1-\delta) - \frac{1.15(1-\delta)}{\sqrt{\mu\beta}}\right] \sqrt{R} - \frac{1}{2} - \frac{2(1-\delta)}{\delta} e^{-R \frac{\delta^2}{1-\delta}}\right).$$

Choosing $\delta = \frac{2}{\sqrt{R}}$ then noting that $\mu\beta \geq 100$ and $R \geq 100$ gives the result.

G.6.1. Proof of (49): Lower Bound on Probability that Current Epoch is Short. Recall that an epoch j is *long* if $\tau_{j+1} - \tau_j > \beta$, that L is the index of the first long epoch, and that, if $L \geq j$, then we learn that $L \geq j$ precisely at time τ_j , i.e. when epoch j begins. Moreover, since the system is Markovian, the behavior of the system from τ_j onwards is completely independent of what happened previously. Thus,

$$\Pr(L > j | L \geq j) = \Pr(L > j | \mathcal{F}_{\tau_j}, L \geq j) = \Pr(\tau_{j+1} - \tau_j \leq \beta | \mathcal{F}_{\tau_j}, L \geq j) = \Pr(\tau_{j+1} - \tau_j \leq \beta).$$

From here, we note that the random time $\tau_{j+1} - \tau_j$ is a stopping time; a hitting time, to be exact. Moreover, since the number of servers $Z(t)$ can not increase before time $\tau_j + \beta$ and can not decrease until τ_{j+1} , we have that the coupled process $\tilde{N}(t)$ defined as

$$\tilde{N}(t - \tau_j) \triangleq 1 + A(\tau_j, t) - \mathcal{D}[R - j]((\tau_j, t))$$

is in correspondence with $N(t)$; in particular,

$$N(t) = \tilde{N}(t - \tau_j) + R - j - 1$$

for any time $t \in [\tau_j, \min(\tau_j + \beta, \tau_{j+1})]$. If we define the coupled hitting time $\gamma_c \triangleq \min\{t > 0 : \tilde{N}(t) \leq 0\}$, then we also have that the hitting time $\gamma_c = \tau_{j+1} - \tau_j$, whenever the event $\{\tau_{j+1} - \tau_j \leq \beta\}$ occurs. From here, we can apply Claim 20 to find that

$$\Pr(\gamma_c \leq \beta) \geq \left(1 - \frac{j}{R}\right) \left(1 - \frac{b_1}{\sqrt{\mu\beta R}}\right), \text{ as desired.}$$

G.6.2. Proof of (50). Having shown the above bound on the conditional extension of the tail, we note that, for $j \leq \delta R$,

$$\begin{aligned} \Pr(L \geq j + 1) &= \Pr(L \geq j + 1 | L \geq j) \Pr(L \geq j | L \geq j - 1) \cdots \Pr(L \geq 1) \\ &\geq \prod_{i=0}^j \left(1 - \frac{i}{R}\right) \left(1 - \frac{b_1}{\sqrt{\mu\beta R}}\right) \\ &\geq \left(1 - \frac{b_1}{\sqrt{\mu\beta R}}\right)^{j+1} \prod_{i=0}^j e^{-\frac{i}{R-i}} \\ &\geq \left(1 - \frac{b_1}{\sqrt{\mu\beta R}}\right)^{j+1} e^{-\sum_{i=0}^j \frac{i}{R-i}} \\ &= \left(1 - \frac{b_1}{\sqrt{\mu\beta R}}\right)^{j+1} e^{-\sum_{i=0}^j \frac{i}{R} \frac{R}{R-j}} \\ &= \left(1 - \frac{b_1}{\sqrt{\mu\beta R}}\right)^{j+1} e^{-\frac{j(j+1)}{2R} \frac{1}{1-\delta}}, \end{aligned}$$

as desired. □

G.6.3. Proof of (G.6): Final Bound on $\mathbb{E}[L]$ using Gaussian Integral. We now complete the proof. Let $a \triangleq \frac{1}{2R} \frac{1}{1-\delta}$ and $\psi \triangleq -\ln\left(1 - \frac{b_1}{\sqrt{\mu\beta R}}\right)$ as a shorthand. Then we can rewrite (50) as

$$\Pr(L \geq j+1) \geq e^{-aj^2 - (\psi+a)j - \psi}.$$

Now, using the sum-of-tails formula for expectations, we find that

$$\begin{aligned} \mathbb{E}[L] &= \sum_{j=0}^{R-1} \Pr(L \geq j+1) \\ &\geq \sum_{j=0}^{\delta R-1} \Pr(L \geq j+1) \\ &\geq \sum_{j=0}^{\delta R-1} e^{-aj^2 - (\psi+a)j - \psi} \\ &\geq \int_0^{\delta R} e^{-aj^2 - (\psi+a)j - \psi} \mathbf{d}j \\ &= \int_0^{\delta R} e^{-a(j^2 + (\frac{\psi}{a}+1)j) - \psi} \mathbf{d}j \\ &= \int_0^{\delta R} e^{-a(j + \frac{1}{2}(\frac{\psi}{a}+1))^2 + \frac{a}{4}(\frac{\psi}{a}+1)^2 - \psi} \mathbf{d}j \\ &= e^{\frac{a}{4}(\frac{\psi}{a}+1)^2 - \psi} \int_0^{\delta R} e^{-a(j + \frac{1}{2}(\frac{\psi}{a}+1))^2} \mathbf{d}j. \end{aligned}$$

Evaluating the integral further, we find that

$$\begin{aligned} \int_0^{\delta R} e^{-a(j + \frac{1}{2}(\frac{\psi}{a}+1))^2} \mathbf{d}j &= \int_{\frac{1}{2}(\frac{\psi}{a}+1)}^{\delta R + \frac{1}{2}(\frac{\psi}{a}+1)} e^{-aj^2} \mathbf{d}j \\ &= \int_0^{\infty} e^{-aj^2} \mathbf{d}j - \int_0^{\frac{1}{2}(\frac{\psi}{a}+1)} e^{-aj^2} \mathbf{d}j - \int_{\delta R + \frac{1}{2}(\frac{\psi}{a}+1)}^{\infty} e^{-aj^2} \mathbf{d}j. \end{aligned}$$

We now bound each of these integrals in turn. First, we know classically that

$$\int_0^{\infty} e^{-aj^2} \mathbf{d}j = \frac{1}{2} \sqrt{\frac{\pi}{a}} = \sqrt{\frac{\pi}{2}} \cdot \sqrt{1-\delta} \sqrt{R} \geq \sqrt{\frac{\pi}{2}} \cdot (1-\delta) \sqrt{R}$$

Next, we note that, since the integrand is ≤ 1 ,

$$\begin{aligned} \int_0^{\frac{1}{2}(\frac{\psi}{a}+1)} e^{-aj^2} \mathbf{d}j &\leq \frac{1}{2} \left(\frac{\psi}{a} + 1 \right) \\ &= \frac{1}{2} \left(2R(1-\delta) \ln\left(\frac{1}{1 - \frac{b_1}{\sqrt{\mu\beta R}}} \right) \right) + \frac{1}{2} \\ &\leq R(1-\delta) \frac{b_1}{\sqrt{\mu\beta R}} \frac{1}{1 - \frac{b_1}{\sqrt{\mu\beta R}}} + \frac{1}{2} \\ &\leq \left(\frac{1}{\sqrt{\beta}} \right) (1-\delta) \cdot \frac{100 \cdot b_1}{100 - b_1} \sqrt{R} + \frac{1}{2} \\ &\leq \frac{1.15(1-\delta)}{\sqrt{\mu\beta}} \sqrt{R} + \frac{1}{2}. \end{aligned}$$

Finally, we have that,

$$\int_{\delta R + \frac{1}{2}(\frac{\psi}{\alpha} + 1)}^{\infty} e^{-aj^2} dj \leq \int_{\delta R}^{\infty} e^{-aj^2} dj \leq \int_{\delta R}^{\infty} e^{-a\delta Rj} dj = \frac{1}{a\delta R} e^{-a\delta^2 R^2} = \frac{2(1-\delta)}{\delta} e^{-R\frac{\delta^2}{1-\delta}}.$$

To complete the proof, we note that $e^{-\psi} = \left(1 - \frac{b_1}{\sqrt{\mu\beta R}}\right)$, thus

$$\begin{aligned} \mathbb{E}[L] &\geq e^{\frac{a}{4}(\frac{\psi}{\alpha} + 1)^2 - \psi} \int_0^{\delta R} e^{-a(j + \frac{1}{2}(\frac{\psi}{\alpha} + 1))^2} dj \\ &\geq e^{-\psi} \int_0^{\delta R} e^{-a(j + \frac{1}{2}(\frac{\psi}{\alpha} + 1))^2} dj \\ &\geq \left(1 - \frac{b_1}{\sqrt{\mu\beta R}}\right) \left[\sqrt{\frac{\pi}{2}}(1-\delta) - \frac{1.15(1-\delta)}{\sqrt{\mu\beta}} \right] \sqrt{R} - \frac{1}{2} - \frac{2(1-\delta)}{\delta} e^{-R\frac{\delta^2}{1-\delta}} \\ &= \left(1 - \frac{b_1}{\sqrt{\mu\beta R}}\right) \left[\sqrt{\frac{\pi}{2}}(1-\delta) - \frac{1.15(1-\delta)}{\sqrt{\mu\beta}} \right] \sqrt{R} - \frac{1}{2} - \frac{2(1-\delta)}{\delta} e^{-R\frac{\delta^2}{1-\delta}} \\ &= \left(1 - \frac{b_1}{\sqrt{\mu\beta R}}\right) \left(\left[\sqrt{\frac{\pi}{2}}(1-\delta) - \frac{1.15(1-\delta)}{\sqrt{\mu\beta}} \right] \sqrt{R} - \frac{1}{2} - \frac{2(1-\delta)}{\delta} e^{-R\frac{\delta^2}{1-\delta}} \right). \end{aligned}$$

From here, we could choose δ to maximize our lower bound further based on system parameters, but a simple choice is $\delta = \frac{2}{\sqrt{R}}$. This gives

$$\begin{aligned} \mathbb{E}[L] &\geq \left(1 - \frac{b_1}{\sqrt{\mu\beta R}}\right) \left[\left(1 - \frac{2}{\sqrt{R}}\right) \left(\sqrt{\frac{\pi}{2}} - \frac{1.15}{\sqrt{\mu\beta}} - 2e^{-4} \right) - \frac{1}{2\sqrt{R}} \right] \sqrt{R} \\ &\geq \frac{2}{3} \sqrt{\frac{\pi}{2}} \sqrt{R}, \end{aligned}$$

as desired.

G.7. Proof of Claim 14.

CLAIM 14 (Upper Bound on $\mathbb{E}[N(T_A)]$). Recall that $T_A \triangleq \min\{t > 0 : Z(t) = R + 1\}$. Then,

$$\mathbb{E}[N(T_A) - R] \leq F_1 \mu \beta \sqrt{R} \left(1 + \frac{F_2}{\sqrt{\mu\beta}}\right) \leq 2.9 \mu \beta \sqrt{R}$$

and

$$\mathbb{E}\left[(N(T_A) - R)^2\right] \leq F_1^2 (\mu\beta)^2 R \left(1 + \frac{F_2}{\sqrt{\mu\beta}}\right)^2 + 2\mu\beta R \leq 8.4(\mu\beta)^2 R + 2\mu\beta R$$

where $F_1 = 2.12$ and $F_2 = 3.645$.

G.7.1. Proof. The beginning of the proof will be the same for both of these inequalities. Using the up-crossing and down-crossing decomposition of Section 6.3, we know that time T_A occurs either during a *rise* or during a *fall*. Since the number of jobs $N(t) \leq R + C_3 \sqrt{\mu\beta R}$ during a rise,

$$[N(T_A) - R] \mathbf{1}_{T_A \text{ during a rise}} \leq C_3 \sqrt{\mu\beta R} \mathbf{1}_{T_A \text{ during a rise}}.$$

If T_A occurs during a fall, we need a more nuanced bound. Writing out the event $\{T_A \text{ during a fall}\}$ in terms of disjoint events, we find

$$\{T_A \text{ during a fall}\} = \bigcup_{j=0}^{R-1} \bigcup_{i=1}^{\infty} \left\{ u_i^{(j)} \leq T_A < d_i^{(j)} \right\},$$

so that, for $c \in \{1, 2\}$,

$$\begin{aligned} \mathbb{E} \left[[N(T_A) - R]^c \mathbf{1}_{T_A \text{ during a fall}} \right] &= \sum_{j=0}^{R-1} \sum_{i=1}^{\infty} \mathbb{E} \left[[N(T_A) - R]^c \mathbf{1}_{u_i^{(j)} \leq T_A < d_i^{(j)}} \right] \\ &= \sum_{j=0}^{R-1} \sum_{i=1}^{\infty} \mathbb{E} \left[[N(T_A) - R]^c \mathbf{1}_{T_A < d_i^{(j)}} \middle| \mathcal{F}_{u_i^{(j)}}, n_u^{(j)} \geq i \right] \Pr(n_u^j \geq i). \end{aligned}$$

To bound this conditional expectation, we apply Claim 18. Notice that $N(u_i^{(j)}) - R = C_3 \sqrt{\mu\beta R}$, the $(R+1)$ -th server starts up at time $T_A = u_i^{(j)} + Y_{R+1}(u_i^{(j)})$ if $T_A < d_i^{(j)}$, the time $d_i^{(j)}$ is a hitting time, and that $Z(t) \geq R - j$ until time $\tau_{j+1} \geq d_i^{(j)}$. It follows that

$$\mathbb{E} \left[[N(T_A) - R] \mathbf{1}_{T_A < d_i^{(j)}} \middle| \mathcal{F}_{u_i^{(j)}}, n_u^{(j)} \geq i \right] \leq C_3 \sqrt{\mu\beta R} + \mu j Y_{R+1}(u_i^{(j)}) \leq C_3 \sqrt{\mu\beta R} + \mu j \beta,$$

and that

$$\begin{aligned} \mathbb{E} \left[[N(T_A) - R]^2 \mathbf{1}_{T_A < d_i^{(j)}} \middle| \mathcal{F}_{u_i^{(j)}}, n_u^{(j)} \geq i \right] &\leq \left(C_3 \sqrt{\mu\beta R} + \mu j Y_{R+1}(u_i^{(j)}) \right)^2 + \mu 2 R Y_{R+1}(u_i^{(j)}) \\ &\leq \left(C_3 \sqrt{\mu\beta R} + \mu j \beta \right)^2 + \mu 2 R \beta \\ &= C_3^2 \mu \beta R + 2 C_3 \sqrt{\mu\beta R} \mu j \beta + (\mu \beta)^2 j^2 + 2 \mu R \beta. \end{aligned}$$

It now suffices to bound $\sum_j \sum_i j^c \Pr(n_u^{(j)} \geq i)$, where $c \in \{0, 1, 2\}$. We do this via the same method used in Section 6.3:

$$\begin{aligned} \sum_{j=1}^R \sum_{i=1}^{\infty} j^c \Pr(n_u^{(j)} \geq i) &\leq \sum_{j=1}^R \sum_{i=1}^{\infty} j^c \Pr(n_e \geq j) p_{\text{rise}}^{(j)} (1 - p_2)^{i-1} \\ &= \frac{1}{p_2} \sum_{j=1}^R j^c p_{\text{rise}}^{(j)} \Pr(n_e \geq j) \\ &\leq \frac{1}{C_4 p_2} \sum_{j=1}^R j^c C_4 p_{\text{rise}}^{(j)} \prod_{\ell=0}^{j-1} \left(1 - C_4 p_{\text{rise}}^{(\ell)} \right). \end{aligned}$$

This is simply the expectation of a time-varying geometric random variable G , with $\Pr(G = j | G \geq j) = C_4 p_{\text{rise}}^{(j)}$. It follows that if one lower-bounds $p_{\text{rise}}^{(j)}$, then an upper bound on the desired expectation is obtained. Applying Claim 6, we note that we are essentially bounding G using $Y \sim \text{Geometric} \left(\frac{0.99 C_4 A_5}{\sqrt{R}} \right)$ and saying $Y + A_5 \sqrt{R}$ stochastically-dominates G . It follows that

$$\mathbb{E}[G] \leq A_5 \sqrt{R} + \frac{1}{0.99 C_4 A_5} \sqrt{R}$$

and that, for any b ,

$$\begin{aligned}\mathbb{E}[(G+b)^2] &\leq \mathbb{E}[(Y + A_5\sqrt{R} + b)^2] = \mathbb{E}[Y^2] + 2(A_5\sqrt{R} + b)\mathbb{E}[Y] + (A_5\sqrt{R} + b)^2 \\ &= 2\mathbb{E}[Y]^2 - \mathbb{E}[Y] + 2(A_5\sqrt{R} + b)\mathbb{E}[Y] + (A_5\sqrt{R} + b)^2 \\ &\leq (\mathbb{E}[Y] + A_5\sqrt{R} + b)^2 + \mathbb{E}[Y]^2.\end{aligned}$$

Defining $B_5 \triangleq \frac{C_3}{C_4 p_2}$, $B_6 \triangleq \frac{1}{C_4 p_2} \left(\frac{1}{0.99 C_4 A_5} + A_5 \right)$, and $B_7 \triangleq \frac{1}{2 C_4 p_2} \left[\frac{1}{(0.99 C_4 A_5)^2} + 2 \right]$, it follows that

$$\begin{aligned}\mathbb{E} \left[[N(T_A) - R]^2 \mathbf{1}_{T_A \text{ during a fall}} \right] &\leq \frac{1}{C_4 p_2} \sum_{j=1}^R C_4 p_{\text{rise}}^{(j)} \left[\left(C_3 \sqrt{\mu \beta R} + \mu j \beta \right)^2 + \mu 2 R \beta \right] \prod_{\ell=0}^{j-1} \left(1 - C_4 p_{\text{rise}}^{(\ell)} \right) \\ &= \frac{1}{C_4 p_2} \mathbb{E} \left[\left(C_3 \sqrt{\mu \beta R} + \mu G \beta \right)^2 + \mu 2 R \beta \right] \\ &\leq \frac{1}{C_4 p_2} \mathbb{E} \left[\left(C_3 \sqrt{\mu \beta R} + \mu \beta Y + \mu \beta A_5 \sqrt{R} \right)^2 + \mu 2 R \beta \right] \\ &= \frac{1}{C_4 p_2} \left[\left(C_3 \sqrt{\mu \beta R} + \mu \beta \frac{1}{0.99 C_4 A_5} \sqrt{R} + \mu \beta A_5 \sqrt{R} \right)^2 + \mu \beta \frac{1}{(0.99 C_4 A_5)^2} R + 2 \mu \beta R \right] \\ &\leq \frac{1}{C_4^2 p_2^2} \left[\left(C_3 \sqrt{\mu \beta R} + \mu \beta \frac{1}{0.99 C_4 A_5} \sqrt{R} + \mu \beta A_5 \sqrt{R} \right)^2 \right] + \frac{1}{C_4 p_2} \left[\frac{1}{(0.99 C_4 A_5)^2} + 2 \right] \mu \beta R \\ &= \left(B_5 \sqrt{\mu \beta R} + B_6 \mu \beta \sqrt{R} \right)^2 + 2 B_7 \mu \beta R.\end{aligned}$$

and that

$$\begin{aligned}\mathbb{E} \left[[N(T_A) - R] \mathbf{1}_{T_A \text{ during a fall}} \right] &\leq \frac{1}{C_4 p_2} \sum_{j=1}^R C_4 p_{\text{rise}}^{(j)} \left[C_3 \sqrt{\mu \beta R} + \mu j \beta \right] \prod_{\ell=0}^{j-1} \left(1 - C_4 p_{\text{rise}}^{(\ell)} \right) \\ &= \frac{1}{C_4 p_2} \mathbb{E} \left[C_3 \sqrt{\mu \beta R} + \mu \beta G \right] \\ &\leq \frac{1}{C_4 p_2} \left[C_3 \sqrt{\mu \beta R} + \mu \beta \left(A_5 \sqrt{R} + \frac{1}{0.99 C_4 A_5} \sqrt{R} \right) \right] \\ &= \left(B_5 \sqrt{\mu \beta R} + B_6 \mu \beta \sqrt{R} \right).\end{aligned}$$

Defining $F_1 \triangleq B_6$ and $F_2 \triangleq \frac{(B_5 + C_3)}{B_6}$, it follows that

$$\mathbb{E} [N(T_A) - R] \leq \left((B_5 + C_3) \sqrt{\mu \beta R} + B_6 \mu \beta \sqrt{R} \right) = F_1 \mu \beta \sqrt{R} \left(1 + \frac{F_2}{\sqrt{\mu \beta}} \right)$$

and that

$$\mathbb{E} \left[[N(T_A) - R]^2 \right] \leq \left(B_5 \sqrt{\mu \beta R} + B_6 \mu \beta \sqrt{R} \right)^2 + 2 B_7 \mu \beta R \leq F_1 (\mu \beta)^2 R \left(1 + \frac{F_2}{\sqrt{\mu \beta}} \right)^2 + 2 \mu \beta R \quad \square$$

Appendix H: Analysis of m -Policies

In this section, we study a natural policy that aims to mitigate the impact of setup times on waiting time. Our result reveals that this policy is inadequate unless it keeps significantly more servers on compared to the original policy described in Section 2; see Figure 9 for an illustration. This highlights the need for further studies on designing policies for turning servers on and off in queueing systems with setup times to achieve better tradeoffs between energy efficiency and performance.

The policies we study in this section are referred to as m -policies; we now describe how these policies operate. Recall that servers are indexed by integers $1, 2, \dots, k$, and that we, without loss of generality, re-index our collection of servers to ensure that each job departure is always from the highest indexed server among the busy servers. Under the original policy (described in Section 2), server i is turned off when the number of jobs in the system drops from i to $i - 1$. In contrast, an m -policy turns off servers less aggressively. Specifically, server i is turned off when the number of jobs in the system drops from $i - m$ to $i - m - 1$, where $0 \leq m \leq k$. Here m is referred to as the buffer size. The special case $m = 0$ corresponds to the original policy. The m -policies with $m > 0$ reduce job waiting times, as keeping m extra servers on ensures that the next m job arrivals do not need to wait for server setup.

The m -policies are more complicated to analyze compared to the original policy. One difficulty comes from the fact, that under an m -policy, servers which are *on* are not necessarily *busy*. To capture this distinction, we introduce a new variable $\text{On}(t)$ which tracks the number of servers that are *on* at time t . This quantity generally is not equal to the number of servers that are *busy* at time t , which we continue to call $Z(t)$. However, despite this difficulty, we are able to analyze m -policies using our MIST technique, with a more complicated coupling construction.

H.1. Main Result: A Lower Bound on m -Policies.

Our main result is that, unless the buffer size m is larger than $\Theta(\sqrt{R})$, the average waiting time under an m -policy will not be much smaller than the original ($m = 0$) policy. More specifically, we show the theorem below.

THEOREM 4 (Lower Bound on m -Policies). *When using an m -policy to control setup in an $M/M/k$ /Setup-Deterministic system with an offered load $R \triangleq k\rho \geq 100$, a setup time $\beta \geq 100\frac{1}{\mu}$, and an m -value with $m \leq \sqrt{R}$, the expected number of jobs in queue in steady state is lower-bounded by*

$$\mathbb{E}[Q(\infty)] \geq F_1 \left(\mu\beta\sqrt{R} + \frac{1}{1-\rho} \right),$$

for some constant F_1 independent of system parameters.




Figure 9 Simulation results showing the average waiting time for the M/M/k/Setup-Deterministic under various m -policies. All results obtained with mean service time $\frac{1}{\mu} = 1$ ms, mean setup time $\beta = 1000$ ms, and load kept at a constant $\rho = 0.5$. Note that a significant performance difference can only be found once the quantity m is growing faster than $R^{0.5}$.

Proof Overview. In the remainder of Section H, we provide a proof of the stated lower bound, Theorem 4. Our argument here is a variation on the argument made in Section G.6. We begin by defining our renewal points to be those moments when the number of *on* servers $\text{On}(t)$ transitions from $(R+1)$ to R ; note that this is slightly distinct from our previous definition since, in general, the number of busy servers $Z(t) \neq \text{On}(t)$. Taking X to be the next renewal point and applying the renewal reward theorem, it suffices to lower bound the renewal integral $\mathbb{E} \left[\int_0^X Q(t) dt \right]$ and upper bound the cycle length $\mathbb{E}[X]$. In what follows, we show the following two lemmas, which, via the above, suffice to prove the lower bound of Theorem 4.

LEMMA 8 (Lower Bound on Renewal Integral). *The renewal integral can be bounded by*

$$\mathbb{E} \left[\int_0^X [N(t) - R] dt \right] \geq F_2 \mu \beta^2 \sqrt{R} + I^{\text{busy}} \left(\left[F_2 \mu \beta \sqrt{R} - (k - R) \right], k - R \right),$$

where $I^{\text{busy}}(n, j) \triangleq \frac{n}{\mu j} \left[\frac{n+1}{2} + \frac{R}{j} + 1 \right]$ and $F_2 \triangleq 0.23$.

LEMMA 9 (Upper Bound on Cycle Length). *The length of a renewal cycle can be bounded by*

$$\mathbb{E}[X] \leq F_3 \beta + F_4 \frac{\mu \beta \sqrt{R}}{\mu \min(k - R, \sqrt{R})},$$

where $F_3 \triangleq 2.6$ and $F_4 \triangleq 7.2$.

Completion of Proof, assuming Lemmas 8 and 9. This result follows directly from the casework used in the proof of Theorem 3.

H.2. Proof of Lemma 8, the Integral Lower Bound.

High-Level Strategy. As mentioned previously, to prove this lower bound, we use a variation on the strategy used in Section G.6. In that proof, we first define a random quantity L which represents how many servers end up turning off at the beginning of a renewal cycle, then derive a bound on the desired integral in terms of its expected value $\mathbb{E}[L]$, then directly derive a bound on $\mathbb{E}[L]$. In this proof, we define an almost equivalent random variable L_{ex} and proceed with more or less the same steps; however, the more-complex behavior of the m -policy requires a more delicate analysis in order for us to find $\mathbb{E}[L_{\text{ex}}]$.

Definition of τ_{poss} . Before defining L_{ex} , we define a related stopping time τ_{poss} and prove a small but critical fact about it. Let $\tau_{\text{poss}} \triangleq \min \{t \geq 0 : Q(s) \geq 1 \text{ for all } s \in [t - \beta, t]\}$ be the first moment that the number of jobs in queue $Q(t)$ has been non-zero for a full setup time. At the beginning of a renewal cycle, since servers are setting up if and only if jobs are in the queue, it follows that the time τ_{poss} is the first time a server could possibly complete setup. Since T_A is the first time a specific server sets up, we have $\tau_{\text{poss}} \leq T_A$.

Definition of L_{ex} . With the stopping time τ_{poss} defined, we now define the random variable L_{ex} and describe its role going forward. Let $L_{\text{ex}} \triangleq \max_{s \in [0, \tau_{\text{poss}}]} [R - \text{On}(t)]$ be the maximum number of servers turned off in the period before time τ_{poss} . Now, via the same martingale arguments given in the proof of (17), we have that

$$\mathbb{E} \left[\int_0^{\tau_{\text{poss}}} Q(t) dt \right] \geq \frac{1}{2} \beta^2 \mu \mathbb{E}[L_{\text{ex}}]. \quad (51)$$

Likewise, in analogy to (20), by making a coupling argument we also have that

$$\mathbb{E} \left[\int_{\tau_{\text{poss}}}^X Q(t) dt \right] \geq I^{\text{busy}} \left([\mathbb{E}[N(\tau_{\text{poss}})] - k]^+, k - R \right), \quad (52)$$

where one should recall that $I^{\text{busy}}(n, j) \triangleq \frac{n}{\mu j} \left[\frac{n+1}{2} + \frac{R}{j} + 1 \right]$. After some martingale analysis (52) becomes

$$\mathbb{E} \left[\int_{\tau_{\text{poss}}}^X Q(t) dt \right] \geq I^{\text{busy}} \left([\beta \mathbb{E}[L_{\text{ex}}] - (k - R)]^+, k - R \right). \quad (53)$$

Thus, outside of some algebra, to prove Lemma 8 it suffices for us to give a bound on the expectation $\mathbb{E}[L_{\text{ex}}]$. We give this lower bound in the following claim, which we will prove immediately.

CLAIM 25 (Bound on L_{ex}). *For any $R \geq 2 \max(\sqrt{R}, m)$, we have*

$$\mathbb{E}[L_{\text{ex}}] \geq F_2 \min \left(\frac{R}{m}, \sqrt{R} \right),$$

where $F_2 \triangleq 0.23$.

H.3. Proof of Claim 25: Lower Bound on $\mathbb{E}[L_{\text{ex}}]$.

To bound $\mathbb{E}[L_{\text{ex}}]$, we upper bound the stopping probability $\Pr(L_{\text{ex}} = j | L_{\text{ex}} \geq j)$, showing that, for any $j \leq \max(\sqrt{R}, m)$,

$$\Pr(L_{\text{ex}} = j | L_{\text{ex}} \geq j) \leq 4 \frac{\max(m, \sqrt{R})}{R} \left(1 + \frac{b_1}{\sqrt{\mu\beta}}\right); \quad (54)$$

we defer the proof of this until the next section, Section H.3.1. From here, we can bound $\mathbb{E}[L_{\text{ex}}]$ using the usual sum-of-tails method, obtaining

$$\begin{aligned} \mathbb{E}[L_{\text{ex}}] &= \sum_{j=0}^{R-m} \Pr(L_{\text{ex}} \geq j+1) \\ &\geq \sum_{j=0}^{\max(\sqrt{R}, m)} \left(1 - 4 \frac{\max(m, \sqrt{R})}{R} \left(1 + \frac{b_1}{\sqrt{\mu\beta}}\right)\right)^j \\ &\geq \left[\min\left(\frac{R}{m}, \sqrt{R}\right)\right] \frac{1}{4\left(1 + \frac{b_1}{\sqrt{\mu\beta}}\right)} \left[1 - \left(1 - 4 \frac{\max(m, \sqrt{R})}{R} \left(1 + \frac{b_1}{\sqrt{\mu\beta}}\right)\right)^{\max(\sqrt{R}, m)}\right] \\ &\geq \left[\min\left(\frac{R}{m}, \sqrt{R}\right)\right] \frac{1}{4\left(1 + \frac{b_1}{\sqrt{\mu\beta}}\right)} [1 - e^{-4}] \\ &\geq F_2 \min\left(\frac{R}{m}, \sqrt{R}\right), \end{aligned}$$

where the constant $F_2 \triangleq 0.23 \leq \frac{1}{4\left(1 + \frac{b_1}{\sqrt{\mu\beta}}\right)} [1 - e^{-4}]$. To complete our proof, it suffices to show (54).

H.3.1. Proof of (54): Upper Bound on $\Pr(L_{\text{ex}} = j | L_{\text{ex}} \geq j)$. To show the inequality (54), we begin by giving an exact analysis of the probability $\Pr(L_{\text{ex}} = j | L_{\text{ex}} \geq j)$ by analyzing the sample paths of $N(t)$ at the beginning of a renewal cycle, then afterwards apply some straightforward Markov chain analysis. In particular, for two probabilities $p_{\text{start}}^{(j)}$ and $p_{\text{long}}^{(j)}$ and a parameter γ (to be defined later), we first show that

$$\Pr(L_{\text{ex}} = j | L_{\text{ex}} \geq j) \leq \min\left(\frac{p_{\text{long}}^{(j)}}{\gamma}, p_{\text{start}}^{(j)}\right). \quad (55)$$

Afterwards, we bound these quantities, showing that, for some $x \in (0, 1)$ and for all $j \leq \max(\sqrt{R}, m)$,

$$\min\left(\frac{p_{\text{long}}^{(j)}}{\gamma}, p_{\text{start}}^{(j)}\right) \leq 2 \frac{\max(m, \sqrt{R})}{R} \left(1 + \frac{b_1}{\sqrt{\mu\beta}}\right) \min\left(\frac{1}{x}, \frac{1}{1-x}\right). \quad (56)$$

Since $\min\left(\frac{1}{x}, \frac{1}{1-x}\right) \leq 2$, we obtain that

$$\Pr(L_{\text{ex}} = j | L_{\text{ex}} \geq j) \leq 4 \left(1 + \frac{b_1}{\sqrt{\mu\beta}}\right) \frac{\max(m, \sqrt{R})}{R},$$

as desired. Thus, it suffices to show (55) and (56).

H.3.2. Proof of (55). As previously stated, we begin by giving an exact analysis of the probability $\Pr(L_{\text{ex}} = j | L_{\text{ex}} \geq j)$. We make a recursive argument, based on a sequence of events (and their implicitly-defined stopping times). Note that, if we were to write this explicitly, then we would need to invoke our main lemma, Lemma 4.

First set of events. Assume that one begins at time τ_j , the first moment where one knows that $L_{\text{ex}} \geq j$; note that, since at time τ_j we have just turned off the $(R - j + 1)$ -th server, we know the number of jobs $N(t) = R - j + 1 - m$. From here, one can wait until one of two events happen: either 1) the number of jobs $N(t)$ decreases to $(R - j - m)$, at which point $L_{\text{ex}} > j$; or 2) the number of jobs $N(t)$ increases to $(R - j + 1)$, at which point the event $[L_{\text{ex}} = j]$ still has a chance to occur. We call the probability of the latter event $p_{\text{start}}^{(j)}$; we think of this as a “starting” probability.

Second set of events. At the moment when $N(t)$ reaches $(R - j + 1)$, the number of jobs in the queue $Q(t) = 1$. As such, we can again partition our sample paths based on whether: 1) $Q(t)$ stays non-zero for a full setup time, at which point $[L_{\text{ex}} = j]$; or 2) $Q(t)$ drops to 0 within a setup time. We call the probability of the first event $p_{\text{long}}^{(j)}$.

Third set of events. In the case where the queue empties before a setup time has passed, the number of jobs $N(t) = (R - j)$ and we again can condition on the next event to happen: either 1) the number of jobs $N(t)$ decreases to $(R - j - m)$, at which point $L_{\text{ex}} > j$; or 2) the number of jobs $N(t)$ increases to $(R - j + 1)$, at which point the event $[L_{\text{ex}} = j]$ still has a chance to occur. We call the latter probability $p_{\text{re}}^{(j)}$, since it can be thought of as a “restart” probability. Note that, in the case where the number of jobs $N(t)$ rises, we have reached an identical state to that which initiates the second set of events; thus we can induct.

Result in terms of probabilities. From this analysis, it follows that

$$\begin{aligned} \Pr(L_{\text{ex}} = j | L_{\text{ex}} \geq j) &= p_{\text{start}}^{(j)} \left(p_{\text{long}}^{(j)} + (1 - p_{\text{long}}^{(j)}) p_{\text{re}}^{(j)} \left(p_{\text{long}}^{(j)} + (1 - p_{\text{long}}^{(j)}) p_{\text{re}}^{(j)} (\dots) \right) \right) \\ &= p_{\text{start}}^{(j)} p_{\text{long}}^{(j)} \frac{1}{1 - \left(1 - p_{\text{long}}^{(j)}\right) p_{\text{re}}^{(j)}}. \end{aligned}$$

By noting that the behavior of $N(t)$ in the first and third scenarios is sample-path-equivalent to the behavior of an $M/M/(R - j)$, and that the behavior of $N(t)$ in the second scenario is sample-path-equivalent to that of an (overloaded) $M/M/1$, we can analyze these probabilities. In Section I.2, we show via basic Markov chain analysis that $p_{\text{re}}^{(j)} = 1 - \gamma p_{\text{start}}^{(j)}$, where $\gamma \triangleq \frac{(R-j)!}{(R-j-m)!} \frac{1}{R^m} \leq 1$. It follows that both

$$p_{\text{start}}^{(j)} p_{\text{long}}^{(j)} \frac{1}{1 - \left(1 - p_{\text{long}}^{(j)}\right) p_{\text{re}}^{(j)}} = p_{\text{start}}^{(j)} \frac{p_{\text{long}}^{(j)}}{p_{\text{long}}^{(j)} + \left(1 - p_{\text{long}}^{(j)}\right) \gamma p_{\text{start}}^{(j)}} \leq p_{\text{start}}^{(j)}$$

and

$$p_{\text{start}}^{(j)} p_{\text{long}}^{(j)} \frac{1}{1 - \left(1 - p_{\text{long}}^{(j)}\right) p_{\text{re}}^{(j)}} = p_{\text{long}}^{(j)} \frac{p_{\text{start}}^{(j)}}{\gamma p_{\text{start}}^{(j)} + \left(1 - \gamma p_{\text{start}}^{(j)}\right) p_{\text{long}}^{(j)}} \leq \frac{p_{\text{long}}^{(j)}}{\gamma},$$

and thus that

$$\Pr(L_{\text{ex}} = j | L_{\text{ex}} \geq j) \leq \min\left(\frac{p_{\text{long}}^{(j)}}{\gamma}, p_{\text{start}}^{(j)}\right),$$

as desired.

H.3.3. Proof of (56): Upper Bound on $p_{\text{start}}^{(j)}$ and $p_{\text{long}}^{(j)}$. To complete our analysis of $\Pr(L_{\text{ex}} = j | L_{\text{ex}} \geq j)$, it suffices to bound the probabilities $p_{\text{long}}^{(j)}$ and $p_{\text{start}}^{(j)}$, and the constant γ . To upper bound the probability $p_{\text{long}}^{(j)}$, we note that, from the analysis in Section G.6.1,

$$\begin{aligned} p_{\text{long}}^{(j)} &\leq 1 - \left(1 - \frac{j}{R}\right) \left(1 - \frac{b_1}{\sqrt{\mu\beta R}}\right) \\ &= \frac{b_1}{\sqrt{\mu\beta R}} + \frac{j}{R} \left(1 - \frac{b_1}{\sqrt{\mu\beta R}}\right) \\ &\leq \frac{\left(j + b_1 \frac{\sqrt{R}}{\sqrt{\mu\beta}}\right)}{R} \\ &\leq \left(1 + \frac{b_1}{\sqrt{\mu\beta}}\right) \frac{\sqrt{R}}{R} \\ &\leq \left(1 + \frac{b_1}{\sqrt{\mu\beta}}\right) \frac{2 \max(\sqrt{R}, m)}{R}. \end{aligned}$$

To bound the constant γ , we note that, defining $x \triangleq \left(1 - \frac{j+m}{R}\right)^m$,

$$\gamma \triangleq \frac{(R-j)!}{(R-j-m)!} \frac{1}{R^m} \geq \left(\frac{R-j-m}{R}\right)^m = \left(1 - \frac{j+m}{R}\right)^m \triangleq x.$$

To upper bound the probability $p_{\text{start}}^{(j)}$, we use the analysis from Section I.2, which shows that

$$\begin{aligned} p_{\text{start}}^{(j)} &\leq \frac{j+m}{R} \frac{1}{1 - \left(1 - \frac{j+m}{R}\right)^{m+1}} \\ &\leq \frac{j+m}{R} \frac{1}{1 - \left(1 - \frac{j+m}{R}\right)^m} \\ &= \frac{j+m}{R} \frac{1}{1-x} \\ &\leq \frac{2 \max(\sqrt{R}, m)}{R} \frac{1}{1-x} \\ &\leq \left(1 + \frac{b_1}{\sqrt{\mu\beta R}}\right) \frac{2 \max(\sqrt{R}, m)}{R} \frac{1}{1-x}. \end{aligned}$$

Thus, from (55), we have that

$$\begin{aligned} \Pr(L_{\text{ex}} = j | L_{\text{ex}} \geq j) &\leq \min\left(\frac{p_{\text{long}}^{(j)}}{\gamma}, p_{\text{start}}^{(j)}\right) \\ &\leq \left(1 + \frac{b_1}{\sqrt{\mu\beta R}}\right) \frac{2 \max(\sqrt{R}, m)}{R} \min\left(\frac{1}{x}, \frac{1}{1-x}\right) \\ &\leq 4 \left(1 + \frac{b_1}{\sqrt{\mu\beta}}\right) \frac{\max(m, \sqrt{R})}{R}, \end{aligned}$$

as desired. Thus, we complete our proof of Claim 25, our lower bound on $\mathbb{E}[L_{\text{ex}}]$.

H.4. Proof of Lemma 9, the Cycle Length Upper Bound.

Preliminaries: Cycle decomposition. We now prove Lemma 9, an upper bound on the expected cycle length $\mathbb{E}[X]$. As before, we begin by breaking a renewal cycle into three parts; we later analyze each part separately. To do so, we define a threshold H and two stopping times; the latter stopping time will depend on H . To begin, noting that R is the minimum number of servers needed to stabilize the system, we define the threshold $H \triangleq \min(k - R, \sqrt{R})$ as the minimum between $(k - R)$, the number of “extra” servers, and \sqrt{R} , a quantity which represents the “typical variation” in the system. Next, we define the accumulation time $T_A \triangleq \min\{t \geq 0 : \text{On}(t) \geq R + 1\}$ as the first moment that the $(R + 1)$ -th server turns on. Finally, we define the draining time $T_B \triangleq \min\{t \geq T_A : N(t) \leq R + H\}$ as the first moment after time T_A that the number of jobs $N(t)$ dips below the threshold $(R + H)$.

Preliminaries: Definition of phases. With these two stopping times defined, our renewal cycle takes on a similar decomposition as before. From time 0 to time T_A , the number of servers that are *on* is $\leq (R + 1)$, meaning that the departure rate of jobs must be less than the arrival rate of jobs. Accordingly, jobs accumulate and thus we call this the *accumulation* phase. From time T_A to time T_B , the number of servers *on* is $\geq (R + 1)$, meaning that, on the whole, the accumulated jobs are *draining* from the system; thus we call this the *draining* phase. From time T_B to the end of the cycle at time X , the behavior of the system roughly resembles that of a balanced random walk; we thus call this period the *balanced* phase.

Proof strategy. Using this decomposition, we prove the following bounds on the expected length of each phase: for the accumulation phase $[0, T_A)$

$$\mathbb{E}[T_A] \leq 1.08\beta, \quad (57)$$

for the draining phase $[T_A, T_B)$

$$\mathbb{E}[T_B - T_A] \leq \beta + E_1 \frac{\mu\beta\sqrt{R}}{\mu H}, \quad (58)$$

and for the balanced phase $[T_B, X)$

$$\mathbb{E}[X - T_B] \leq E_3 \frac{1}{\mu} + E_{10}\beta. \quad (59)$$

Putting these together, we confirm that

$$\mathbb{E}[X] \leq F_3\beta + F_4 \frac{\mu\beta\sqrt{R}}{\mu(k - R)}.$$

In what follows, we briefly discuss the proof of each of these inequalities in turn.

H.4.1. Proof of (57): Bound on $\mathbb{E}[T_A]$. For the bound on $\mathbb{E}[T_A]$, our previous proof still applies: since $\text{On}(t) \leq R$ until time T_A , we can couple the system to an M/M/R queue and bound the length of an analogous version of T_A in that system. To review the argument, we can alternatively define $T_A = \min \{t \geq 0 : \min_{s \in [t-\beta, t]} N(t) \geq R+1\}$ as the first moment that the number of jobs $N(t)$ stays above $(R+1)$ for a full setup time β . If we let $\tilde{N}(t)$ denote the number of jobs in a coupled M/M/R queue starting with $\tilde{N}(0) = (R-m)$, then we can accordingly define \tilde{T}_A as the first moment that the coupled number of jobs $\tilde{N}(t)$ stays above $(R+1)$ for a full setup time β . Since the departure rate of the original system is upper-bounded by the departure rate of this coupled M/M/R system, the original system always contains more jobs, i.e. $N(t) \geq \tilde{N}(t)$. It follows that $T_A \leq \tilde{T}_A$ and thus it suffices to bound $\mathbb{E}[\tilde{T}_A]$; Section B.5 tells us that

$$\mathbb{E}[T_A] \leq \mathbb{E}[\tilde{T}_A] \leq 1.08\beta,$$

which completes the proof.

H.4.2. Proof of (58): Bound on $\mathbb{E}[T_B - T_A]$. For the bound on $\mathbb{E}[T_B - T_A]$, our strategy again is to derive a bound on the conditional expectation of $[T_B - T_A]$, where we condition on the number of jobs in the system at time T_A , and afterwards give a bound on both the first and second moments of this number of jobs $N(T_A)$. Actually, in this case, our conditional bound from Section 6.4 still applies; by making worst-case assumptions about the system's setup state at time T_A and noting that, if the number of jobs stays large for a full setup time then many servers must be *on*, we can use a combination of random walk and busy period analysis to obtain

$$\mathbb{E}[T_B - T_A | \mathcal{F}_{T_A}] \leq \beta + \frac{N(T_A) - R}{\mu H} + \frac{1}{\mu} \ln \left(\frac{N(T_A) - R}{H} \right); \quad (60)$$

see Section 6.4 for details. To complete our bound of the draining phase, we must bound the expected value $\mathbb{E}[N(T_A)]$; we defer the proof of this bound to Section I.1. There, we find that

$$\mathbb{E}[N(T_A) - R] \leq E_1 \mu \beta \sqrt{R}, \quad (61)$$

where $E_1 \triangleq 5.21$. Thus, combining (61) with (60) (and noting that $\frac{x+\ln(x)}{x} \leq 1 + \frac{1}{e}$), we find that

$$\begin{aligned} \mathbb{E}[T_B - T_A] &\leq \beta + \mathbb{E} \left[\frac{N(T_A) - R}{\mu H} \right] + \frac{1}{\mu} \mathbb{E} \left[\ln \left(\frac{N(T_A) - R}{H} \right) \right] \\ &\leq \beta + \frac{\mathbb{E}[N(T_A) - R]}{\mu H} + \frac{1}{\mu} \ln \left(\frac{\mathbb{E}[N(T_A) - R]}{H} \right) \\ &\leq \beta + \left(1 + \frac{1}{e} \right) \frac{\mathbb{E}[N(T_A) - R]}{\mu H} \\ &\leq \beta + 1.37 E_1 \frac{\mu \beta \sqrt{R}}{\mu H}. \end{aligned}$$

H.4.3. Proof of (59): Bound on $\mathbb{E}[X - T_B]$.

Main idea: Threshold-based division. To bound the final portion of the renewal cycle, $\mathbb{E}[X - T_B]$, we divide time based on the threshold H and make a coupling argument. More specifically, we make an argument based on whether the number of jobs $N(t)$ is *below* or *above* the threshold $(R + H)$. First, we note that, when the number of jobs $N(t)$ is *below* $(R + H)$, the departure rate of jobs is lower-bounded by $\mu \min(R, N(t))$. Said differently, while below $(R + H)$, the number of jobs $N(t)$ is upper-bounded by the number of jobs in an appropriately-coupled M/M/R/($R + H$) queue. On the other hand, when the number of jobs $N(t)$ becomes *larger* than $(R + H)$, we can use the “wait-busy” period idea of the previous phase to bound the amount of time it takes for number of jobs $N(t)$ to come back down to $(R + H)$; we call this an excursion.

The below-threshold case. To analyze the below-threshold portion of time, we make a coupling argument based on our above departure rate bound. If we let $T_{(R+H) \rightarrow (R-m)}^{M/M/R/(R+H)}$ to be the time it takes for an M/M/R/($R + H$) to go from $(R + H)$ jobs in the system to $(R - m)$ jobs in the system, then, from a straightforward coupling argument, the time spent below $(R + H)$ can be upper-bounded by $\mathbb{E}\left[T_{(R+H) \rightarrow (R-m)}^{M/M/R/(R+H)}\right]$. As such, it suffices to bound the expectation of this first passage time, which, for $m \leq \sqrt{R}$, was done in Claim 22. This gives

$$\begin{aligned} \mathbb{E}[\text{time spent with } N(t) \leq R + H] &\leq \mathbb{E}\left[T_{(R+H) \rightarrow (R-m)}^{M/M/R/(R+H)}\right] \\ &\leq 7 \frac{m + H}{\mu \sqrt{R}} \\ &\leq 7 \frac{\sqrt{R} + \min(\sqrt{R}, k - R)}{\sqrt{R}} \frac{1}{\mu} \\ &\leq 14 \frac{1}{\mu}. \end{aligned}$$

The above-threshold case. To analyze the above-threshold portion of time, we bound the number of excursions and use the “wait-busy” idea to bound the length of an excursion. To bound the number of excursions, we note (by our below-threshold coupling argument) that the probability of having an additional excursion is upper-bounded by the probability that, in an M/M/R queue started with $(R + H)$ jobs in the system, the queue reaches $(R + H + 1)$ jobs in system before it reaches $(R - m)$ jobs in system; we call this probability p_e . Accordingly, the number of excursions is stochastically dominated by a $\text{Geo}(1 - p_e) - 1$ random variable, and thus the expected number of excursions is bounded by $\frac{1}{1 - p_e}$; we give an explicit bound on p_e later, showing that

$$\frac{1}{1 - p_e} \leq \sqrt{R} \left(1 + e^{\frac{m(m+1)}{R-m}}\right), \quad (62)$$

which, for $m \leq \sqrt{R}$ and $R \geq 100$, can be bounded by

$$\frac{1}{1 - p_e} \leq \sqrt{R} \left(1 + e^{\frac{\sqrt{R}+1}{\sqrt{R}-1}}\right) \leq 4.4\sqrt{R}.$$

Next, by using the “wait-busy” idea of the previous phase, previous arguments tell us that, for $b_1 = \sqrt{\frac{2}{\pi}}$,

$$\mathbb{E}[\text{length of an excursion}] \leq \frac{1}{\mu} b_1 \sqrt{\frac{\mu\beta}{R}} + \frac{1}{\mu H}. \quad (63)$$

Thus, we obtain that for $\mu\beta \geq 100$,

$$\begin{aligned} \mathbb{E}[\text{time spent with } N(t) > R + H] &\leq \frac{1}{1 - p_e} \cdot \left[\frac{1}{\mu} b_1 \sqrt{\frac{\mu\beta}{R}} + \frac{1}{\mu H} \right] \\ &\leq \beta \frac{4.4b_1}{\sqrt{\mu\beta}} + \frac{4.4\sqrt{R}}{\mu H} \\ &\leq 0.36\beta + \frac{4.4\sqrt{R}}{\mu H}. \end{aligned}$$

We arrive at the following bound on $\mathbb{E}[X - (T_A + T_B)]$:

$$\begin{aligned} \mathbb{E}[X - T_B] &\leq \mathbb{E}[\text{time spent with } N(t) \leq R + H] + \mathbb{E}[\text{time spent with } N(t) > R + H] \\ &\leq 14 \frac{1}{\mu} + 0.36\beta + \frac{4.4\sqrt{R}}{\mu H} \\ &\leq 0.5\beta + \frac{4.4\sqrt{R}}{\mu H}. \end{aligned}$$

H.4.4. Proof of Lemma 9: Closing. As mentioned, with (57),(58), and (59) proven, we complete our bound on $\mathbb{E}[X]$ by simple addition:

$$\begin{aligned} \mathbb{E}[X] &= \mathbb{E}[T_A] + \mathbb{E}[T_B - T_A] + \mathbb{E}[X - T_B] \\ &\leq 1.08\beta + \beta + 1.37E_1 \frac{\mu\beta\sqrt{R}}{\mu H} + 0.5\beta + \frac{4.4\sqrt{R}}{\mu H} \\ &\leq 2.6\beta + 7.2 \frac{\mu\beta\sqrt{R}}{\mu H} \\ &= F_3\beta + F_4 \frac{\mu\beta\sqrt{R}}{\mu H}, \end{aligned}$$

where we have taken $F_3 \triangleq 2.6$ and $F_4 \triangleq 7.2$.

Appendix I: Miscellaneous Claims: Analysis of m -Policies

I.1. Proof of (61): Upper Bound on $\mathbb{E}[N(T_A) - R]$ for m -Policies

Here, we discuss our bound on the expected value of the number of jobs $\mathbb{E}[N(T_A) - R]$ for m -policies. In particular, we show that

$$\mathbb{E}[N(T_A) - R] \leq E_1 \mu\beta\sqrt{R},$$

where $E_1 \triangleq 5.21$.

Definition of epoch. To do so, we again make use of the notion of *epochs*. Recall that time $\tau_j \triangleq \min \{t \geq 0 : \text{On}(t) \leq R - j\}$ is the first moment that only $(R - j)$ servers are on. We call the period $[\tau_j, \min(\tau_{j+1}, T_A))$ the j -th epoch, and we say epoch j occurs if the time τ_j happens before time T_A . As a shorthand, we use n_e to denote the number of epochs that occur; note that the condition [epoch j occurs] is equivalent to the condition $[n_e \geq j]$.

Definition of a trial. To analyze the time T_A more closely, we define a series of trials. The first trial in epoch j (that is, the $(1, j)$ -th trial) begins at the first moment that $N(t) \geq R + C_3\sqrt{\mu\beta R}$ during epoch j . We say this trial *succeeds* if the number of jobs $N(t)$ does not drop below $(R + 1)$ for a full setup time; otherwise, we say that this trial *fails*. If, in the same epoch, the number of jobs $N(t)$ again reaches $(R + 1)$, then the next trial begins and may either succeed or fail; this continues until either the next epoch begins or a trial succeeds. We refer to the i -th trial to occur in epoch j as trial (i, j) .

We now note that

$$\mathbb{E}[N(T_A) - R] \leq C_3\sqrt{\mu\beta R} + \sum_{i,j} \Pr(\text{trial } (i, j) \text{ occurs}) \mathbb{E}[[N(T_A) - R] \mathbf{1}_{\text{trial } (i, j) \text{ succeeds}} | \text{trial } (i, j) \text{ occurs}]. \quad (64)$$

To analyze this conditional expectation, we note that, if trial (i, j) occurs at time $\phi_{i,j}$ and this trial is successful, then time $T_A = (\phi_{i,j} + \beta)$. On the other hand, let $\psi_{i,j} \triangleq \min \{t > 0 : N(\phi_{i,j} + t) < R + 1\}$ be the next moment the number of jobs $N(t)$ falls below $(R + 1)$; we make two points. First, the trial succeeds if and only if $\psi_{i,j} \geq \beta$. Second, if the trial fails, then $[N(\psi_{i,j}) - R] = 0$. As such,

$$[N(T_A) - R] \mathbf{1}_{\text{trial } (i, j) \text{ succeeds}} = [N(\phi_{i,j} + \min(\psi_{i,j}, \beta)) - R].$$

Noting that $V(t) \triangleq N(\phi_{i,j} + t) - R - \mu jt$ is a super-martingale (since we are in epoch j), it follows from Doob's Optional Stopping Theorem that

$$\begin{aligned} V(0) &= C_3\sqrt{\mu\beta R} \\ &\geq \mathbb{E}[V(\min(\psi_{i,j}, \beta))] \\ &= \mathbb{E}[[N(\phi_{i,j} + \min(\psi_{i,j}, \beta)) - R] | \mathcal{F}_{\phi_{i,j}}] - \mu j \mathbb{E}[\min(\psi_{i,j}, \beta) | \mathcal{F}_{\phi_{i,j}}], \end{aligned}$$

and thus that

$$\mathbb{E}[[N(T_A) - R] \mathbf{1}_{\text{trial } (i, j) \text{ succeeds}} | \text{trial } (i, j) \text{ occurs}] \leq C_3\sqrt{\mu\beta R} + \mu j \mathbb{E}[\min(\psi_{i,j}, \beta) | \text{trial } (i, j) \text{ occurs}].$$

Reduction to trial quantities. Applying this to (64),

$$\begin{aligned} \mathbb{E}[N(T_A) - R] &= \sum_{i,j} \Pr(\text{trial } (i, j) \text{ occurs}) \mathbb{E}[[N(T_A) - R] \mathbf{1}_{\text{trial } (i, j) \text{ succeeds}} | \text{trial } (i, j) \text{ occurs}] \\ &\leq \sum_{i,j} \Pr(\text{trial } (i, j) \text{ occurs}) \left[C_3\sqrt{\mu\beta R} + \mu j \mathbb{E}[\min(\psi_{i,j}, \beta) | \text{trial } (i, j) \text{ occurs}] \right] \\ &\leq C_3\sqrt{\mu\beta R} \mathbb{E}[\text{total number of trials}] + \sum_j \mu j \mathbb{E}[\text{total time spent on trials during epoch } j]. \end{aligned}$$

To address the first term, we note that for each trial which occurs, according to Claim 3, with probability at least $\left[1 - 2\Phi\left(-\frac{3}{\sqrt{2}}\right) - \frac{1}{100}\right] \geq 0.95 \triangleq C_4$ the first phase ends. It follows that $\mathbb{E}[\text{total number of trials}] \leq \frac{1}{0.95}$.

To address the second term, we note that

$$\begin{aligned} & \mathbb{E}[\text{total time spent on trials during epoch } j] \\ &= \Pr(\text{trial } (1, j) \text{ occurs}) \mathbb{E}[\text{time spent on trials during epoch } j | \text{trial } (1, j) \text{ occurs}] \\ &\leq \Pr(\text{trial } (1, j) \text{ occurs}) \mathbb{E}[\text{time remaining in first phase} | \text{trial } (1, j) \text{ occurs}] \\ &\leq \Pr(\text{trial } (1, j) \text{ occurs}) \mathbb{E}[\tilde{T}_A] \\ &\leq \Pr(\text{trial } (1, j) \text{ occurs}) 1.08\beta, \end{aligned}$$

where the final inequality is via coupling the system to an $M/M/R$ queue in the same way as in Section B.5.

Defining $G_1 \triangleq 3.3 > \frac{1}{0.95}C_3$ and $G_2 \triangleq 1.08$, we find that

$$\mathbb{E}[N(T_A) - R] \leq G_1 \sqrt{\mu\beta R} + G_2 \mu\beta \sum_j j \Pr(\text{trial } (1, j) \text{ occurs}).$$

Development using epochs. From here, we note that $p_{\text{rise}}^{(j)} \triangleq \Pr(\text{trial } (1, j) \text{ occurs} | n_e \geq j)$ and thus it suffices to bound $\Pr(n_e \geq j)$. To provide an upper bound on the probability that epoch j , it suffices to show a lower bound on the conditional probability $\Pr(n_e = j | n_e \geq j)$. In Section G.3, we have already shown that

$$\Pr(n_e = j | n_e \geq j) \geq (1 - C_4 p_{\text{rise}}^{(j)});$$

applying this bound, one finds that

$$\Pr(n_e \geq j) \leq \prod_{i=0}^{j-1} (1 - C_4 p_{\text{rise}}^{(i)}).$$

Thus, we obtain

$$\begin{aligned} \sum_j j \Pr(\text{trial } (1, j) \text{ occurs}) &\leq \frac{1}{C_4} \sum_j C_4 p_{\text{rise}}^{(j)} j \prod_{i=0}^{j-1} (1 - p_{\text{rise}}^{(i)}) \\ &= \frac{1}{C_4} \sum_j \prod_{i=0}^j (1 - C_4 p_{\text{rise}}^{(i)}). \end{aligned}$$

From here, it suffices to lower bound $p_{\text{rise}}^{(i)}$.

I.1.1. Lower bound on $p_{\text{rise}}^{(j)}$. To bound the probability that the number of jobs $N(t)$ reaches the threshold $R + C_3 \sqrt{\mu\beta R}$ at some point during epoch j , we focus on one particular manner in which that event might happen. To define this manner, we introduce the notion of an *excursion*. We do so recursively: the first excursion in epoch j occurs the moment the number of jobs reaches $(R - j + 1)$, and ends when either

1) the number of jobs $N(t)$ dips back down to $(R - j)$ or 2) a full setup time has passed. If a full setup time passes during an excursion, we call it *long*; otherwise, we call it *short*. Successive excursions occur whenever an up-crossing of $(R - j)$ occurs, but only if every previous excursion in the epoch has been *short*. Thus, we can lower bound the probability of reaching the desired threshold by requiring that it occurs before the epoch's *first long excursion*, i.e.

$$\begin{aligned} p_{\text{rise}}^{(j)} &\geq p_{\text{start}}^{(j)} \left(p_{\text{th}} + p_f p_{\text{re}}^{(j)} p_{\text{th}} + (p_f p_{\text{re}}^{(j)})^2 p_{\text{th}} + \dots \right) \\ &= p_{\text{th}} p_{\text{start}}^{(j)} \frac{1}{1 - p_f p_{\text{re}}^{(j)}} \end{aligned}$$

where we have used the fact that the system state is the same at the beginning of every excursion, and where we have defined p_{th} as the probability that the number of jobs $N(t)$ crosses the threshold during excursion $(1, j)$, defined $p_{\text{re}}^{(j)}$ as the probability that, after a short excursion, another excursion occurs, defined $p_{\text{start}}^{(j)}$ as the probability that an excursion occurs, and defined p_f as the probability that the excursion is both short and ends without $N(t)$ having reached the desired threshold. We arrive at the following claim, proven in Section I.4:

CLAIM 26. *Recall the definitions of p_{th} , p_f , $p_{\text{re}}^{(j)}$, and $p_{\text{start}}^{(j)}$. Then, for $j \geq \sqrt{R}$, we have the following bound*

$$p_{\text{th}} p_{\text{start}}^{(j)} \frac{1}{1 - p_f p_{\text{re}}^{(j)}} \geq E_6 \frac{j}{R} \geq E_6 \frac{1}{\sqrt{R}},$$

where $E_6 = 0.32$.

Accordingly, we find that

$$\begin{aligned} \mathbb{E}[N(T_A) - R] &\leq G_1 \sqrt{\mu\beta R} + G_2 \mu\beta \frac{1}{C_4} \sum_{j=0}^{R-m} \prod_{i=0}^j (1 - C_4 p_{\text{rise}}^{(i)}) \\ &\leq G_1 \sqrt{\mu\beta R} + G_2 \mu\beta \frac{1}{C_4} \left[\sqrt{R} \sum_{j=\sqrt{R}}^{R-m} \prod_{i=0}^j (1 - C_4 p_{\text{rise}}^{(i)}) \right] \\ &\leq G_1 \sqrt{\mu\beta R} + G_2 \mu\beta \frac{1}{C_4} \left[\sqrt{R} + \sum_{j=\sqrt{R}} \left(1 - C_4 E_6 \frac{1}{\sqrt{R}} \right)^{[j-\sqrt{R}]^+} \right] \\ &= G_1 \sqrt{\mu\beta R} + \mu\beta \sqrt{R} \frac{G_2}{C_4} \left[1 + \frac{1}{C_4 E_6} \right] \\ &= \mu\beta \sqrt{R} \left(\frac{G_2}{C_4} \left[1 + \frac{1}{C_4 E_6} \right] + \frac{G_1}{\sqrt{\mu\beta}} \right) \\ &\leq E_1 \mu\beta \sqrt{R} \end{aligned}$$

as desired, where we have taken $E_1 \triangleq 5.21 > \left(\frac{G_2}{C_4} \left[1 + \frac{1}{C_4 E_6} \right] + \frac{G_1}{\sqrt{100}} \right)$. Thus, we have shown (61), an upper bound on the quantity $\mathbb{E}[N(T_A) - R]$ for m -policies.

I.2. Bounds on hitting probabilities for an M/M/ ∞ queue.

We now analyze a general hitting probability problem in the M/M/ ∞ queue. In particular, we consider the case where the arrival rate is $k\lambda = \mu R$, and define an upper and lower boundary of b and a respectively. We then find the probability that, starting from state i , we hit the upper boundary state b before we hit the lower boundary point a ; we call this probability p_i^{up} .

I.2.1. Solving for p_i^{up} . We solve for these probabilities using the usual finite difference method. We first note our boundary conditions: the probabilities at the boundaries are $p_a^{\text{up}} = 0$ and $p_b^{\text{up}} = 1$. From here, we note that, by a conditioning argument,

$$p_i^{\text{up}} = \frac{R}{R+i} p_{i+1}^{\text{up}} + \frac{i}{R+i} p_{i-1}^{\text{up}}.$$

After some algebra, we obtain that, defining $\Delta_i = p_{i+1}^{\text{up}} - p_i^{\text{up}}$,

$$\frac{p_{i+1}^{\text{up}} - p_i^{\text{up}}}{p_i^{\text{up}} - p_{i-1}^{\text{up}}} = \frac{i}{R} = \frac{\Delta_i}{\Delta_{i-1}}.$$

This shows that

$$\Delta_i = \frac{i}{R} \cdot \frac{i-1}{R} \cdots \frac{a+1}{R} \cdot \Delta_a = \frac{(i)!}{(a)! R^{i-a}} \Delta_a.$$

Using our upper boundary condition, we have

$$\begin{aligned} 1 &= p_b^{\text{up}} = (p_b^{\text{up}} - p_{b-1}^{\text{up}}) + (p_{b-1}^{\text{up}} - p_{b-2}^{\text{up}}) + \cdots + (p_{a+2}^{\text{up}} - p_{a+1}^{\text{up}}) + (p_{a+1}^{\text{up}}) \\ &= \sum_{i=a}^{b-1} \Delta_i \\ &= \Delta_a \sum_{i=a}^{b-1} \frac{(i)!}{(a)! R^{i-a}}. \end{aligned}$$

Thus we have that

$$\Delta_a = p_{a+1}^{\text{up}} = \frac{1}{\sum_{i=a}^{b-1} \frac{(i)!}{(a)! R^{i-a}}} = \frac{1}{\sum_{i=0}^{b-1-a} \frac{(i+a)!}{(a)! R^i}}. \quad (65)$$

Notably, we also have that

$$p_{b-1}^{\text{up}} = 1 - \Delta_{b-1} = 1 - \frac{(b-1)!}{(a)! R^{b-1-a}} \Delta_a. \quad (66)$$

I.2.2. Bounds on Δ_a . We now give two upper bounds and a lower bound on Δ_a by bounding the sum in its denominator, using $a = R - j - m$ and $b = R - j + 1$. For the first upper bound on the denominator,

$$\begin{aligned} \sum_{i=0}^{b-1-a} \frac{(i+a)!}{(a)! R^i} &= \sum_{i=0}^m \frac{(R-j-m+i)!}{(R-j-m)! R^i} \\ &= \sum_{i=0}^m \prod_{\ell=1}^i \frac{R-j-m+\ell}{R} \\ &\leq \sum_{i=0}^m \prod_{\ell=1}^i e^{-\frac{(j+m)}{R} + \frac{\ell}{R}} \end{aligned}$$

$$\begin{aligned}
 &= \sum_{i=0}^m e^{-\frac{(j+m)i}{R} + \frac{i(i+1)}{2R}} \\
 &= \sum_{i=0}^m e^{-\frac{i}{R} [j+m - \frac{i+1}{2}]} \\
 &\leq \sum_{i=0}^m e^{-\frac{i}{R} [j + \frac{1}{2}m - \frac{1}{2}]} \\
 &\leq \frac{1}{1 - e^{-\frac{j + \frac{1}{2}(m-1)}{R}}}.
 \end{aligned}$$

Thus, we find that

$$p_{a+1}^{\text{up}} \geq 1 - e^{-\frac{j + \frac{1}{2}(m-1)}{R}}. \quad (67)$$

For the second upper bound on the denominator,

$$\begin{aligned}
 \sum_{i=0}^{b-1-a} \frac{(i+a)!}{(a)!R^i} &= \sum_{i=0}^m \frac{(R-j-m+i)!}{(R-j-m)!R^i} \\
 &= \sum_{i=0}^m \prod_{\ell=1}^i \frac{R-j-m+\ell}{R} \\
 &\leq \sum_{i=0}^m \left(\frac{R-j}{R}\right)^i \\
 &= \frac{R}{j} \left(1 - \left(1 - \frac{j}{R}\right)^{m+1}\right) \\
 &\leq \frac{R}{j}.
 \end{aligned}$$

This in turn gives

$$p_{a+1}^{\text{up}} \geq \frac{j}{R}. \quad (68)$$

Likewise, for a lower bound on the denominator,

$$\begin{aligned}
 \sum_{i=0}^m \prod_{\ell=1}^i \frac{R-j-m+\ell}{R} &\geq \sum_{i=0}^m \left(1 - \frac{(j+m)}{R}\right)^i \\
 &= \frac{R}{j+m} \left(1 - \left(1 - \frac{(j+m)}{R}\right)^{m+1}\right).
 \end{aligned}$$

Thus, we find that

$$p_{a+1}^{\text{up}} \leq \frac{j+m}{R} \frac{1}{\left(1 - \left(1 - \frac{(j+m)}{R}\right)^{m+1}\right)}. \quad (69)$$

I.2.3. Bounds on p_{b-1}^{up} . Likewise, since $p_{b-1}^{\text{up}} = 1 - \delta_{b-1} = 1 - \frac{(R-j)!}{(R-j-m)!R^m} \Delta_a$, we also have that

$$[1 - p_{b-1}^{\text{up}}]^{-1} = \frac{R^m (R-j-m)!}{(R-j)!} \sum_{i=0}^m \frac{(R-j-m+i)!}{(R-j-m)!R^i}$$

$$\begin{aligned}
&= \sum_{i=0}^m \frac{(R-j-m+i)!R^{m-i}}{(R-j)!} \\
&= \sum_{i=0}^m \frac{(R-j-i)!R^i}{(R-j)!} \\
&\leq \sum_{i=0}^m \left(\frac{R}{R-j-m} \right)^i \\
&= \frac{R-j-m}{j+m} \left[\left(1 + \frac{j+m}{R-j-m} \right)^{m+1} - 1 \right] \\
&\leq \frac{R-j-m}{j+m} \left[e^{\left(\frac{(j+m)(m+1)}{R-j-m} \right)} - 1 \right].
\end{aligned}$$

This tells us that $1 - p_{b-1}^{\text{up}} \geq \frac{j+m}{R-j-m} \left[e^{\left(\frac{(j+m)(m+1)}{R-j-m} \right)} - 1 \right]^{-1}$, or, equivalently,

$$p_{b-1}^{\text{up}} \leq 1 - \frac{j+m}{R-j-m} \left[e^{\left(\frac{(j+m)(m+1)}{R-j-m} \right)} - 1 \right]^{-1}.$$

We likewise have that

$$\begin{aligned}
[1 - p_{b-1}^{\text{up}}]^{-1} &\leq (m+1) \left(\frac{R}{R-j-m} \right)^m \\
&\leq (m+1) e^{\frac{(j+m)m}{R-j-m}} \\
&\leq (m+1) e^{\frac{(j+m)(m+1)}{R-j-m}}.
\end{aligned}$$

Combining these two results, we obtain that

$$p_{b-1}^{\text{up}} \leq 1 - \max \left(\frac{1}{m+1}, \frac{j+m}{R} \right) e^{-\frac{(j+m)(m+1)}{R-j-m}}.$$

Continuing on for the lower bound, we find that

$$\begin{aligned}
[1 - p_{b-1}^{\text{up}}]^{-1} &\geq \sum_{i=0}^m \left(\frac{R}{R-j} \right)^i \\
&= \frac{R-j}{j} \left[\left(1 + \frac{j}{R-j} \right)^{m+1} - 1 \right],
\end{aligned}$$

so that,

$$p_{b-1}^{\text{up}} \geq 1 - \frac{j}{R-j} \left[\left(1 + \frac{j}{R-j} \right)^{m+1} - 1 \right].$$

I.3. Proof of Claim 27: Bound on p_e , an M/M/R hitting probability.

We now show the following claim.

CLAIM 27. *Let p_e be the probability that, in an M/M/R queue started with $(R+H)$ jobs in the system, the number of jobs reaches $(R+H+1)$ before it reaches $(R-m)$. Then,*

$$\frac{1}{1-p_e} \leq H + \sqrt{R} e^{\frac{m(m+1)}{R-m}} \leq \sqrt{R} \left(1 + e^{\frac{m(m+1)}{R-m}} \right).$$

We prove this claim by analyzing particular state transitions. Let $T_{i \rightarrow j}$ be the first passage time from state i to state j , and when evaluating the event $[T_{i \rightarrow a} < T_{i \rightarrow b}]$, consider the first passage times to be evaluated on the same sample path. Then $p_e = \Pr(T_{(R+H) \rightarrow (R-m)} < T_{(R+H) \rightarrow (R+H+1)})$. By conditioning, we have that

$$\begin{aligned} 1 - p_e &= \Pr(T_{(R+H) \rightarrow (R-m)} < T_{(R+H) \rightarrow (R+H+1)}) \\ &= \Pr(T_{(R+H) \rightarrow R} < T_{(R+H) \rightarrow (R+H+1)}) [\Pr(T_{R \rightarrow (R-m)} < T_{R \rightarrow (R+H+1)})]. \end{aligned}$$

From here, note that

$$\begin{aligned} &\Pr(T_{R \rightarrow (R-m)} < T_{R \rightarrow (R+H+1)}) \\ &= \Pr(T_{R \rightarrow (R-m)} < T_{R \rightarrow (R+1)}) + \Pr(T_{(R+1) \rightarrow (R-m)} < T_{(R+1) \rightarrow (R+H+1)}) \\ &= \Pr(T_{R \rightarrow (R-m)} < T_{R \rightarrow (R+1)}) + \Pr(T_{(R+1) \rightarrow R} < T_{(R+1) \rightarrow (R+H+1)}) \Pr(T_{R \rightarrow (R-m)} < T_{R \rightarrow (R+H+1)}). \end{aligned}$$

Note also that, from basic results on simple random walks, we have that $\Pr(T_{(R+H) \rightarrow R} < T_{(R+H) \rightarrow (R+H+1)}) = \frac{1}{H+1}$ and $\Pr(T_{(R+1) \rightarrow R} < T_{(R+1) \rightarrow (R+H+1)}) = \frac{H}{H+1}$. Let $p_d \triangleq \Pr(T_{R \rightarrow (R-m)} < T_{R \rightarrow (R+1)})$. Then we have

$$\begin{aligned} 1 - p_e &= \frac{1}{H+1} \left[p_d + (1 - p_d) \frac{H}{H+1} (p_d + \dots) \right] \\ &= \frac{p_d}{H+1} \frac{1}{1 - (1 - p_d) \frac{H}{H+1}} \\ &= \frac{p_d}{H+1 - H + p_d H} \\ &= \frac{p_d}{1 + p_d H}. \end{aligned}$$

This tells us that

$$\frac{1}{1 - p_e} = H + \frac{1}{p_d}.$$

I.3.1. Analyzing p_d . Since the $M/M/R$ queue acts equivalently to the $M/M/\infty$ queue between the states R and $(R - m)$, to analyze the probability $p_d \triangleq \Pr(T_{R \rightarrow (R-m)} < T_{R \rightarrow (R+1)})$, it suffices to analyze the appropriate hitting probability in the $M/M/\infty$. Applying the result from Section I.2, noting that the probability in question is $1 - p_{b-1}^{\text{up}}$ in the case where $j = 0$, we find that

$$p_d \geq \max\left(\frac{1}{m+1}, \frac{m}{R}\right) e^{-\frac{m(m+1)}{R-m}} \geq \frac{1}{\sqrt{R}} e^{-\frac{m(m+1)}{R-m}}.$$

Final bound on $(1 - p_e)^{-1}$. Putting these together and noting that $H \triangleq \min(\sqrt{R}, k - R) \leq \sqrt{R}$, we find

$$\begin{aligned} \frac{1}{1 - p_e} &= H + \frac{1}{p_d} \\ &\leq \sqrt{R} + \sqrt{R} e^{\frac{m(m+1)}{R-m}} \\ &= \sqrt{R} \left(1 + e^{\frac{m(m+1)}{R-m}}\right). \end{aligned}$$

I.4. Proof of Claim 26.

We now prove Claim 26; see Section I.1.1 for details.

First observations. To begin, we first recall that p_f is the probability that a given excursion is both short (that is, lasts for less than a setup time) and does not reach the appropriate threshold. From a union bound, it follows that $p_f \geq 1 - (p_{\text{th}} + p_{\text{long}}^{(j)})$, where $p_{\text{long}}^{(j)}$ is the probability that the excursion lasts for (at least) a setup time. We also note that the “retrial” probability $p_{\text{re}}^{(j)} = 1 - \gamma p_{\text{start}}^{(j)} \geq 1 - p_{\text{start}}^{(j)}$. This follows from a straightforward symmetry argument: since the associated Markov chain is biased towards *gaining* jobs, it follows that the probability that the number of jobs increases by one before it decreases by m is higher than the probability the number of jobs decreases by one before it increases by m . Putting these together, we find that

$$\begin{aligned} p_{\text{th}} p_{\text{start}}^{(j)} \frac{1}{1 - p_f p_{\text{re}}^{(j)}} &\geq \frac{p_{\text{th}} p_{\text{start}}^{(j)}}{1 - (1 - p_{\text{long}}^{(j)} - p_{\text{th}})(1 - p_{\text{start}}^{(j)})} \\ &= \frac{p_{\text{th}} p_{\text{start}}^{(j)}}{1 - (1 - p_{\text{start}}^{(j)}) + (p_{\text{long}}^{(j)} + p_{\text{th}})(1 - p_{\text{start}}^{(j)})} \\ &\geq \frac{p_{\text{th}} p_{\text{start}}^{(j)}}{p_{\text{start}}^{(j)} + p_{\text{long}}^{(j)} + p_{\text{th}}}. \end{aligned}$$

Bounding the probabilities. From here, we note that, in Claim 6, we showed that $p_{\text{th}} \geq 0.99 \frac{j}{R}$ for any $j \geq \sqrt{R}$. Likewise, applying the results from Section I.2, we find that $p_{\text{start}}^{(j)} \geq \frac{j}{R}$. To bound $p_{\text{long}}^{(j)}$, we use the arguments from Section G.6, which show that

$$\begin{aligned} p_{\text{long}}^{(j)} &\leq 1 - \left(1 - \frac{j}{R}\right) \left(1 - \frac{b_1}{\sqrt{\mu\beta R}}\right) \\ &\leq \frac{j}{R} + \frac{b_1}{\sqrt{\mu\beta R}} \\ &\leq \frac{j}{R} \left(1 + \frac{b_1}{\sqrt{\mu\beta}}\right). \end{aligned}$$

Completion of the proof of Claim 26. Combining these, we find that

$$\frac{p_{\text{th}} p_{\text{start}}^{(j)}}{1 - p_f p_{\text{re}}^{(j)}} \geq \frac{0.99 \frac{j}{R} \cdot \frac{j}{R}}{\frac{j}{R} + 0.99 \frac{j}{R} + \frac{j}{R} \left(1 + \frac{b_1}{\sqrt{\mu\beta}}\right)} = \frac{j}{R} \left(\frac{0.99}{2.99 + \frac{b_1}{\sqrt{\mu\beta}}} \right) \geq \frac{j}{R} (0.32).$$