

Performance of the Gittins policy in the $G/G/1$ and $G/G/k$, with and without setup times

Yige Hong^{a,*}, Ziv Scully^b

^a Carnegie Mellon University, Computer Science Department, 5000 Forbes Ave, Pittsburgh, PA, 15213, USA

^b Cornell University, School of Operations Research and Information Engineering, 206 Rhodes Hall, Ithaca, NY, 14853-3801, USA

ARTICLE INFO

Keywords:

Queueing theory
 $G/G/k$
 Setup times
 Gittins policy
 SRPT
 Multiserver jobs

ABSTRACT

How should we schedule jobs to minimize mean queue length? In the preemptive $M/G/1$ queue, we know the optimal policy is the Gittins policy, which uses any available information about jobs' remaining service times to dynamically prioritize jobs. For models more complex than the $M/G/1$, optimal scheduling is generally intractable. This leads us to ask: beyond the $M/G/1$, does Gittins still perform well?

Recent results show Gittins performs well in the $M/G/k$, meaning that its additive suboptimality gap is bounded by an expression which is negligible in heavy traffic. But allowing multiple servers is just one way to extend the $M/G/1$, and most other extensions remain open. Does Gittins still perform well with non-Poisson arrival processes? Or if servers require setup times when transitioning from idle to busy?

In this paper, we give the first analysis of the Gittins policy that can handle any combination of (a) multiple servers, (b) non-Poisson arrivals, and (c) setup times. Our results thus cover the $G/G/1$ and $G/G/k$, with and without setup times, bounding Gittins's suboptimality gap in each case. Each of (a), (b), and (c) adds a term to our bound, but all the terms are negligible in heavy traffic, thus implying Gittins's heavy-traffic optimality in all the systems we consider. Another consequence of our results is that Gittins is optimal in the $M/G/1$ with setup times at all loads.

1. Introduction

We consider the classic problem of preemptively scheduling jobs in a queue to minimize mean number-in-system, or equivalently mean response time (a.k.a. sojourn time). Even in single-server queueing models, this can be a nontrivial problem whose answer depends on the information available to the scheduler. The simplest case is when the scheduler knows each job's size (a.k.a. service time), for which the optimal policy is Shortest Remaining Processing Time (SRPT) [1]: always serve the job of least remaining work.

In the more realistic case of scheduling with unknown or partially known job sizes, the optimal policy is only known for the $M/G/1$. It is called the *Gittins* policy (a.k.a. Gittins index policy) [2–5]. Based on whatever service time information is available for each job, Gittins assigns each job a scalar *rank* (i.e. priority), then serves the job of least rank. For example, SRPT is the special case of Gittins where job sizes are known exactly, and a job's rank is its remaining work. More generally, a job's rank is, roughly speaking, an estimate of its remaining work based on whatever information is available.

The Gittins policy is known to be optimal in the $M/G/1$ [4,5]. But plenty of systems and models have more complex features, including:

* Corresponding author.

E-mail addresses: yigeh@cs.cmu.edu, yigeh@andrew.cmu.edu (Y. Hong), zivscully@cornell.edu (Z. Scully).

<https://doi.org/10.1016/j.peva.2023.102377>

Available online 29 October 2023

0166-5316/© 2023 The Author(s).

Published by Elsevier B.V. This is an open access article under the CC BY license

(<http://creativecommons.org/licenses/by/4.0/>).

- (a) *Multiple servers*, such as the M/G/k.
- (b) *Non-Poisson arrival processes*, such as the G/G/1 (more specifically, the GI/GI/1).
- (c) *Periods of server unavailability*, such as models with setup times.

Either (a) or (b) alone makes optimal scheduling intractable. Combining all three, as in the G/G/k with setup times (G/G/k/setup), only adds to the challenge.

With optimality out of reach, we are left to find a tractable near-optimal policy. We thus ask:

How well does Gittins perform in systems with features (a), (b), and (c) like the G/G/k/setup?

Gittins is a natural candidate because its definition naturally generalizes beyond the M/G/1, even if its optimality proof does not [4]. For instance, in a G/G/k, Gittins simply serves the k jobs of k least ranks, or all jobs if there are fewer than k .

Only feature (a) has been addressed in full generality in prior work [6–8]. Specifically, it is known that in the M/G/k, the additive suboptimality gap of Gittins (abbreviated in equations to “Gtn”) is bounded by [8]¹

$$E[N]_{M/G/k}^{Gtn} - \inf_{\text{policies } \pi} E[N]_{M/G/k}^{\pi} \leq C(k-1) \log \frac{1}{1-\rho}. \tag{1.1}$$

Let us briefly explain the notation used in (1.1):

- $E[N]_{M/G/k}^{\pi}$ is the mean number-in-system under policy π in M/G/k.
- k is the number of servers.
- $\rho \in [0, 1)$ is the *load* (a.k.a. utilization), namely the average fraction of servers that are busy.
- $C \approx 3.775$ is a constant.

A notable feature of (1.1) is that under mild conditions [6], the right-hand side is dominated by $\inf_{\pi} E[N]_{M/G/k}^{\pi}$, the performance of the optimal policy, in the *heavy-traffic limit*, meaning as $\rho \rightarrow 1$. That is, as the M/G/k gets busier and busier, the difference between Gittins’s performance and that of the optimal policy becomes negligible. Gittins is thus considered *heavy-traffic optimal* in the M/G/k.

The above progress on analyzing Gittins in the multiserver M/G/k is certainly promising for handling (a). But, as we explain in more detail in Section 1.2, key steps of the prior M/G/k analysis rely on Poisson arrivals and uninterrupted server availability, so they cannot handle (b) and (c).

1.1. Results: Performance bounds and heavy-traffic optimality

We give the first analysis of the Gittins policy for systems with any combination of (a) multiple servers, (b) G/G arrivals, and (c) setup times. We frame our results in terms of the G/G/k with setup times (G/G/k/setup), which can have all three features. But they also apply to systems with a subset of the features, such as the G/G/1 or M/G/k/setup, because we can view them as special cases of the G/G/k/setup.

Our main results, presented in full in Section 4, can be summarized as

$$E[N]_{G/G/k/setup}^{Gtn} - \inf_{\text{policies } \pi} E[N]_{G/G/k/setup}^{\pi} \leq \ell_{(a)} + \ell_{(b)} + \ell_{(a)\&(c)}, \tag{1.2}$$

where each term on the right-hand side is a “suboptimality loss” that is non-zero when the model has the features in the subscript. For example, the M/G/k/setup has only features (a) and (c), so $\ell_{(b)} = 0$. A particularly notable case is the M/G/1/setup, which has only feature (c), so all three loss terms in (1.2) are zero. Indeed, we show that Gittins is optimal among non-idling policies in the M/G/1/setup, a previously unknown result.

Our result generalizes prior work on Gittins in the M/G/k in the sense that $\ell_{(a)}$ turns out to be the right-hand side of (1.1). Remarkably, the other loss terms, $\ell_{(b)}$ and $\ell_{(a)\&(c)}$, are uniformly bounded at all loads. This implies that, under mild conditions, Gittins is heavy-traffic optimal in the G/G/k/setup.

Stability of the G/G/k/setup under complex scheduling policies

Perhaps a more basic question than optimizing mean response time is stability: under which scheduling policies is the G/G/k/setup stable? This is an open question which is outside the scope of our work. As such, our results hold under an assumption on the stability region of the G/G/k/setup (Assumption 3.1). We believe the assumption always holds, giving a partial proof sketch in Appendix D.

¹ Throughout this paper, \log is the natural logarithm.

Beyond the G/G/k/setup

The techniques underlying our results are very general, applying even beyond the G/G/k/setup. In Section 10, we sketch how our results could be extended to other systems.

- Building on the theme of multiple servers, we consider systems with *multiserver jobs*, which must be simultaneously served by multiple servers [9]. Due to the prevalence of multiserver jobs in cloud computing, these models have received lots of recent attention [7,10–15].
- Building on the theme of non-Poisson arrivals, we consider *batch arrivals* of jobs [16].
- Building on the theme of setup times, we consider *generalized vacations*, which model a variety of scenarios where servers are temporarily unavailable [17,18].

1.2. *Main obstacles and key ideas*

While there is a substantial literature on scheduling in the M/G/1 [19, Part VII], much less is known as soon as we introduce features (a), (b), and (c). Any two of these, let alone all three, yields a system where optimal scheduling has never been studied. This is perhaps unsurprising in light of the fact analyzing these systems under First-Come First-Served (FCFS) is already very difficult. See Li and Goldberg [20] and references therein for a review of the G/G/k, and similarly for Williams et al. [21] for the M/G/k/setup. Even in the G/G/1, we only know the optimal scheduling policy for known job sizes, when it is SRPT [1].

Fortunately, recent advances analyzing Gittins in the M/G/k [6–8] give us hope in the form of a new avenue for analyzing performance. Scully et al. [6] introduce a new queueing identity, now known as *WINE* [8] (Section 6),² which relates the number of jobs N in the system to, roughly speaking, the amount of *work* in the system. This is helpful because bounding the amount of work in an M/G/k, which WINE turns into a bound on $E[N]$, turns out to be significantly easier than directly bounding $E[N]$.

WINE holds in any queueing system, including the G/G/k/setup, so we can and do use the same overall strategy of bounding work, then using WINE to turn the work bound into an $E[N]$ bound. However, there are significant obstacles to carrying out this strategy in the G/G/k/setup.

Non-Poisson arrivals

The first step of our strategy is to analyze the amount of work in the system. The approach that prior work takes to analyze the M/G/k is to use a *work decomposition law*. This is a result which, in its most general form, relates the amount of work in a generic system with M/G arrivals to the amount of work in a “resource-pooled” M/G/1 experiencing the same arrivals. The prior M/G/k bound in (1.1) comes from the fact that the M/G/k and resource-pooled M/G/1 turn out to have similar amounts of work. We would like to take a similar approach with the G/G/k/setup. Unfortunately, the combination of G/G arrivals and multiple servers rules out using existing work decomposition laws (Section 2.3).

To overcome this, we prove a *new work decomposition law for G/G arrivals* (Section 7). We view this as the main technical contribution that makes our results possible. Indeed, by combining WINE and our new work decomposition law, the $\ell_{(a)}$ and $\ell_{(b)}$ terms of (1.2) follows relatively easily. But the $\ell_{(c)}$ term and heavy-traffic analysis present additional obstacles, as discussed below.

Setup times

One of the key observations behind the prior M/G/k analysis is that whenever there are k jobs in the system, all servers are occupied. This implies that in terms of work, an M/G/k never falls too far behind an M/G/1, where the M/G/1 experiences the same arrivals and has the same total service capacity. But in an M/G/k/setup or G/G/k/setup, there is no analogous limit to how far behind an M/G/1/setup or G/G/k/setup system can be, because there is no limit on the number of jobs that might arrive during a setup time.

To overcome this, we perform a novel analysis of setup times to bound the number of arrivals during one setup time *in expectation*. This analysis is the basis of the $\ell_{(a)\&(c)}$ term of (1.2).

Heavy traffic analysis

The above ideas are enough to prove the bound in (1.2). But the question remains: is the right-hand side of (1.2) small or large relative to $\inf_{\pi} E[N]_{G/G/k/setup}^{\pi}$, the performance of the optimal policy? If the latter dominates the former in the $\rho \rightarrow 1$ limit, then Gittins is optimal in heavy traffic. The right-hand side grows as $O(\log \frac{1}{1-\rho})$, so the main challenge is to give a lower bound on the performance of the optimal policy. In prior work on the M/G/k, one can use SRPT in a resource-pooled M/G/1 as a lower bound on the optimal policy, which is helpful because the SRPT has been analyzed in heavy traffic [26]. We would like to use the same approach with the G/G/1 as the lower bound, but SRPT has never been analyzed in the heavy-traffic G/G/1.

To overcome this, we give the *first heavy-traffic analysis of SRPT in the G/G/1*. This provides a lower bound on $\inf_{\pi} E[N]_{G/G/k/setup}^{\pi}$, which turns out to be enough for our purposes. The key ingredient of our heavy-traffic analysis is, once again, our new work decomposition law for G/G arrivals, underscoring its importance as our key technical contribution.

² Scully [8, Section 2.2.3] notes that WINE builds upon several similar identities that precede it [22–25].

1.3. Outline

The rest of the paper is organized as follows:

- Section 2 reviews related work.
- Section 3 describes our $G/G/k$ /setup model, and in particular details of the setup times.
- Section 4 presents our main results on Gittins: suboptimality gap bounds (Theorems 4.1 and 4.2) and heavy-traffic optimality (Theorem 4.3).
- Section 5 gives a high-level overview of how we prove our main results.
- Section 6 reviews necessary background on Gittins and WINE.
- Section 7 proves a *new work decomposition law for systems with G/G arrivals*. This is the key technical contribution that underlies all of our other results.
- Section 8 proves the suboptimality gap bounds (Theorems 4.1 and 4.2).
- Section 9 proves heavy-traffic optimality (Theorem 4.3). The key step involves giving *first heavy-traffic analysis of SRPT in the $G/G/1$* , a result of independent interest.

2. Related work

2.1. Optimal scheduling in queues

Gittins in single-server systems

The Gittins policy was originally conceived to solve the Markovian multi-armed bandit problem [4,27], but it was soon adapted to also solve the problem of scheduling in an $M/G/1$ to minimize mean number of jobs and similar metrics. See Scully and Harchol-Balter [5] and the references therein for a review of Gittins in the $M/G/1$. However, aside from some particular cases [1,28], the degree to which Gittins performs well in the $G/G/1$ or $G/G/1$ /setup was previously unknown.

The “SOAP” technique of Scully et al. [29] can be used to analyze the performance of the Gittins policy in the $M/G/1$. However, while SOAP is convenient for analyzing any fixed size distribution (e.g. numerically), using it to prove theorems that hold for all size distributions is cumbersome [30, Section 1.1]. Moreover, SOAP is limited to the $M/G/1$ and, thanks to an extension by van Vreumingen [31], the $M/G/1$ /setup. Analyzing Gittins with G/G arrivals or multiple servers seems to be beyond SOAP [32, Appendix B].

Gittins in multi-server systems

Gittins is known to be suboptimal with multiple servers [4], but researchers have studied the extent to which the suboptimality gap is large or small. The earliest results of this type analyzed an $M/M/k$ with Bernoulli feedback [22] and *nonpreemptive* $M/G/k$ with Bernoulli feedback [23]. These results proved (in the latter case, under an additional assumption) constant suboptimality gaps for Gittins in these systems. But both models are somewhat restrictive, excluding, for instance, heavy-tailed job size distributions that are common in computer systems [33–37]. More recent work, which we discussed in Section 1, overcomes these limitations to bound the performance of Gittins in the $M/G/k$ for general job sizes, including heavy-tailed sizes [7,8,38]. However, all of the above work assumes M/G arrivals with no server unavailability.

2.2. Setup times

Multiserver models

A significant line of previous work has studied the $M/M/k$ /setup with exponential setup times and FCFS scheduling [39–44]. Among those works, Gandhi and Harchol-Balter [40] and Gandhi and Harchol-Balter [42] also demonstrate that their results generalize to $M/G/k$ /setup with exponential setup times via simulation or analyzing special examples. Recently, Williams et al. [21] go beyond exponential setup times, studying $M/M/k$ /setup with deterministic setup times and FCFS scheduling. However, none of these prior works apply to general setup times, non-Poisson arrivals, or scheduling policies beyond FCFS.

We note that Glazebrook [23], who studies the nonpreemptive $M/G/k$ with Bernoulli feedback, actually studies a more general model that allows for certain types of server unavailability, such as server breakdowns. However, setup times are not covered by Glazebrook [23]. It is likely that more general future work could simultaneously cover setup times, server breakdowns, and other types of server unavailability. See Section 10.3 for additional discussion.

Single-server models

Compared with multiserver models, single-server models with setup times are better understood [17,18,45–53]. See Doshi [52] for a survey of the work before 1986 and [53] for a more recent survey. These works consider various arrival and service processes, as well as other types of server unavailability in addition to setup times.

However, none of the above works discuss optimal scheduling in the presence of setup times. Progress was made by van Vreumingen [31], who obtains the mean response time of Gittins in the $M/G/1$ /setup as a special case of a more general analysis (Section 2.1). But the analysis does not show that other policies might outperform Gittins, nor does it apply to the $G/G/1$ /setup.

2.3. Decomposition laws in queues

There is a long tradition of proving work decomposition laws for queueing systems [6,8,17,18,22,23,51,54,55]. Most of these laws take the form

$$E[\text{work in complex system with M/G arrivals}] = E[\text{work in M/G/1}] + E[\text{extra work due to complexity}].$$

For example, if the complex system is an M/G/1/setup, the extra work from complexity depends on the setup time distribution. Most work decomposition laws are actually even stronger, holding *distributionally* instead of just in expectation.

We need a work decomposition law where the complexity includes, among other factors, having multiple servers. Such a result for M/G arrivals is relatively recent [6,8], and no such result exists for G/G arrivals. While there are work decomposition laws for G/G arrivals in the literature [46,51,55], to the best of our knowledge, they apply only to single-server models with vacations. To the best of our knowledge, we prove the first work decomposition law for G/G arrivals that holds for multiserver systems like the G/G/k.

3. Model

3.1. Core queueing models: G/G/k, G/G/1, M/G/k, and M/G/1

We consider a G/G/k queueing model with a single central queue and k identical servers. The system experiences *G/G arrivals*: jobs arrive one-by-one with i.i.d. *interarrival times*, and each job has an i.i.d. *size*, or service requirement. Interarrival times and job sizes are independent of each other. We denote a generic random interarrival time by A and a generic random job size by S .³

At any moment of time, a job in the system can be served by one server. Any jobs not in service wait in the queue. Once a job's service is finished, it departs. We follow the convention that each of the k servers has service rate $1/k$. A job of size S thus requires kS time in service to finish. This convention gives all systems we study the same maximum total service rate, namely $k \cdot 1/k = 1$, and thereby the same stability condition.

The name "G/G/k" denotes the fact that the system has G/G arrivals and k servers. When A is exponentially distributed, we write M/G in place of G/G, as in "M/G/k".

Scheduling policies

The scheduling policy decides, at every moment in time, which job is in service at which server. We consider a preempt-resume model where preemption occurs without delay or loss of work.

The scheduling objective is minimizing the *mean number of jobs* in the system. We denote the mean number of jobs in system SYS under scheduling policy π by $E[N]_{\text{SYS}}^{\pi}$, omitting the "SYS" and/or " π " if there is no ambiguity. By Little's law [56], minimizing $E[N]$ is equivalent to minimizing *mean response time*, the average amount of time a job spends between its arrival and departure.

We use a flexible model of how much the scheduler knows about each job's size (Section 3.3). We restrict attention to *non-idling* policies, which are those that never unnecessarily leave servers idle. Nevertheless, our results have implications even for idling policies (Section 4.1).

As a consequence of frequent preemption, the server can share one server between multiple jobs. We formalize this in [Appendix B.1](#), but our presentation does not depend on the formal details.

Load and stability

We write $\lambda = 1/E[A]$ for the average arrival rate and $\rho = \lambda E[S]$ for the system's *load*, or utilization. One can think of ρ as the average fraction of servers that are busy. It is clear that $\rho < 1$ is a necessary condition for stability (unless both A and S are deterministic), so we assume this throughout.

Some of our results are stated for the *heavy-traffic limit*. For our purposes, this limit, denoted $\rho \rightarrow 1$, refers to a limit as the job size distribution S remains constant, and the interarrival time distribution A is scaled uniformly down with its mean approaching the mean job size. That is, the system with load ρ has interarrival time $A_{\rho} = A_1/\rho$ for some fixed distribution A_1 , where $E[A_1] = E[S]$.

It seems intuitive that $\rho < 1$ should be sufficient for stability under non-idling policies, and it is in the G/G/1 [57]. But to the best of knowledge, there are no results characterizing stability of the G/G/k under complex scheduling policies. Even under FCFS, proving stability of the G/G/k is not simple, because the system can be stable even when it never empties [58–61]. Setup times further complicate the matter.

We consider the question of proving stability of the G/G/k/setup under arbitrary non-idling scheduling policies to be outside the scope of this paper, so we simply assume stability when $\rho < 1$. We expect this is indeed the case, giving the initial steps of a proof sketch in [Appendix D](#).

Assumption 3.1. For all $\rho < 1$, the G/G/k/setup is stable under all non-idling scheduling policies, including the Gittins policy.

³ This arrival process is often referred to more specifically as GI/GI arrivals, with the "I" emphasizing the independence assumption. Under this convention, G/G arrivals include even more general stationary arrival processes where independence does not hold. In this work, we focus only on the independent case, so we write simply "G/G" instead of "GI/GI" for brevity.

Additional assumption on interarrival times

Our results for G/G arrivals depend on “how non-Poisson” arrival times are, which we quantify using the following assumption.

Assumption 3.2. There exist $A_{\min}, A_{\max} \in \mathbb{R}_{\geq 0}$ such that $\mathbf{E}[A - a \mid A > a] \in [A_{\min}, A_{\max}]$ for all $a \geq 0$. That is, letting the *interarrival age* A_{age} be the time since the last arrival and *residual interarrival time* A_{res} be the amount of time until the next arrival, we have

$$\mathbf{E}[A_{\text{res}} \mid A_{\text{age}}] \in [A_{\min}, A_{\max}] \quad \text{with probability 1.}$$

One may always use $A_{\min} = \inf_{a \geq 0} \mathbf{E}[A - a \mid A > a]$ and $A_{\max} = \sup_{a \geq 0} \mathbf{E}[A - a \mid A > a]$, so this assumption boils down to the latter being finite.

Our results use [Assumption 3.2](#) via the quantity $\lambda(A_{\max} - A_{\min})$, which we can think of as measuring “how non-Poisson” arrival times are. In the Poisson case, one may use $A_{\min} = A_{\max} = 1/\lambda$, so $\lambda(A_{\max} - A_{\min}) = 0$.

Many interarrival distributions A satisfy [Assumption 3.2](#), such as all phase-type distributions. One can also think of [Assumption 3.2](#) as a relaxation of the well-known *New Better than Used in Expectation* (NBUE) property, which is the special case where $A_{\max} = \mathbf{E}[A]$. The main distributions ruled out by [Assumption 3.2](#) are various classes of heavy-tailed distributions, e.g. power-law tails.

3.2. Setup times

In addition to the basic G/G/k model defined above, we also consider models in which servers require *setup times* to transition from idle to busy. We denote these models with an extra “/setup”, as in G/G/k/setup. Whenever a server switches from idle to busy, it must first complete an i.i.d. amount of *setup work*, denoted U . Like work from jobs, servers complete setup work at rate $1/k$, so setup work U results in setup time kU . Setup work amounts are independent of interarrival times and job sizes.

For the purposes of stating our results and proofs in a unified manner, we consider the G/G/k without setup times to be the special case of the G/G/k/setup where $U = 0$ with probability 1.

In our model, a server can be in one of three states:

- *Setting up*, i.e. doing setup work.
- *Busy*, i.e. serving a job.
- *Idle*, i.e. neither serving a job nor doing setup work.

In the G/G/1/setup, state transitions are straightforward: the server goes from setting up to busy when it finishes its setup work, from busy to idle when no jobs remain in the system, and from idle to setting up when a job arrives to an empty system. But in the G/G/k/setup, the transitions are more complicated. This is because there are several design choices to make, and thus multiple models that might be studied. For example, if we already have one busy server, how many jobs should there be in the queue before we start setting up a second server? For concreteness, we study one particular setup time model, described below, but our work still has implications for alternative models ([Sections 4.1 and 10.3](#)).

In the G/G/k/setup, we use the following setup time model: a server transitions

- from setting up to busy when it finishes its setup work,
- from busy to idle when the system has fewer jobs than busy servers, and
- from idle to setting up when the system has fewer busy or setting up servers than jobs.

Thus, transitions to setting up are triggered by arrivals, and transitions to idle are generally triggered by departures. Servers transition “one at a time”, e.g. an arrival triggers at most server to start setting up.

Note that once a setup time begins, it is never canceled, even if the job whose arrival triggered the setup time begins service at another server. Unless another job arrives during the setup time, the server will transition from setting up to busy, then immediately back to idle. Not canceling setup times is a natural modeling choice for some systems, e.g. computer systems where cutting power during startup is undesirable. But our techniques could also be used to analyze setup times that can be canceled ([Section 10.3](#)).

3.3. What the scheduler knows about jobs’ sizes

We consider a flexible model of the scheduler’s knowledge called the *Markov-process job model* [[5,6,8](#)]. In this model, each job has a *state*, which inhabits some *job state space* \mathbb{X} , representing what the scheduler knows about that job. Each job’s state evolves as an i.i.d. absorbing continuous-time Markov-process $\{X(t)\}_{t \geq 0}$ on some state space \mathbb{X} , where $X(t)$ is the state of the job after it has received $t \geq 0$ service. That is, a job’s state evolves while it is in service but stays static while it is in the queue. There is an extra absorbing state $\top \notin \mathbb{X}$, corresponding to the job finishing, i.e. jobs exit the system when their state becomes \top .

We call $\{X(t)\}_{t \geq 0}$ the *job Markov process*. We can recover the job size from the job Markov process as

$$S = \inf\{t \geq 0 \mid X(t) = \top\}.$$

To clarify, the amount of service t in $X(t)$ is measured in *work* rather than *time*, so jobs evolve at rate $1/k$ when served in a k -server system ([Section 3.1](#)). As discussed in [Appendix B.3](#), we make some purely technical assumptions on the job Markov process (e.g. r.c.l.l.) to ensure Gittins is well defined.

We assume that the scheduler always knows the state of all jobs in the system, which we denote by (X_1, \dots, X_N) . We also assume the scheduler knows the dynamics of the job Markov process, e.g. the size distribution S . A job’s state thus encodes everything the scheduler knows about the job. For example, given a job in state x , the scheduler knows its *remaining work*, namely the amount of service the job needs to complete, is distributed as⁴

$$S(x) = (\inf\{t \geq 0 \mid X(t) = \top\} \mid X(0) = x).$$

Below are two concrete examples of the Markov-process job model. These are extremes: the first is perfect size information, and the other is zero size information beyond knowing the distribution S . For additional examples, including cases where the scheduler has partial size information, see Scully et al. [6, Section 3].

Example 3.3. The case of *known sizes* is when a job’s state is its remaining work. The state space is $\mathbb{X} = (0, \infty)$, the initial state is distributed as $X(0) \sim S$, and the absorbing state is $\top = 0$. During service, the job’s state decreases at rate 1. That is, $X(t) = (X(0) - t)^+$. In state x , the remaining work is $S(x) = x$.

Example 3.4. The case of *unknown sizes* is when a job’s state is the amount of service it has received so far. The state space is $\mathbb{X} = [0, \infty)$, the initial state is $X(0) = 0$, and the absorbing state is an isolated point \top . During service, the job’s state increases at rate 1 and has a chance to jump to \top , with the exact chance depending on the distribution of S . That is, $X(t) = t$ until the job completes, after which $X(t) = \top$. In state x , the remaining work is the conditional distribution $S(x) \sim (S - t \mid S > t)$.

3.4. The Gittins policy

The scheduling policy we focus on in this work is the *Gittins policy* (a.k.a. *Gittins index policy*). Gittins is primarily known for the fact that it minimizes $\mathbf{E}[N]$ in the M/G/1 [4,5]. In formulas, we abbreviate Gittins to “Gtn”, as in $\mathbf{E}[N]_{G/G/k/setup}^{\text{Gtn}}$.

The Gittins policy has a relatively simple form. It assigns each job a numerical priority, called a *rank*, where lower rank is better. Gittins always serves the job or jobs of least rank,⁵ and it is non-idling, serving as many jobs as the number of available servers allows. Gittins determines ranks using a *rank function*

$$\text{rank}_{\text{Gtn}} : \mathbb{X} \rightarrow \mathbb{R}_{\geq 0},$$

assigning $\text{rank}_{\text{Gtn}}(x)$ to a job in state $x \in \mathbb{X}$. A job’s rank thus depends only on its own state.

It turns out that our proofs do not directly use the definition of Gittins’s rank function. As such, we specify the Gittins rank function for the concrete job Markov processes in [Examples 3.3](#) and [3.4](#), the latter of which in particular explains the key intuition. We refer the curious reader to [Appendix B.2](#) for the general definition, though we emphasize it does not play a direct role in our proofs.

Example 3.5. In the case of *known job sizes*, it turns out that Gittins reduces to SRPT, which always serves the job of least remaining work. A job’s rank is thus its remaining work. Recalling from [Example 3.3](#) that a job’s state is its remaining work under known sizes, we simply have $\text{rank}_{\text{Gtn}}(x) = x$.

Example 3.6. In the case of *unknown job sizes*, recall from [Example 3.4](#) that a job’s state x is the amount of service it has already received. In this case, the Gittins rank function is [4]

$$\text{rank}_{\text{Gtn}}(x) = \inf_{y > x} \frac{\mathbf{E}[\min\{S, y\} - x \mid S > x]}{\mathbf{P}[S \leq y \mid S > x]}.$$

The intuition for this formula is as follows. Consider a job in state x , and suppose we start serving the job, but decide to “give up” if it reaches state y . On the right-hand side, the numerator is the expected amount of service until we either complete the job or give up, and the denominator is the probability the job completes before we give up. The right-hand side is thus a “service-per-completion” ratio, giving an expected amount of effort it would take to finish one job in expectation. A job’s rank under Gittins is the best service-per-completion ratio one can obtain by optimally choosing the state y in which to give up.

4. Main results

We now state our main results. All of our results hold under the assumptions of [Section 3](#), and in particular [Assumptions 3.1](#) and [3.2](#). As in [Section 1](#), we can view a G/G/k/setup system, or any special case thereof, by whether it has (a) multiple servers,

⁴ Abusing notation slightly, we interpret conditioning $X(0) = x$ as the usual notion of starting the job Markov process from state x . This gets around the corner case where the initial state $X(0)$ is never x .

⁵ Much literature on the Gittins policy uses the opposite convention, where higher numbers are better. These works typically call a job’s priority its *index* [2–4], which is the reciprocal of its rank [5].

(b) non-Poisson arrivals, and (c) setup times. Our bounds use the quantities

$$\begin{aligned} \ell_{(a)} &= C(k-1) \log \frac{1}{1-\rho}, \\ \ell_{(b)} &= \lambda(A_{\max} - A_{\min}), \\ \ell_{(c)} &= \mathbf{1}(\mathbf{P}[U > 0] > 0) (2(k-1) + \lambda(A_{\max} + k\mathbf{E}[U_e])), \end{aligned}$$

where $C = \frac{9}{8 \log 1.5} + 1 \approx 3.775$. The idea is that $\ell_{(a)}$ is the loss due to feature (a), as it is nonzero only for systems with $k \geq 2$ servers, and similarly for $\ell_{(b)}$ and $\ell_{(c)}$,⁶

Theorem 4.1. *The performance gap between the Gittins policy in G/G/k/setup and the optimal policy in G/G/1 is bounded by*

$$\mathbf{E}[N]_{G/G/k/setup}^{\text{Gtn}} - \inf_{\pi} \mathbf{E}[N]_{G/G/1}^{\pi} \leq \ell_{(a)} + \ell_{(b)} + \ell_{(c)}.$$

Note that although [Theorem 4.1](#) is not directly about the suboptimality gap of Gittins policy in G/G/k/setup, it still provides an upper bound on the suboptimality gap, because the optimal performance of G/G/1 is a lower bound to G/G/k/setup. This is because servers in G/G/k/setup have speed $1/k$ ([Section 3.1](#)), so the G/G/1 can mimic any policy in the G/G/k/setup through processor sharing and idling.

With that said, in the special case of the non-idling G/G/1/setup, we can prove a stronger result that drops the $\ell_{(c)}$ term by comparing to a G/G/1/setup instead of a G/G/1.

Theorem 4.2. *In the G/G/1/setup, the performance gap between the Gittins policy and the optimal non-idling policy is bounded by*

$$\mathbf{E}[N]_{G/G/1/setup}^{\text{Gtn}} - \inf_{\pi} \mathbf{E}[N]_{G/G/1/setup}^{\pi} \leq \ell_{(b)}.$$

In particular, in the M/G/1/setup, the Gittins policy minimizes $\mathbf{E}[N]$ among non-idling policies.

The suboptimality gap in [Theorem 4.1](#) is constant when $k = 1$ and $O(\log \frac{1}{1-\rho})$ when $k \geq 2$. In both cases, the gap grows more slowly in the $\rho \rightarrow 1$ limit than $\mathbf{E}[N]_{G/G/1}^{\pi}$, implying heavy-traffic optimality.

Theorem 4.3. *In the G/G/k/setup, if either $k = 1$ or $\mathbf{E}[S^2(\log S)^+] < \infty$, and if either S or A is not deterministic, the Gittins policy is heavy-traffic optimal. Specifically, $\lim_{\rho \rightarrow 1} \mathbf{E}[N]_{G/G/k/setup}^{\text{Gtn}} / \inf_{\pi} \mathbf{E}[N]_{G/G/1}^{\pi} = 1$.*

We prove this result in [Section 9](#). The main obstacle is showing a lower bound on $\mathbf{E}[N]_{G/G/1}^{\pi}$. We use SRPT as a lower bound, so the first step of the proof is to analyze SRPT in the heavy-traffic G/G/1 ([Theorem 9.1](#)). We find its performance is within a constant factor of SRPT in the heavy-traffic M/G/1.

4.1. Remarks on main results

Alternative setup time models

Because [Theorems 4.1](#) and [4.3](#) compare Gittins in the G/G/k/setup to the optimal policy in a G/G/1, it also effectively compares Gittins under our setup time model to Gittins in the G/G/k with essentially any other setup time model. This is because the G/G/1 serves as a lower bound for alternative setup time models, not just our specific G/G/k/setup. The takeaway is that changing the setup time model does not significantly impact performance in heavy traffic, which makes intuitive sense: servers seldom set up if they are usually busy.

Idling policies

We say in [Section 3.1](#) that we only consider non-idling policies, but [Theorem 4.1](#) still compares Gittins to idling policies. This is because the optimal policy in a G/G/1 is clearly non-idling, and, by the discussion above, it gives a lower bound on any policy, idling or non-idling, in the G/G/k/setup.

Why, then, does [Theorem 4.2](#) only compare to the optimal non-idling policy? This is because in the G/G/1/setup, idling the server can change when setup times occur. By idling with one job in the queue, one can effectively control when setup times occur by choosing when to start the job, without waiting for an arrival. Hypothetically, this could improve performance in the G/G/1/setup. That is, idling effectively allows a policy to use an alternative setup time model, so we rule it out.

⁶ The reason [\(1.2\)](#) has an $\ell_{(a)\&(c)}$ term instead of an $\ell_{(c)}$ term is because it summarizes both [Theorems 4.1](#) and [4.2](#).

Opportunities for a tighter bound

The bound shown in [Theorem 4.1](#) represents a trade-off between proving a tight bound and stating the result simply. We prioritized making the statement as simple as possible while ensuring the $\ell_{(a)}$ term matches the bound for Gittins in the M/G/k from prior work [\[8\]](#). But there are at least two clear avenues for tightening our bounds.

First, there are other bounds on Gittins in the M/G/k [\[6,7\]](#), which can be better than $\ell_{(a)}$ in some cases. We believe that one may take $\ell_{(a)}$ to be the minimum of these bounds, but doing so would complicate the result and proof without substantially changing the main takeaway.

Second, our bound is loose in light traffic. We should have $\mathbf{E}[N] \rightarrow 0$ as $\rho \rightarrow 0$, but our $\ell_{(b)}$ and $\ell_{(c)}$ terms remain nonzero at all loads. One can sharpen our analysis for the special case of the M/G/k/setup to obtain a suboptimality gap that becomes zero in the $\rho \rightarrow 0$ limit. But we doubt even this improved bound is very tight at low loads, so we omit the extra casework.

5. Proof overview

In this section, we give an overview of the proofs of our main results: bounds on Gittins’s suboptimality gap ([Theorems 4.1 and 4.2](#)) and Gittins’s heavy-traffic optimality ([Theorem 4.3](#)). At a high level, our proofs work by combining two queuing identities: WINE, which is from prior work; and a novel *work decomposition law*, which is built on similar decomposition results from prior work ([Section 2.3](#)).

The first tool, WINE ([Lemma 6.2](#)), expresses the mean number-in-system in terms of *mean r-work* $\mathbf{E}[W_r]$ [\[5,6,8\]](#):

$$\mathbf{E}[N] = \int_0^\infty \frac{\mathbf{E}[W_r]}{r^2} dr.$$

A system’s r-work W_r is the total service required to serve all jobs in the system until they all either complete or reach rank greater than r , as determined by rank_{Gtn} ([Section 3.3](#)). For example, ∞ -work is the total remaining work of all jobs, which we call *total work* or simply *work*. See [Section 6](#) for details.

The second tool, the work decomposition law ([Theorem 7.2](#)), implies bounds on $\mathbf{E}[W_r]$ under any policy, including Gittins. Combining this with WINE yields bounds on $\mathbf{E}[N]$. Our proof thus boils down to three steps:

- Proving the work decomposition law ([Section 5.1](#)).
- Using the work decomposition law to bound Gittins’s suboptimality gap ([Section 5.2](#)).
- Using the suboptimality gap bounds to show Gittins is heavy-traffic optimal ([Section 5.3](#)).

5.1. New tool: Work decomposition law for G/G arrivals

Our work decomposition law, [Theorem 7.2](#), characterizes mean r-work $\mathbf{E}[W_r]$ in the G/G/k/setup. For simplicity of presentation, below we focus on the special case of mean total work $\mathbf{E}[W]$.

[Theorem 7.2](#) implies that in the G/G/k/setup under any policy π ,

$$\mathbf{E}[W]^\pi - \mathbf{E}[W]_{G/G/1} \leq \frac{\mathbf{E}[J_{\text{idle}}W]^\pi}{1 - \rho} + \frac{\mathbf{E}[J_{\text{setup}}W]^\pi}{1 - \rho} + \rho(A_{\text{max}} - A_{\text{min}}). \tag{5.1}$$

Above, $\mathbf{E}[W]_{G/G/1}$ is the mean work in a non-idling G/G/1, which is policy-invariant; and J_{idle} and J_{setup} are the fraction of idle and setting-up servers, respectively. Flipping the sign on the $\rho(A_{\text{max}} - A_{\text{min}})$ term yields a lower bound instead of an upper bound.

The work decomposition law decomposes work $\mathbf{E}[W]^\pi$ into the policy-invariant term $\mathbf{E}[W]_{G/G/1}$, plus error terms that can depend on the policy π . Each error term characterizes the consequence of a complicating factor that G/G/k/setup has on the top of the G/G/1 system:

- (a) The first term is due to having multiple servers. It vanishes when $k = 1$, as then $J_{\text{idle}} = 0$ if $W > 0$.⁷
- (b) The second term is due to the setup time. It vanishes if servers do not need setup, as then $J_{\text{setup}} = 0$.
- (c) The third term is due to non-Poisson arrivals. It vanishes for Poisson arrivals, as then $A_{\text{max}} = A_{\text{min}}$.

How we prove the work decomposition law

The proof of work decomposition laws in prior work involves viewing W as a process in the steady state and analyzing its continuous changes and jumps. This strategy works well in M/G systems, because all times have an equal chance of seeing W jump up due to an arrival. But in G/G systems, the chance of having an arrival in the next moment depends on A_{age} , the amount of time since the previous arrival. The jumps of W are thus more complicated to analyze.

The key idea in our proof is to smooth out the non-constant jumping rate of W . Specifically, we consider the process $W - \rho A_{\text{res}}$, which only differs from W by one interarrival time. This process decreases at a constant rate of $1 - \rho$. When an arrival happens, the process jumps, but the expected change is $\mathbf{E}[S] - \rho\mathbf{E}[A] = 0$. Therefore, arrivals only have a “second-order” effect on W , which makes them easier to analyze. This idea builds upon similar smoothing approaches in recent queueing literature [\[62,63\]](#).

⁷ When generalizing [\(5.1\)](#) from total work to r-work, there are actually two terms due to having multiple servers. But both vanish when $k = 1$.

5.2. From work decomposition to suboptimality gap bounds

We focus here on proving [Theorem 4.1](#), commenting only briefly on the similar proof of [Theorem 4.2](#).

Combining our work decomposition law with WINE gives a formula for Gittins’s suboptimality gap that has the same types of error terms as [\(5.1\)](#). Each error term in the work decomposition law will result in one term in the suboptimality gap $\ell_{(a)} + \ell_{(b)} + \ell_{(c)}$ in [Theorem 4.1](#), after doing the integration and applying some additional treatments that are specific to each term.

Among the three error terms, $\ell_{(a)}$ can be derived similarly to prior work on the M/G/k [6–8], and $\ell_{(b)}$ follows from [Assumption 3.2](#). But the term corresponding to setup, $\ell_{(c)}$, requires a new analysis. We demonstrate the intuition by bounding $\mathbf{E}[J_{\text{setup}}W]$ in [\(5.1\)](#). First, we write $\mathbf{E}[J_{\text{setup}}W] = \frac{1}{k} \sum_{i=1}^k \mathbf{E}[J_{\text{setup},i}W]$, where $J_{\text{setup},i} = \mathbb{1}(\text{server } i \text{ is setting up})$. Observe that

$$\mathbf{E}[J_{\text{setup},i}W] = \mathbf{P}[J_{\text{setup},i} = 1] \mathbf{E}[W \mid J_{\text{setup},i} = 1].$$

Intuitively, $\mathbf{P}[J_{\text{setup},i} = 1]$ should be diminishing as the load gets heavy because the queue length will get longer the server i will be turned off less frequently. The second factor, $\mathbf{E}[W \mid J_{\text{setup},i} = 1]$, should be bounded because given that the server i is setting up, the work in the system should be no more than the work that arrives during the setup, plus the work that already exists when the setup happens.

For the proof of [Theorem 4.2](#), which gives a tighter bound for the single-server case, we apply WINE and work decomposition law in the same way as above. We will get an expression for $\mathbf{E}[N]_{G/G/1/\text{setup}}^{\text{Gin}}$ in terms of one G/G/1 term, and two error terms corresponding to non-Poisson arrivals and setup times. Instead of analyzing the setup term as in the proof of [Theorem 4.1](#), we make the simple observation that the setup term is the same for all non-idling policies, so it does not contribute to the suboptimality gap.

5.3. From suboptimality gap bounds to heavy-traffic optimality

[Theorem 4.1](#) provides an upper bound on the suboptimality gap of Gittins policy in G/G/k/setup. To show that the suboptimality gap is small compared with $\inf_{\pi} \mathbf{E}[N]_{G/G/k/\text{setup}}^{\pi}$ and establish heavy-traffic optimality of the Gittins policy, we need a lower bound on $\inf_{\pi} \mathbf{E}[N]_{G/G/k/\text{setup}}^{\pi}$. This lower bound can be obtained by analyzing $\mathbf{E}[N]_{G/G/1}^{\text{SRPT}}$ because SRPT gives the optimal number-in-system in G/G/1 with known job sizes [1], which is no more than the optimal number-in-system in G/G/k/setup achievable by a policy that does not know the job size.

We use WINE and work-decomposition law, in a similar way as in the proofs in the suboptimality gaps, to connect SRPT’s performance in the G/G/1 to its performance in the M/G/1. Our end result ([Theorem 9.1](#)) shows that $\mathbf{E}[N]_{G/G/1}^{\text{SRPT}}$ is a constant factor away from $\mathbf{E}[N]_{M/G/1}^{\text{SRPT}}$ as $\rho \rightarrow 1$. This lets us to use the known heavy-traffic asymptotics of SRPT in the M/G/1 [26] to lower bound $\mathbf{E}[N]_{G/G/1}^{\text{SRPT}}$ and thus show Gittins’s heavy-traffic optimality in the G/G/k/setup.

6. Background on WINE and r-work

A queueing system’s work W is the total remaining work of all jobs in the system: $W = \sum_{i=1}^N S(X_i)$, where $S(X_i)$ is the remaining work of job i ([Section 3.3](#)). We define r-work similarly: we first define the remaining r-work of a job, then define the system’s r-work to be the sum of all jobs’ remaining r-work.

Definition 6.1. Let $r \geq 0$. The remaining r-work of a job in state x , denoted $S_r(x)$, is the amount of service it needs until it either finishes or reaches a state whose rank is at least r :

$$\begin{aligned} S_r(x) &= \text{amount of service a job starting at } x \text{ needs to finish or reach rank at least } r \\ &= (\inf \{t \geq 0 \mid X(t) = \top \text{ or } \text{rank}_{\text{Gin}}(X(t)) \geq r\} \mid X(0) = x). \end{aligned}$$

A system’s r-work, denoted W_r , is the sum of the remaining r-work of all jobs in the system: $W_r = \sum_{i=1}^N S_r(X_i)$.

We now present the WINE identity. It holds for any scheduling policy that has access to only the current and past system states ([Section 3.3](#)). For concreteness, we state WINE for our specific queueing model, but it holds in essentially any system which uses the Markov-process job model.

Lemma 6.2 (WINE [[6](#), [Theorem 6.3](#)]). *In the G/G/k/setup under any scheduling policy,*

$$N = \int_0^{\infty} \frac{\mathbf{E}[W_r \mid X_1, \dots, X_N]}{r^2} dr, \quad \mathbf{E}[N] = \int_0^{\infty} \frac{\mathbf{E}[W_r]}{r^2} dr.$$

WINE, which integrates the entire system’s r-work to get the number of jobs, follows from a more basic identity, sometimes called “single-job WINE” [[8](#)], which integrates a single job’s remaining r-work.

Lemma 6.3 (Single-Job WINE [[6](#), [Lemma 6.2](#)]). *For any job state $x \in \mathbb{X}$, we have $\int_0^{\infty} \frac{\mathbf{E}[S_r(x)]}{r^2} dr = 1$.*

One subtlety about WINE is that while it applies to any scheduling policy, the definition of r-work uses Gittins’s rank function. As a general rule, this makes analyzing Gittins’s performance using WINE easier than analyzing other policies’ performance using WINE, particularly when proving upper bounds, though there are some exceptions [[8,64](#)]. Our work is no exception: we prove our main results by upper bounding Gittins’s r-work and lower-bounding the optimal policies’ r-work.

6.1. Additional definitions for reasoning about r-work

WINE reduces the problem of analyzing the steady-state mean number of jobs $E[N]$ to the problem of analyzing steady-state mean r-work $E[W_r]$. In order to analyze r-work, we need to understand the means by which the amount of r-work in the system changes over time. This section introduces the standard concepts and vocabulary used to discuss r-work [6,8,29].

The definitions in this section are parameterized by a rank $r \geq 0$, as denoted by a prefix “r-”. We often drop this prefix when the rank r is clear from context or not important to the discussion.

Relevant, irrelevant, fresh, and recycled jobs

We call a job *r-relevant* whenever its rank is less than r . Otherwise, the job is *r-irrelevant*.

Whether a job is r-relevant or r-irrelevant varies over time. Consider one job’s journey through the system. When the job arrives, it may be either, depending on its initial state $X(0)$. As the job is served, its rank can go up and down, so it may alternate between r-relevant and r-irrelevant, possibly multiple times, before eventually finishing and exiting the system.

From the above discussion, it is evident that there are two ways for the amount of r-work in a system to increase. Both are important, so we introduce terminology for discussing both.

Definition 6.4. We call an r-relevant job *r-fresh* if it has been r-relevant ever since its arrival. That is, new arrivals that are initially r-relevant are r-fresh until they either finish or become r-irrelevant.

- We write $S_r = S_r(X(0))$ for the random amount of service during which a newly arrived job is r-fresh. Arriving jobs may be r-irrelevant, so it may be that $S_r = 0$ with nonzero probability.
- We call $\rho_r = \lambda E[S_r]$ the *r-fresh load*. It is the average rate r-work is added by new arrivals.

Definition 6.5. We call an r-relevant job *r-recycled* if it was r-irrelevant at some point in the past. We refer to the moment a job switches from r-irrelevant to r-recycled as an *r-recycling*.

- We write $\lambda_{r\text{-rcy}}$ for the average rate of r-recyclings.
- We write $S_{r\text{-rcy}}$ for the random amount of r-work added by a single r-recycling.
- We call $\rho_{r\text{-rcy}} = \lambda_{r\text{-rcy}} E[S_{r\text{-rcy}}]$ the *r-recycled load*. It is the average rate r-work is added by r-recyclings.

Note that r-recyclings can only occur when an r-irrelevant job is in service. This is because a job’s rank depends only on its state, which only changes during service (Sections 3.3 and 3.4).

We assume for ease of presentation that $\lambda_{r\text{-rcy}} < \infty$ for all $r \geq 0$. This is actually not always the case, but the assumption can be straightforwardly relaxed by appealing to more advanced Palm calculus techniques. See Scully et al. [6, Appendix E] for details.

Server states

Finally, we need notations for discussing how many servers are doing r-work or setting up.

Definition 6.6.

- We write J_r for the fraction of servers that are serving r-relevant jobs.⁸ Note that $E[J_r] = \rho_r + \rho_{r\text{-rcy}}$.
- We write J_{setup} for the fraction of servers that are setting up.

7. Work decomposition law for G/G arrivals

In this section, we introduce a new formula for the expected r-work in systems with G/G arrivals, which we call the *work decomposition law*. The formula decomposes mean r-work into multiple terms, where each term is either fixed or easy to analyze. To present the formula and its proof, we need to first review a concept called *Palm expectation*. We give an intuitive review that suffices for our purposes below, referring the reader to Miyazawa [65] and Baccelli and Brémaud [66] for more formal treatments.

7.1. Palm expectation and notation

Consider the amount of work in the system as seen by arriving jobs. Intuitively speaking, to define the long-run average as seen by arrivals, we would like to start with the steady-state average, then condition on the event “an arrival is occurring”. However, the probability that an arrival occurs at any given point in time is 0, so we cannot naively apply the usual conditional expectation formula. Palm expectations solve this problem, giving a way of formally defining a notion that corresponds to conditioning on events like “an arrival is occurring”. We use two main Palm expectations in our analysis of r-work:

- The *Palm expectation at arrivals*, denoted $E_{\text{arv}}[\cdot]$.

⁸ If processor sharing is occurring, we interpret J_r as the total service rate assigned to r-relevant jobs (Appendix B.1).

- The Palm expectation at r-recyclings, denoted $\mathbf{E}_{r\text{-rcy}}[\cdot]$.

That is, $\mathbf{E}_{\text{arv}}[V]$ is the expectation of random quantity V , which may depend on the system state, sampled at the time “just before” an arrival occurs; and $\mathbf{E}_{r\text{-rcy}}[V]$ is the same, but for r-recyclings instead of arrivals.

Although we define $\mathbf{E}_{\text{arv}}[\cdot]$ and $\mathbf{E}_{r\text{-rcy}}[\cdot]$ as referring to the moment “just before” an event, it is also helpful to be able to refer to the quantity “just after” the event. To facilitate this, we use the following notation:

- Within $\mathbf{E}_{\text{arv}}[\cdot]$, we denote the remaining r-work of the arriving job by S_r .
 - By our independence assumptions (Section 3.1), S_r is independent of the system state.
- Within $\mathbf{E}_{r\text{-rcy}}[\cdot]$, we denote the remaining r-work of the job being recycled by $S_{r\text{-rcy}}$.
 - In general, $S_{r\text{-rcy}}$ is *not* independent of the system state, because recyclings are caused by events happening within the system.

A notation shortcut

In addition to the Palm expectation, we also define the notation $\mathbf{E}_{r\text{-acc}}[\cdot]$ as

$$\mathbf{E}_{r\text{-acc}}[V] = \frac{\mathbf{E}[(1 - J_r)V] + \lambda_{r\text{-rcy}}\mathbf{E}_{r\text{-rcy}}[S_{r\text{-rcy}}V]}{1 - \rho_r}, \tag{7.1}$$

where J_r is the fraction of servers that are busy with r-relevant jobs (Section 6.1). One can interpret $\mathbf{E}_{r\text{-acc}}[\cdot]$ as a type of Palm expectation. For instance, it behaves like an expectation in the sense that $\mathbf{E}_{r\text{-acc}}[v] = v$ for deterministic values v . But for our purposes, it suffices to understand $\mathbf{E}_{r\text{-acc}}[\cdot]$ as simply a notation shortcut.

Excess distributions

We introduce a piece of notation that occurs frequently in queuing and renewal theory [19,67], including in the statement of our work decomposition law below.

Definition 7.1. Given a nonnegative distribution V , we define its *excess*, denoted V_e , to be the distribution with tail⁹

$$\mathbf{P}[V_e > t] = \int_t^\infty \frac{\mathbf{P}[V > u]}{\mathbf{E}[V]} du, \quad \text{which has mean} \quad \mathbf{E}[V_e] = \frac{\mathbf{E}[V^2]}{2\mathbf{E}[V]}.$$

7.2. Statement and proof of work decomposition law

Now we are ready to present the work decomposition law for systems with G/G arrivals. We state the result for r-work W_r , but we can apply the result to total work W by taking an $r \rightarrow \infty$ limit.

Theorem 7.2 (Work Decomposition Law for G/G Arrivals). *In the G/G/k/setup under any policy π ,*

$$\mathbf{E}[W_r]^\pi = \frac{\rho_r(\mathbf{E}[(S_r)_e] - \mathbf{E}[S_r] + \mathbf{E}[A_e]) + \rho_{r\text{-rcy}}\mathbf{E}[(S_{r\text{-rcy}})_e]}{1 - \rho_r} + \mathbf{E}_{r\text{-acc}}[W_r]^\pi - \rho_r\mathbf{E}_{r\text{-acc}}[A_{\text{res}}]^\pi.$$

Proof. We drop the superscript π throughout. For each r , define

$$Z_r = \frac{1}{2}(W_r - \rho_r A_{\text{res}})^2. \tag{7.2}$$

We use Z_r as a “test function” and extract information about the G/G system by looking at how Z_r changes and applying Miyazawa’s Rate Conservation Law (RCL) [65]. We discuss why Z_r is the right choice of test function in Remark 7.4 below.

Over time, the quantity Z_r changes in the following ways: continuous change as r-work and remaining arrival time decrease over time, jump when a new job arrives, and jump at an r-recycling event. We use Z_r' to denote the continuous change of Z_r , use $\Delta_{\text{arv}}Z_r$ to denote the jumps of Z_r at arrival times, and use $\Delta_{r\text{-rcy}}Z_r$ to denote jumps of Z_r at recycling times. Analogous notations are used for W_r and A_{res} .

Miyazawa’s RCL [65] implies

$$\mathbf{E}[Z_r'] + \lambda\mathbf{E}_{\text{arv}}[\Delta_{\text{arv}}Z_r] + \lambda_{r\text{-rcy}}\mathbf{E}_{r\text{-rcy}}[\Delta_{r\text{-rcy}}Z_r] = 0. \tag{7.3}$$

This RCL is simply describing the fact that the contribution of continuous changes ($\mathbf{E}[Z_r']$) and jumps ($\mathbf{E}_{\text{arv}}[\Delta_{\text{arv}}Z_r]$ and $\lambda_{r\text{-rcy}}\mathbf{E}_{r\text{-rcy}}[\Delta_{r\text{-rcy}}Z_r]$) in Z_r cancel out in the long-run average sense.

To extract information about the G/G system, we analyze Z_r' , $\Delta_{\text{arv}}Z_r$, and $\Delta_{r\text{-rcy}}Z_r$ as below.

⁹ As a corner case, if $V = 0$ with probability 1, we let $V_e = 0$ with probability 1.

- At all times, W_r decreases at rate $-W'_r = J_r$, and A_{res} decreases at rate $-A'_{\text{res}} = 1$, so Z_r decreases at rate $-Z'_r = (J_r - \rho_r)(W_r - \rho_r A_{\text{res}})$.
- When a new arrival happens, Z_r jumps as follows. The new job contributes r -work S_r , so W_r jumps up by $\Delta_{\text{arv}} W_r = S_r$. And by definition, the new arrival happens just as A_{res} reaches 0, at which point it jumps up to a newly sampled interarrival time A , so $\Delta_{\text{arv}} A_{\text{res}} = A$. This means that when a new arrival happens, Z_r jumps by $\Delta_{\text{arv}} Z_r = (S_r - \rho_r A)W_r + \frac{1}{2}(S_r - \rho_r A)^2$.
- When an r -recycling happens, the r -work W_r increases by $\Delta_{\text{r-recy}} W_r = S_{\text{r-recy}}$, and A_{res} is unaffected, so $\Delta_{\text{r-recy}} Z_r = S_{\text{r-recy}}(W_r - \rho_r A_{\text{res}}) + \frac{1}{2}S_{\text{r-recy}}^2$.

Given the above formulas for Z'_r , $\Delta_{\text{arv}} Z_r$ and $\Delta_{\text{r-recy}} Z_r$, the terms in (7.3) can be computed one-by-one as follows. For $\mathbf{E}[Z'_r]$, we have

$$\mathbf{E}[Z'_r] = -(1 - \rho_r)\mathbf{E}[W_r] + \rho_r(1 - \rho_r)\mathbf{E}[A_e] + \mathbf{E}[(1 - J_r)W_r] - \rho_r\mathbf{E}[(1 - J_r)A_{\text{res}}],$$

where we have used the fact that by basic renewal theory, $\mathbf{E}[A_{\text{res}}] = \mathbf{E}[A_e]$. For $\lambda\mathbf{E}_{\text{arv}}[\Delta_{\text{arv}} Z_r]$, we have

$$\begin{aligned} \lambda\mathbf{E}_{\text{arv}}[\Delta_{\text{arv}} Z_r] &= \lambda\mathbf{E}_{\text{arv}}[(S_r - \rho_r A)W_r + \frac{1}{2}(S_r - \rho_r A)^2] \\ &= \lambda\mathbf{E}[S_r - \rho_r A]\mathbf{E}_{\text{arv}}[W_r] + \frac{1}{2}\lambda\mathbf{E}[(S_r - \rho_r A)^2] \\ &= \rho_r\mathbf{E}[(S_r)_e] + \rho_r^2\mathbf{E}[A_e] - \rho_r\mathbf{E}[S_r], \end{aligned} \tag{7.4}$$

where the second equality is due to the fact that the new job's r -work S_r and next interarrival time A are independent of the previous amount of r -work W_r , and the third equality uses Definition 7.1 and the fact that $\mathbf{E}[S_r] = \rho_r\mathbf{E}[A]$. Finally, for $\lambda_{\text{r-recy}}\mathbf{E}_{\text{r-recy}}[\Delta_{\text{r-recy}} Z_r]$, we have

$$\lambda_{\text{r-recy}}\mathbf{E}_{\text{r-recy}}[\Delta_{\text{r-recy}} Z_r] = \lambda_{\text{r-recy}}\mathbf{E}_{\text{r-recy}}[S_{\text{r-recy}}W_r] - \rho_r\lambda_{\text{r-recy}}\mathbf{E}_{\text{r-recy}}[S_{\text{r-recy}}A_{\text{res}}] + \rho_r\lambda_{\text{r-recy}}\mathbf{E}[(S_{\text{r-recy}})_e].$$

Combining the three terms with (7.3) completes the proof. \square

Remark 7.3. The proof of Theorem 7.2 does not depend on the details of the Gittins policy. It relies only partitioning the job states, with one part playing the role of states with rank less than r . See Scully [8, Sections 7.2 and 8.3] for an example of this with M/G arrivals.

Remark 7.4. The basic reason for multiplying A_{res} by ρ_r in the definition of Z_r is that we want to avoid having a $\mathbf{E}_{\text{arv}}[W_r]$ term in (7.3), as it is likely to be a large and intractable term. As we can see from (7.4), when computing $\lambda\mathbf{E}_{\text{arv}}[\Delta_{\text{arv}} Z_r]$, the term involving $\mathbf{E}_{\text{arv}}[W_r]$ vanishes because $\mathbf{E}[S_r - \rho_r A] = \mathbf{E}[S_r](1 - \lambda\mathbf{E}[A]) = 0$. Intuitively, ensuring $W_r - \rho_r A_{\text{res}}$ has zero change in expectation when a new job arrives prevents the Palm expectation of arrivals from appearing in the RCL equation. This trick appears throughout the literature on applying the RCL to queues [62,63].

8. Bounding Gittins's suboptimality gap

In this section, we prove the main results using WINE and the work decomposition law introduced in Sections 6 and 7. We first derive a general formula that decomposes the suboptimality gap into four terms and analyze each term based on the specific settings of each theorem. We express the formula in terms of the following quantities.

Definition 8.1. We define *residual interarrival cost* m_{res} , *recycling cost* m_{rcy} , *idleness cost* m_{idle} , and *setup cost* m_{setup} as follows:

$$\begin{aligned} m_{\text{res}} &= \int_0^\infty \frac{-\rho_r\mathbf{E}_{\text{r-acc}}[A_{\text{res}}]}{r^2} dr, & m_{\text{idle}} &= \int_0^\infty \frac{\mathbf{E}[(1 - J_r - J_{\text{setup}})W_r]}{r^2(1 - \rho_r)} dr, \\ m_{\text{rcy}} &= \int_0^\infty \frac{\lambda_{\text{r-recy}}\mathbf{E}_{\text{r-recy}}[S_{\text{r-recy}}W_r]}{r^2(1 - \rho_r)} dr, & m_{\text{setup}} &= \int_0^\infty \frac{\mathbf{E}[J_{\text{setup}}W_r]}{r^2(1 - \rho_r)} dr. \end{aligned}$$

Lemma 8.2 (Decomposition of Performance Difference). *The performance difference between the Gittins policy in G/G/k/setup and any policy π in the G/G/1 (or G/G/1/setup) can be decomposed as below:*

$$\mathbf{E}[N]_{\text{G/G/k/setup}}^{\text{Gtn}} - \mathbf{E}[N]_{\text{G/G/1}}^\pi = (m_{\text{res}}^{\text{Gtn}} - m_{\text{res}}^\pi) + m_{\text{rcy}}^{\text{Gtn}} + m_{\text{idle}}^{\text{Gtn}} + (m_{\text{setup}}^{\text{Gtn}} - m_{\text{setup}}^\pi).$$

Proof. Only the last two terms in Theorem 7.2, namely $\mathbf{E}_{\text{r-acc}}[W_r]$ and $\rho_r\mathbf{E}_{\text{r-acc}}[A_{\text{res}}]$, depend on the specific scheduling policy. After expanding the definitions of the cost terms (Definition 8.1) and $\mathbf{E}_{\text{r-acc}}[\cdot]$ (Section 7.1 and (7.1)), the result follows immediately from WINE (Lemma 6.2) and the fact that m_{rcy}^π and m_{idle}^π are nonnegative. \square

Note that we state Lemma 8.2 in terms of the Gittins policy only because our focus is the optimality of Gittins. The lemma is still true if we replace Gittins with any other policy.

Bounding Gittins's suboptimality gap thus reduces to bounding the four terms in Lemma 8.2. We address one in each of Sections 8.1–8.4, combining the bounds to prove our main results in Section 8.5.

8.1. Analysis of the residual interarrival cost

Proposition 8.3 (Residual Interarrival Cost). For any policy π ,

$$m_{\text{res}}^{\text{Gtn}} - m_{\text{res}}^{\pi} \leq \lambda(A_{\text{max}} - A_{\text{min}}).$$

Proof. Observe that for deterministic v , we have $\mathbf{E}_{\text{r-acc}}[v] = v$ by the computation

$$\mathbf{E}_{\text{r-acc}}[v] = \frac{\mathbf{E}[(1 - J_r)v] + \lambda_{\text{r-rcy}} \mathbf{E}[S_{\text{r-rcy}}v]}{1 - \rho_r} = \frac{1 - \rho_r - \rho_{\text{r-rcy}} + \rho_{\text{r-rcy}}}{1 - \rho_r} v = v.$$

The result follows from the fact that $\mathbf{E}_{\text{r-acc}}[A_{\text{res}}] = \mathbf{E}_{\text{r-acc}}[\mathbf{E}[A_{\text{res}} \mid A_{\text{age}}]]$ and [Assumption 3.2](#). \square

8.2. Analysis of the recycling cost

Proposition 8.4 (Recycling Cost). In the $G/G/k/\text{setup}$, under the Gittins policy, we have

$$m_{\text{rcy}}^{\text{Gtn}} \leq (k - 1) \log \frac{1}{1 - \rho}.$$

Proof. The same bound has been shown for $M/G/k$ without setup times, e.g. [\[8, Proposition 17.9\]](#) and [\[7, Lemma B.5\]](#). It turns out that the prior proofs rely only on the following fact:

Immediately before an r -recycling, the number of r -relevant jobs is at most $k - 1$.

This fact still holds in the $G/G/k/\text{setup}$ under Gittins, so the same proof goes through. The fact holds because immediately before the recycling, the job that is about to be recycled is in service but is r -irrelevant. If there were k jobs that were r -relevant, they would have priority under Gittins, preventing the r -irrelevant job from being in service and thus preventing the r -recycling. \square

8.3. Analysis of the idleness cost

Proposition 8.5 (Idleness Cost). In the $G/G/k/\text{setup}$, under the Gittins policy, we have

$$m_{\text{idle}}^{\text{Gtn}} \leq (C - 1)(k - 1) \log \frac{1}{1 - \rho} + (k - 1)\mathbb{1}(\mathbf{P}[U > 0] > 0),$$

where $C = \frac{9}{8 \log 1.5} + 1 \approx 3.775$.

The proof of [Proposition 8.5](#) proceeds similarly to the proof of [\[8, Proposition 17.6\]](#), but with a small modification to account for setup times. Given the similarity to prior work, we defer it to [Appendix C](#).

8.4. Analysis of the setup cost

Proposition 8.6 (Single-Server Setup Cost). In the $G/G/1/\text{setup}$, the setup cost is fixed for any setup-non-idling policy. In particular, since Gtn is also a setup-non-idling policy, we have

$$m_{\text{setup}}^{\text{Gtn}} - m_{\text{setup}}^{\pi} = 0,$$

for any other setup-non-idling policy π .

Proof of Proposition 8.6. Because $J_{\text{setup}} \in \{0, 1\}$, we have

$$\mathbf{E}[J_{\text{setup}} W_r]^{\pi} = \mathbf{E}[W_r \mid J_{\text{setup}} = 1]^{\pi} \mathbf{P}[J_{\text{setup}} = 1]^{\pi}.$$

Therefore, recalling [Definition 8.1](#), it suffices to show that both $\mathbf{P}[J_{\text{setup}} = 1]^{\pi}$ and $\mathbf{E}[W_r \mid J_{\text{setup}} = 1]^{\pi}$ do not depend on setup-non-idling policy π . Observe that under any setup-non-idling policy, the distributions of busy periods (the continuous periods when there is work in the system) are unaffected by the order of serving specific jobs, and J_{setup} equals 1 only during the first S_{setup} unit of time in each busy period, so $\mathbf{P}[J_{\text{setup}} = 1]^{\pi}$ do not depend on π . As for $\mathbf{E}[W_r \mid J_{\text{setup}} = 1]^{\pi}$, because the server cannot serve jobs when setting up and there is no r -work in the system when the setup begins, W_r is determined by the amount of r -work that has arrived since the setup begins, whose distribution is independent of the policy. \square

Proposition 8.7 (Multiserver Setup Cost). In the $G/G/k/\text{setup}$, the setup cost under any setup-non-idling π has the following bound:

$$m_{\text{setup}}^{\pi} \leq \mathbb{1}(\mathbf{P}[U > 0] > 0) (\lambda k \mathbf{E}[U_c] + \lambda A_{\text{max}} + k - 1). \tag{8.1}$$

In particular, the bound holds for Gittins.

To prove [Proposition 8.7](#), we require a helper lemma. The lemma bounds the expected number of jobs in the system during a setup time. To state the lemma, we let $J_{\text{setup},i}$ be the indicator of whether server i is setting up and let $U_{\text{age},i}$ be the age of server i 's setup process for each $i = 1, 2, \dots, n$. We set $U_{\text{age},i}$ to zero if server i is not setting up.

Lemma 8.8. *In the G/G/k/setup, for any server i and all a ≥ 0, we have*

$$\mathbf{E}[N \mid J_{\text{setup},i} = 1, U_{\text{age},i} = a] \leq \lambda a + \lambda A_{\text{max}} + k - 1. \tag{8.2}$$

The proof of Lemma 8.8 is nontrivial but uses standard techniques, so we defer it to Appendix C and move on to proving Proposition 8.7.

Proof of Proposition 8.7. The case where U = 0 is clear, so we assume that U > 0 with nonzero probability. We first bound r-setup error using the fact that ρ_r ≤ ρ.

$$m_{\text{setup}} = \int_0^\infty \frac{\mathbf{E}[J_{\text{setup}} W_r]}{r^2(1-\rho_r)} dr \leq \int_0^\infty \frac{\mathbf{E}[J_{\text{setup}} W_r]}{r^2(1-\rho)} dr = \int_0^\infty \frac{\mathbf{E}[\mathbf{E}[J_{\text{setup}} W_r \mid X_1, \dots, X_N]]}{r^2(1-\rho)} dr, \tag{8.3}$$

where X₁, X₂, ... X_N are the states of all the jobs in the system (Section 3.3). Using Tonelli's theorem and WINE (Lemma 6.2), the last expression can be rewritten as

$$\int_0^\infty \frac{\mathbf{E}[\mathbf{E}[J_{\text{setup}} W_r \mid X_1, \dots, X_N]]}{r^2(1-\rho)} dr = \mathbf{E} \left[\frac{J_{\text{setup}}}{1-\rho} \int_0^\infty \frac{\mathbf{E}[W_r \mid X_1, \dots, X_N]}{r^2} dr \right] = \frac{\mathbf{E}[J_{\text{setup}} N]}{1-\rho}. \tag{8.4}$$

By J_{setup} = 1/k ∑_{i=1}^k J_{setup,i}, E[J_{setup,i} N] = E[E[N | J_{setup,i} = 1, U_{age,i}]] and Lemma 8.8,

$$\frac{\mathbf{E}[J_{\text{setup}} N]}{1-\rho} \leq \frac{\lambda}{k(1-\rho)} \sum_{i=1}^k \mathbf{E}[J_{\text{setup},i} U_{\text{age},i}] + \frac{\lambda A_{\text{max}} + k - 1}{k(1-\rho)} \sum_{i=1}^k \mathbf{E}[J_{\text{setup},i}]. \tag{8.5}$$

Now it remains to compute ∑_{i=1}^k E[J_{setup,i}] and ∑_{i=1}^k E[J_{setup,i} U_{age,i}]. The mean fraction of servers setting up is no more than the mean fraction of non-busy servers, which is 1 - ρ, so

$$\frac{1}{k} \sum_{i=1}^k \mathbf{E}[J_{\text{setup},i}] = \mathbf{E}[J_{\text{setup}}] \leq 1 - \rho, \tag{8.6}$$

Basic renewal theory and the 1/k service rate (Section 3.1) imply the average age of a setup time is kE[U_e] (Definition 7.1), so

$$\frac{1}{k} \sum_{i=1}^k \mathbf{E}[J_{\text{setup},i} U_{\text{age},i}] = \sum_{i=1}^k \mathbf{E}[J_{\text{setup},i}] \mathbf{E}[U_e]. \tag{8.7}$$

Combining (8.3)–(8.7) finishes the proof. □

8.5. Proofs of main results

Proof of Theorem 4.1. After expressing the suboptimality gap using Lemma 8.2, we apply Propositions 8.3–8.5 and 8.7 and use the fact that m^π_{setup} is non-negative. Grouping the log 1/(1-ρ) terms to form ℓ_(a) and grouping the 1(P[U > 0] > 0) terms to form ℓ_(c) yields the result. □

Proof of Theorem 4.2. After expressing the suboptimality gap using Lemma 8.2, we apply Propositions 8.3–8.6. The only nonzero contribution comes from m^{Gtn}_{res} - m^π_{res}, which Proposition 8.3 shows to be at most ℓ_(b).

9. Heavy-traffic optimality

We now turn to prove Theorem 4.3, which amounts to showing that Gittins's suboptimality gap, namely E[N]^{Gtn} - inf_π E[N]^π is small relative to the performance of the optimal policy, namely inf_π E[N]^π. It turns out that this is indeed the case: the suboptimality gap is small relative to the performance of SRPT in the G/G/1, which is a lower bound on the performance of any policy in the G/G/k/setup.

We first relate SRPT's performance in the G/G/1 to its performance in the M/G/1, which is known from prior work [26]. We then use this result to prove Theorem 4.3.

Theorem 9.1. *Given Assumption 3.2, in the heavy traffic limit, we have*

$$\lim_{\rho \rightarrow 1} \frac{\mathbf{E}[N]_{\text{G/G/1}}^{\text{SRPT}}}{\mathbf{E}[N]_{\text{M/G/1}}^{\text{SRPT}}} = \frac{c_S^2 + c_A^2}{c_S^2 + 1},$$

where c_v² = Var[V]/E[V]² and the two systems have the same service time distribution and average arrival rate. If c_S² = ∞, then we interpret the right-hand side as 1.

Proof of Theorem 9.1. If c_S² = c_A² = 0, the result holds because E[N] → ∞ in the M/G/1 but not the G/G/1, which in this case are the M/D/1 and D/D/1, respectively (D is for “deterministic”), implying the result. So we focus on the case where c_S² + c_A² > 0. We present the full proof for the c_S² < ∞ case first, briefly sketching how to adapt the argument to the c_S² = ∞ case at the end.

Since SRPT is a special case of Gittins (Example 3.5), we can analyze $E[N]$ in the G/G/1 under SRPT using WINE and the work decomposition law. By Lemma 6.2, Theorem 7.2, and Definition 8.1, we have

$$E[N] = \int_0^\infty \frac{\rho_r(E[(S_r)_e] - E[S_r] + E[A_e]) + \rho_{r\text{-rcy}}E[(S_{r\text{-rcy}})_e]}{r^2(1 - \rho_r)} dr + m_{\text{setup}} + m_{\text{idle}} + m_{\text{rcy}} + m_{\text{res}}.$$

Since SRPT is a Gittins policy, by Propositions 8.3–8.6, we have that in the G/G/1, $m_{\text{setup}} = m_{\text{idle}} = m_{\text{rcy}} = 0$ and $m_{\text{res}} \in [-\lambda A_{\text{max}}, -\lambda A_{\text{min}}]$. We also note that under SRPT, the rank of a job is its remaining size, which together with Definition 7.1 implies

$$E[S_r] = E[S\mathbb{1}(S \leq r)], \quad \rho_r E[(S_r)_e] = \frac{\lambda}{2} E[S^2\mathbb{1}(S \leq r)], \quad \rho_{r\text{-rcy}} E[(S_{r\text{-rcy}})_e] = \frac{\lambda}{2} r^2 \mathbf{P}[S > r].$$

Borrowing the notation of Harchol-Balter [19], we write

$$S_{\bar{r}} = \min\{S, r\} \quad \rho_{\bar{r}} = \lambda E[S_{\bar{r}}],$$

giving us

$$E[N] = \int_0^\infty \frac{\rho_{\bar{r}} E[(S_{\bar{r}})_e] + \rho_r(-E[S_r] + E[A_e])}{r^2(1 - \rho_r)} dr + O(1). \tag{9.1}$$

in the G/G/1 under SRPT. Using Definition 7.1 and the fact that $E[S] = \rho E[A]$, we can compute

$$\rho(E[S_e] - E[S] + E[A_e]) = \frac{\lambda}{2} (\mathbf{Var}[S] + \mathbf{Var}[A]) + O(1 - \rho).$$

Reasoning similarly and using the fact that $\rho_{\bar{r}} - \rho_r \leq 1 - \rho_r$, we have

$$\rho_{\bar{r}} E[(S_{\bar{r}})_e] + \rho_r(-E[S_r] + E[A_e]) = \frac{\lambda}{2} (\mathbf{Var}[S_{\bar{r}}] + \mathbf{Var}[A]) + O(1 - \rho_r).$$

Because $S_{\bar{r}} = \min\{S, r\}$, we have $\lim_{r \rightarrow \infty} \mathbf{Var}[S_{\bar{r}}] = \mathbf{Var}[S]$. Thus, for all $\epsilon > 0$, there exists r^* such that for any $r \geq r^*$, we have $|\mathbf{Var}[S_{\bar{r}}] - \mathbf{Var}[S]| \leq \epsilon$. Since S does not depend on ρ , we can fix a sufficiently small constant ϵ so that r^* is also a constant independent of ρ . Applying (9.1), we can write $E[N]$ as

$$E[N] = \int_0^{r^*} \frac{\rho_{\bar{r}} E[(S_{\bar{r}})_e] + \rho_r(-E[S_r] + E[A_e])}{r^2(1 - \rho_r)} dr + \int_{r^*}^\infty \left(\frac{\frac{\lambda}{2} (\mathbf{Var}[S_{\bar{r}}] + \mathbf{Var}[A])}{1 - \rho_r} + O(1) \right) \frac{1}{r^2} dr + O(1).$$

Observe that the first integral is non-negative, and it can be uniformly bounded at all loads by substituting $\rho_r \mapsto E[S_r]/E[S]$ and $\rho_{\bar{r}} \mapsto E[S_{\bar{r}}]/E[S]$, so it is $O(1)$.¹⁰ As for the second integral, note first that we can ignore the $O(1)$ since $\int_{r^*}^\infty \frac{O(1)}{r^2} dr = O(1)$. Moreover, by our choice of r^* , we have

$$\int_{r^*}^\infty \frac{\frac{\lambda}{2} (\mathbf{Var}[S_{\bar{r}}] + \mathbf{Var}[A])}{r^2(1 - \rho_r)} \frac{1}{r^2} dr = \frac{\lambda}{2} (\mathbf{Var}[S] + \mathbf{Var}[A] + \delta) \int_{r^*}^\infty \frac{1}{r^2(1 - \rho_r)} dr,$$

where $\delta \in [-\epsilon, \epsilon]$. Therefore, we have that in the G/G/1 under SRPT, for some $\delta \in [-\epsilon, \epsilon]$,

$$E[N]_{\text{G/G/1}}^{\text{SRPT}} = \frac{\lambda}{2} (\mathbf{Var}[S] + \mathbf{Var}[A] + \delta) \int_{r^*}^\infty \frac{1}{r^2(1 - \rho_r)} dr + O(1). \tag{9.2}$$

Of course, the M/G/1 is a special case of the G/G/1, so (9.2) also holds for the M/G/1, with $\mathbf{Var}[A] = E[A]^2$ and δ replaced by some other $\delta' \in [-\epsilon, \epsilon]$. Wierman et al. [68, Theorem 5.8] show $E[N]_{\text{M/G/1}}^{\text{SRPT}} = \Omega(\log \frac{1}{1-\rho})$, so

$$\frac{E[N]_{\text{G/G/1-SRPT}}}{E[N]_{\text{M/G/1-SRPT}}} = \frac{\mathbf{Var}[S] + \mathbf{Var}[A] + \delta}{\mathbf{Var}[S] + E[A]^2 + \delta'} + O\left(\frac{1}{\log \frac{1}{1-\rho}}\right) = \frac{\rho^2 c_S^2 + c_A^2 + \delta \lambda^2}{\rho^2 c_S^2 + 1 + \delta' \lambda^2} + O\left(\frac{1}{\log \frac{1}{1-\rho}}\right).$$

The result follows because c_S^2 and c_A^2 are independent of ρ , and $\delta, \delta' \in [-\epsilon, \epsilon]$ for arbitrarily small ϵ .

We have proven the result assuming $c_S^2 < \infty$. If instead $c_S^2 = \infty$, then $\mathbf{Var}[S_r] \rightarrow \infty$ as $r \rightarrow \infty$. This means that for sufficiently large r , the dominant term of W_r is simply $\mathbf{Var}[S_r]/(1 - \rho_r)$, which does not depend on A and is thus the same in the G/G/1 and M/G/1. One can use this fact to show that the performance ratio approaches 1 in heavy traffic by, as in the $c_S^2 < \infty$ case, splitting the WINE integral at large r^* , then observing that the $r > r^*$ part is dominant in heavy traffic. \square

Proof of Theorem 4.3. It suffices to show that the suboptimality gap in Theorem 4.1, which is $O(1)$ for the $k = 1$ case and $O(\frac{1}{1-\rho})$ for the $k \geq 2$ case, is dominated by $E[N]^\pi$ for any policy π .

We begin by observing that for any G/G arrival process, $E[N]_{\text{G/G/1}}^{\text{SRPT}}$ is a lower bound on $E[N]^\pi$ for any policy π . This is because we can view the G/G/k/setup as a version of a G/G/1 that imposes extra constraints on the scheduler (Section 3.3), and SRPT minimizes $E[N]$ in the G/G/1 [1].

Next, we observe that by Theorem 9.1 and our assumption that $c_S^2 + c_A^2 > 0$, SRPT's G/G/1 heavy-traffic performance is within a constant factor of its M/G/1 heavy-traffic performance.

¹⁰ It is not a priori obvious that the integral converges due to the $r \rightarrow 0$ behavior. This can be verified by direct computation, but for our purposes, it suffices to use the prior knowledge that $E[N]$ is finite under SRPT.

It thus suffices to show a lower bound on $E[N]_{M/G/1}^{SRPT}$. When $k = 1$, we need only an $\omega(1)$ bound, which always holds [68, Theorem 5.8]. When $k \geq 2$, we need an $\omega(\log \frac{1}{1-\rho})$ bound, which prior work [6, Proof of Theorem 1.3 in Appendix B.2] shows to hold if $E[S^2(\log S)^+] < \infty$. \square

10. Potential extensions

We have seen that combining our new work decomposition law (Theorem 7.2) with WINE (Lemma 6.2) enables the analysis of systems with many complex features, such as the G/G/k/setup. Thanks to the generality of both systems, we could apply the same technique even beyond the G/G/k/setup. This section sketches how this can be done for three features: multiserver jobs, batch arrivals, and generalized vacations. We emphasize that our goal here is not to give full proofs, but rather to demonstrate the applicability of our technique to additional systems. We thus say “should be” rather than “is” when stating the end results.

10.1. Multiserver jobs

We study a variation of the model of Grosf et al. [7]. We consider a variant of our G/G/k model where each job has a server need $m(x)$, which is a function of its state x . Whenever a job in state x runs, it must occupy exactly $m(x)$ servers. It is thus served at rate $m(x)/k$, thanks to our convention of servers operating at rate $1/k$ (Section 3.1). We refer to this model as the G/G/k/MSJ, where MSJ stands for “multiserver job”. Grosf et al. [7] study what we would call the M/G/k/MSJ.¹¹

One of the main challenges when scheduling in MSJ systems is that it is no longer clear how to stabilize the system. Indeed, analyzing stability even in M/M/k/MSJ systems is an area of current research [13,69], and optimal scheduling in these systems is an open problem. However, Grosf et al. [7] show that if every possible server need $m(x)$ is a divisor of the number of servers k , then one can ensure stability with a procedure called *DivisorFilling*.

The DivisorFilling procedure takes as input any set of k jobs, then outputs a subset of those jobs whose server needs sum to exactly k . DivisorFilling can be combined with the Gittins policy by passing the k jobs of least Gittins rank to DivisorFilling, resulting in a policy called *DivisorFiling-Gittins* [7].

The G/G/k/MSJ under DivisorFiling-Gittins can be analyzed in much the same way as the G/G/k under Gittins. Recalling the structure of the latter analysis from Section 8, we encounter the same four “cost terms” (Definition 8.1) to bound.

- The residual interarrival cost m_{setup} can be bounded exactly as in Proposition 8.3.
- The recycling cost m_{rcy} can be bounded exactly as in Proposition 8.4, because DivisorFiling-Gittins ensures the key property that makes the proof work [7, Lemma B.5].
- The idleness cost m_{idle} can be bounded by following the same steps as [7, Lemma B.3], because their proof does not rely on Poisson arrivals. The resulting bound is $m_{\text{idle}} \leq e(k-1) \lceil \log \frac{1}{1-\rho} \rceil$.¹²
- The setup cost m_{setup} is zero, as there are no setup times.

We thus see that the suboptimality gap of DivisorFiling-Gittins in the G/G/k/MSJ should be at most

$$E[N]_{G/G/k/MSJ}^{\text{DivisorFiling-Gittins}} - \inf_{\pi} E[N]_{G/G/1}^{\pi} \leq e(k-1) \left\lceil \log \frac{1}{1-\rho} \right\rceil + \lambda(A_{\max} - A_{\min}).$$

This is analogous to what Theorem 4.1 says about the G/G/k, as $\ell_{(a)}$ is replaced by $e(k-1) \lceil \log \frac{1}{1-\rho} \rceil$.

It is likely that, for a suitable definition of setup times in an MSJ model, one could analyze the same system with setup times, obtaining a result analogous to the one we obtain for the G/G/k/setup.

10.2. Batch arrivals

Scully and Harchol-Balter [5] introduce a general model of batch arrivals that we call *batch-M/G arrivals*. The main notable feature of the model is that it makes few assumptions about what batches look like. For example, it may be that the initial states of jobs in the same batch are correlated with each other. We consider the same type of batches but allow general batch interarrival times, resulting in *batch-G/G arrivals*. All of the results in Section 4 should generalize to batch-G/G arrivals. There are three changes needed for the proof, and only the third impacts the end results.

First, one needs to modify the definitions of load, r-fresh load, and other concepts related to the arrival process. But these are straightforward changes. The most important note here is that Assumption 3.2 should refer to the batch interarrival time.

Second, batch arrivals affect the work decomposition law (Theorem 7.2). But they only affect the term that is common to all systems, namely the numerator of the $1/(1-\rho_r)$ term. By Lemma 8.2, this change does not affect suboptimality gaps, which is the basis of all of the results in Section 4.

Third, batch arrivals affect the setup cost m_{setup} in multiserver systems. Specifically, we need to slightly modify the statement and proof of Proposition 8.7 to account for the fact that multiple jobs can arrive at once. If more jobs arrive than there are idle

¹¹ Grosf et al. [7] actually consider a slightly more restrictive case in which jobs’ server needs remain constant throughout service, though their proofs could be straightforwardly generalized to handle dynamically changing server needs. The main novelty of our discussion is thus the extension to G/G arrivals.

¹² In the $\rho \rightarrow 1$ limit, this bound below is actually slightly better than that in Proposition 8.5. See Section 4.1 for discussion.

servers, this means a single setup is effectively triggered by multiple jobs. The end result is that one must incorporate a term related to the batch size distribution. In contrast, the setup cost in single-server systems is unaffected.

Taken together, these observations imply that our G/G/k and G/G/1/setup results should generalize immediately to batch arrivals. With some effort, the G/G/k/setup results should also generalize.

10.3. Generalized vacations

The term *generalized vacations* refers to a range of models where servers may be unavailable, including:

- Setup times, as studied in this work. This includes models beyond ours, e.g. where we make different decisions about when to start setting up a server, or where a setup time can be canceled if the job that triggered it enters service at another server.
- *Vacations*, where whenever the server goes idle, it goes on vacation for a given amount of time, then only begins serving jobs again when it returns.
- *Server breakdowns*, where servers can become unavailable in the middle of serving a job.
- *Threshold policies*, where servers stay idle until there are a given number of jobs in the system.

These are only a few examples of what generalized vacations can model [18,51,52].

One can, in principle, bound Gittins's suboptimality gap in the G/G/k with generalized vacations using essentially the same approach we take for the G/G/k/setup. The main change is that we now interpret J_{setup} as the fraction of servers that are unavailable, so we now think of m_{setup} as an *unavailability cost*.

Of course, whether bounding m_{setup} is tractable to bound depends on the specifics of the model. As in the proofs of Propositions 8.6 and 8.7, the key question is: how many jobs might there be in the system while a server is unavailable? Sometimes, this will be very hard to bound, e.g. for server breakdowns. But in other cases, the bound is nearly immediate. For example, consider a threshold policy that does not start serving jobs until there are n jobs present, at which point it serves jobs until the system empties. We would then have $m_{\text{setup}} \in [0, n]$ under any scheduling policy.

One important application of generalized vacations is to more general setup time models. For instance, in practice, it is helpful to not turn servers off as soon as they become idle. One can imagine a wide range of *power management* policies controlling when servers turn on and off. Provided we do not wait too long to set up servers while there are jobs in the queue, m_{setup} should not be too large, in which case Gittins would have a small suboptimality gap. This means that, in some sense, the power management and job scheduling problems are orthogonal, because a single scheduling policy, namely Gittins, performs well for a wide range of power management policies.

11. Conclusion

This work presents the first analysis of the Gittins policy in the G/G/k/setup. We prove simple and explicit bounds on Gittins's suboptimality gap, which are tight enough to imply that Gittins is optimal in heavy traffic in the G/G/k/setup. As a corollary, we find that Gittins is optimal in the M/G/1/setup. Prior to these results, Gittins had not been analyzed in even the G/G/1, let alone the G/G/k/setup.

There are several ways in which one might hope to improve our bounds. This is especially true in light traffic, namely the $\rho \rightarrow 0$ limit. Here we have a constant suboptimality gap for mean number of jobs, but by Little's law [56], this corresponds to an *infinite* suboptimality gap for mean response time. We conjecture that Gittins's mean response time suboptimality gap remains bounded in light traffic, but there are significant obstacles to doing so, related to the notorious problem of analyzing the idle period of the G/G/1 [70,71].

Our theoretical results also raise several questions that could be studied with simulations. One such question is related to the additive structure of our suboptimality gap bound in Theorem 4.1, in which each of (a) multiple servers, (b) non-Poisson arrivals, and (c) setup times contributing to the bound via a separate term. If we simulate Gittins in systems with various mixtures of (a), (b), and (c), do we observe an analogous (approximate) additive structure in its empirical performance? We hypothesize the answer is "yes", because each of the terms in Theorem 4.1 has a distinct cause, and we suspect the interactions between these causes are relatively weak. Investigating this is an interesting direction for future work.

Taking a step back, we might ask: should one use Gittins to minimize mean number of jobs in practice, even beyond the G/G/k/setup modeling assumptions? While this is clearly a question larger than we can definitively answer, we believe that our main results, the potential extensions sketched in Section 10, and other recent work on Gittins and SRPT in multiserver systems [6–8,72,73] point towards "yes". Even though the currently known theoretical bounds on Gittins and SRPT are not tight, we have no comparable bounds for other policies, aside from a few close relatives of SRPT [72]. The mere existence of these bounds is thus a point in favor of Gittins. But we are still in the early years of understanding multiserver scheduling.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Acknowledgments

This research was done in part while Ziv Scully was visiting the Simons Institute for Theoretical Computer Science at UC Berkeley, and in part while he was a FODSI postdoc at Harvard and MIT supported by National Science Foundation, USA grant nos. DMS-2023528 and DMS-2022448. Yige Hong was supported by National Science Foundation, USA grant no. ECCS-2145713.

Appendix A. Summary of notation

Core G/G/k/setup model (Sections 3.1 and 3.2).

- k : number of servers
 - Each server has speed $1/k$, so the total service capacity is always $k \cdot 1/k = 1$
- S : job size random variable
 - With k servers, job of size S results in kS service time
- U : setup work random variable (Section 3.2)
 - With k servers, setup work U results in kU setup time
- A : interarrival time random variable
- $\lambda = 1/\mathbf{E}[A]$: arrival rate
- $\rho = \lambda\mathbf{E}[S] < 1$: load
- N : number of jobs in the system
- $\mathbf{E}[N]_{\text{SYS}}^{\pi}$: mean number of jobs in system SYS under policy π
 - To reduce clutter, we omit “SYS” and/or “ π ” when there is no ambiguity

Remaining interarrival time (Assumption 3.2).

- A_{res} : residual interarrival time, i.e. time until next arrival
- A_{age} : interarrival age, i.e. time since last arrival
- $A_{\text{min}}, A_{\text{max}}$: bounds on $\mathbf{E}[A_{\text{res}} | A_{\text{age}}]$

Markov-process job model (Section 3.3).

- See Examples 3.3 and 3.4 for concrete examples of the Markov-process job model
- \mathbb{X} : space of possible job states
- $X(t)$: state of a job that has received t service so far
- X_i : state of the i th job currently in the system ($i \in \{1, \dots, N\}$)
- \mathbb{T} : terminal state, i.e. a job completes once $X(t) = \mathbb{T}$
- $S(x)$: remaining work of a job in state x
- $W = \sum_{i=1}^N S(X_i)$: (total) work in the system

Gittins policy (Section 3.4).

- Gtn: abbreviation for Gittins in formulas, e.g. $\mathbf{E}[N]^{\text{Gtn}}$
- $\text{rank}_{\text{Gtn}}(x)$: rank, i.e. priority (lower is better), of job in state x
 - See Examples 3.5 and 3.6 for concrete examples of the rank function

Main results (Section 4).

- $\ell_{(a)}, \ell_{(c)}, \ell_{(c)}$: suboptimality due to (a) multiple servers, (b) non-Poisson arrivals, (c) setup times
- $c_V^2 = \frac{\text{Var}[V]}{\mathbf{E}[V]^2}$: squared coefficient of variation of V

WINE and r-work (Section 6).

- $S_r(x)$: remaining r-work of a job in state x
- $W_r = \sum_{i=1}^N S_r(X_i)$: r-work in the system
- r-fresh, r-recycled: see Definitions 6.4 and 6.5
- $S_r = S_r(X(0))$: remaining r-work of a newly arrived job
- $\rho_r = \lambda\mathbf{E}[S_r]$: r-fresh load

- $\lambda_{r\text{-rcy}}$: remaining r -work of a newly r -recycled job
- $S_{r\text{-rcy}}$: remaining r -work of a newly r -recycled job
- $\rho_{r\text{-rcy}} = \lambda_{r\text{-rcy}} \mathbf{E}[S_{r\text{-rcy}}]$: r -recycled load
- J_r : fraction of servers that are serving jobs of rank at most r
- J_{setup} : fraction of servers that are setting up

Stating work decomposition law (Section 7.1).

- $\mathbf{E}_{\text{arv}}[\cdot]$: expectation sampled immediately before an arrival
- $\mathbf{E}_{r\text{-rcy}}[\cdot]$: expectation sampled immediately before an r -recycling
- $\mathbf{E}_{r\text{-acc}}[\cdot]$: an expectation defined in terms of $\mathbf{E}_{r\text{-rcy}}[\cdot]$ in (7.1)
- V_e : excess of random variable V

– Key property: $\mathbf{E}[V_e] = \frac{\mathbf{E}[V^2]}{2\mathbf{E}[V]}$

Proofs of main results (Section 8).

- m_{res}^π : contribution to $\mathbf{E}[N]^\pi$ due to residual interarrival time, bounded in Proposition 8.3
- m_{rcy}^π : contribution to $\mathbf{E}[N]^\pi$ due to recyclings, bounded in Proposition 8.4
- m_{idle}^π : contribution to $\mathbf{E}[N]^\pi$ due to idle servers, bounded in Proposition 8.5
- m_{setup}^π : contribution to $\mathbf{E}[N]^\pi$ due to setup times, bounded in Propositions 8.6 and 8.7

Appendix B. Additional model details

B.1. Processor sharing

Thus far, we have discussed the scheduler as assigning jobs to servers in one-to-one fashion. However, because we are in a preemptive setting, it is possible for the scheduler to effectively share a server between multiple jobs by rapidly switching between them. As such, we explicitly allow the scheduler to use *processor sharing*, simultaneously serving multiple jobs at reduced rates. Processor sharing arises naturally when using Gittins [2,3].

Exactly how processor sharing works with multiple servers and setup times has some subtle corner cases, so for completeness, we give the details below. But for ease of discussion, we do not concern ourselves with these details during the proofs, because they complicate the arguments without adding insight. With that said, modifying the proofs to account for processor sharing is a straightforward exercise. For instance, instead of referring to the job at server i , we might refer to an appropriately weighted random choice among jobs currently sharing server i .

In a system without setup times, the scheduler’s action is a *service rate vector* $\alpha = (\alpha_1, \dots, \alpha_N)$, where recall that N is the number of jobs. Here α_i is the rate at which job i is being served. In the $G/G/k$, the fact that we have k servers of service rate $1/k$ corresponds to α obeying the constraints $\sum_{i=1}^N \alpha_i \leq 1$ and $\alpha_i \in [0, 1/k]$. From this perspective, a non-idling policy is one such that no α_i can be increased without violating a constraint. To clarify, the service rate vector α is a *gradual* control: it is chosen at every moment in time and can be changed at will.

In a system with setup times, the situation is similar, but we can only assign jobs to busy servers. The first constraint on α thus changes to $\sum_{i=1}^N \alpha_i \leq K_{\text{busy}}/k$, where K_{busy} is the number of busy servers.

B.2. General definition of the Gittins rank function

How does Gittins decide what rank to assign each job? The main idea is to assign jobs low rank if they are likely to finish with a small amount of service. To formalize this intuition, consider a job in state x , and suppose we were to serve the job until it enters some state in set $\mathbb{Y} \subseteq \mathbb{X}$. This would involve serving the job for some amount of time, and there would be some chance the job completes. Let

$$S_{\mathbb{Y}}(x) = \text{amount of service a job starting at } x \text{ needs to finish or enter } \mathbb{Y}$$

$$= (\inf \{s \geq 0 \mid X(s) \in \mathbb{Y} \cup \{\text{T}\}\} \mid X(0) = x), \tag{B.1}$$

$$C_{\mathbb{Y}}(x) = \text{event that a job starting at } x \text{ finishes before entering } \mathbb{Y}$$

$$= [X(\inf \{s \geq 0 \mid X(s) \in \mathbb{Y} \cup \{\text{T}\}\}) = \text{T} \mid X(0) = x]. \tag{B.2}$$

We can think of $E[S_{\mathbb{Y}}(x)]/P[C_{\mathbb{Y}}(x)]$ as a time-per-completion ratio for jobs in state x , as measured by serving them until they finish or enter \mathbb{Y} . The key idea of the Gittins policy is to let a job’s rank be its best possible time-per-completion ratio¹³:

$$\text{rank}_{\text{Gtn}}(x) = \inf_{\mathbb{Y} \subseteq \mathbb{X}} \frac{E[S_{\mathbb{Y}}(x)]}{P[C_{\mathbb{Y}}(x)]}. \tag{B.3}$$

As an example, consider the case of known job sizes, where a job’s state x is its remaining work. The best completion ratio is achieved by always running the job until it finishes, i.e. $\mathbb{Y} = \emptyset$. This results in $S_{\emptyset}(x) = x$ and the event $C_{\emptyset}(x)$ occurring with probability 1, so $\text{rank}_{\text{Gtn}}(x) = x$. That is, Gittins always serves the jobs of least remaining work, so it reduces to SRPT (Example 3.5).

B.3. Technical considerations for the job Markov process

We have omitted some technical details and assumptions in our definition of the Gittins policy. For example, in (B.1)–(B.3), we must restrict attention to sets $\mathbb{Y} \subseteq \mathbb{X}$ such that $S_{\mathbb{Y}}(x)$ and $C_{\mathbb{Y}}(x)$ are measurable with respect to the natural filtration on the job Markov process. More generally, the theory of the Gittins policy relies on being able to solve an optimal stopping problem which is known in the Markov-process job model literature as the “Gittins game” [5,6,8].

Following the convention of prior literature, we consider the technical foundations of the Gittins game to be outside the scope of this paper. Our results apply to any job Markov process where the foundations can be established. As explained by Scully et al. [6, Appendix D], this is not a restrictive assumption.

Appendix C. Deferred proofs

Proposition 8.5 (Idleness Cost). *In the G/G/k/setup, under the Gittins policy, we have*

$$m_{\text{idle}}^{\text{Gtn}} \leq (C - 1)(k - 1) \log \frac{1}{1 - \rho} + (k - 1)\mathbb{1}(P[U > 0] > 0),$$

where $C = \frac{9}{8 \log 1.5} + 1 \approx 3.775$.

Proof. At a high level, we follow the same main as the proof of [8, Proposition 17.6], but one of the steps requires modification to account for setup times.

By following the same argument as [8, Lemma 17.4], one can show that

$$m_{\text{idle}}^{\text{Gtn}} = \int_0^\infty \frac{E[(1 - J_r - J_{\text{setup}})W_r]}{r^2(1 - \rho_r)} dr \leq (k - 1) \inf_{n \in \mathbb{Z}_{\geq 1}} n \left(\frac{1}{1 - \rho} \right)^{1/n}. \tag{C.1}$$

The key fact is that when $1 - J_r - J_{\text{setup}} > 0$, there is an idle server, so there are at most $k - 1$ jobs in the system, as otherwise a server would start setting up. Given this fact, the rest of the proof carries through verbatim. We can also show an alternative bound for $m_{\text{idle}}^{\text{Gtn}}$ that is tighter in light traffic:

$$m_{\text{idle}}^{\text{Gtn}} = \int_0^\infty \frac{E[(1 - J_r - J_{\text{setup}})W_r]}{r^2(1 - \rho_r)} dr \leq (k - 1) \frac{\rho}{1 - \rho} + (k - 1)\mathbb{1}(P[U > 0] > 0). \tag{C.2}$$

We prove (C.2) below. Combining (C.1) and (C.2) yields

$$\begin{aligned} m_{\text{idle}}^{\text{Gtn}} &\leq (k - 1) \min \left\{ \inf_{n \in \mathbb{Z}_{\geq 1}} n \left(\frac{1}{1 - \rho} \right)^{1/n}, \frac{\rho}{1 - \rho} + \mathbb{1}(P[U > 0] > 0) \right\} \\ &\leq (k - 1) \min \left\{ \inf_{n \in \mathbb{Z}_{\geq 1}} n \left(\frac{1}{1 - \rho} \right)^{1/n}, \frac{\rho}{1 - \rho} \right\} + (k - 1)\mathbb{1}(P[U > 0] > 0) \\ &\leq (C - 1)(k - 1) \log \frac{1}{1 - \rho} + (k - 1)\mathbb{1}(P[U > 0] > 0), \end{aligned}$$

where the last inequality follows from the same computation as in [8, Proposition 17.6].

It remains only to prove (C.2). The proof follows the same steps as [8, Lemma 17.5], but setup times add some additional considerations. For the purposes of this proof, we abuse notation slightly by defining $X_i = \top$ for the state of job i if there is no i th job. We then define the remaining r -work of \top to be $S_r(\top) = 0$.

We begin by noting that if $J_r + J_{\text{setup}} < 1$, then there are less than k jobs in the system, so

$$E[(1 - J_r - J_{\text{setup}})W_r] = \sum_{i=1}^k E[S_r(X_i)(1 - J_r - J_{\text{setup}})] \leq \frac{k - 1}{k} \sum_{i=1}^k E[S_r(X_i)]. \tag{C.3}$$

¹³ We consider the ratio to be ∞ , and thus not the infimum, if $P[C_{\mathbb{Y}}(x)] = 0$.

where the last equality follows from the fact that $J_r + J_{\text{setup}} < 1$ is an integer multiple of $1/k$. Applying Lemma 6.3 and (C.3) to the definition of $m_{\text{idle}}^{\text{Gtn}}$ yields

$$\begin{aligned} m_{\text{idle}}^{\text{Gtn}} &= \int_0^\infty \frac{\mathbf{E}[(1 - J_r - J_{\text{setup}})W_r]}{r^2(1 - \rho_r)} dr \\ &\leq \frac{k - 1}{k(1 - \rho)} \int_0^\infty \sum_{i=1}^k \frac{\mathbf{E}[S_r(X_i)]}{r^2} dr \\ &\leq \frac{k - 1}{k(1 - \rho)} \sum_{i=1}^k \mathbf{P}[X_i \neq \top] \\ &= \frac{(k - 1)\mathbf{E}[\min\{N, k\}]}{k(1 - \rho)}. \end{aligned} \tag{C.4}$$

The fact that we are considering the setup-non-idling version of Gittins (Section 3.2) implies that the number of non-idle servers (i.e. busy or setting-up) is always at least $\min\{N, k\}$. So it suffices to bound the mean fraction of non-idle servers.

- The mean fraction of busy servers is ρ .
- The mean fraction of setting-up servers is at most $(1 - \rho)\mathbb{1}(\mathbf{P}[U > 0] > 0)$.
 - If there are no setup times, it is clearly zero.
 - If there are setup times, it is at most the fraction of non-busy servers, namely $1 - \rho$.

This means

$$\frac{(k - 1)\mathbf{E}[\min\{N, k\}]}{k(1 - \rho)} \leq (k - 1)\frac{\rho}{1 - \rho} + (k - 1)\mathbb{1}(\mathbf{P}[U > 0] > 0). \tag{C.5}$$

Combining (C.4) and (C.5) yields (C.2), as desired. \square

Lemma 8.8. *In the G/G/k/setup, for any server i and all $a \geq 0$, we have*

$$\mathbf{E}[N \mid J_{\text{setup},i} = 1, U_{\text{age},i} = a] \leq \lambda a + \lambda A_{\text{max}} + k - 1. \tag{8.2}$$

Proof. Observe that during setup, the number of jobs in the system can be divided into three sets: jobs present when setup started, jobs triggered the setup and jobs that have arrived since setup started. Therefore, to bound $\mathbf{E}[N \mid J_{\text{setup},i} = 1, U_{\text{age},i} = a]$, it is sufficient to bound the total number of jobs in the three sets on the right. Just before the setup starts, there is at least one idle server, so there are at most $k - 1$ jobs present. And there is exactly 1 job that triggers the setup. Let $N_{\text{arv}}(a)$ be the number of jobs that have arrived since a unit of time ago, then the above argument shows that

$$\mathbf{E}[N \mid J_{\text{setup},i} = 1, U_{\text{age},i} = a] \leq \mathbf{E}[N_{\text{arv}}(a) \mid J_{\text{setup},i} = 1, U_{\text{age},i} = a] + k. \tag{C.6}$$

Therefore, to prove the lemma, it is sufficient to show that for any $a \geq 0$,

$$\mathbf{E}[N_{\text{arv}}(a) \mid J_{\text{setup},i} = 1, U_{\text{age},i} = a] \leq \lambda a + \lambda A_{\text{max}} - 1. \tag{C.7}$$

Conditioned on $J_{\text{setup},i} = 1$ and $U_{\text{age},i} = a$, we know that server i started setup process a unit of time ago. Since the setup process is triggered by an arrival, if we denote the sequence of interarrival times since the setup as $A_1, A_2, \dots, A_n, \dots$, then we have the equation

$$\mathbf{E}\left[\sum_{n=1}^{N_{\text{arv}}(a)+1} A_n \mid J_{\text{setup},i} = 1, U_{\text{age},i} = a\right] = \mathbf{E}[a + A_{\text{res}} \mid J_{\text{setup},i} = 1, U_{\text{age},i} = a]. \tag{C.8}$$

The right-hand side of (C.8) can be upper bounded using Assumption 3.2 and the fact that A_{res} is conditionally independent of the past given the age:

$$\mathbf{E}[a + A_{\text{res}} \mid J_{\text{setup},i} = 1, U_{\text{age},i} = a] = \mathbf{E}[a + A_{\text{res}} \mid U_{\text{age},i} = a] \leq a + A_{\text{max}}. \tag{C.9}$$

To bound the left-hand side of (C.8), observe that $J_{\text{setup},i} = 1$ and $U_{\text{age},i} = a$ is equivalent to having a setup process start a unit of time ago and last until now, which is determined by the setup time and is independent of arrival process in the recent a unit of time, so we have

$$\mathbf{E}\left[\sum_{n=1}^{N_{\text{arv}}(a)+1} A_n \mid J_{\text{setup},i} = 1, U_{\text{age},i} = a\right] = \mathbf{E}\left[\sum_{n=1}^{N_{\text{arv}}(a)+1} A_n\right]. \tag{C.10}$$

We further observe that $A_1, A_2, \dots, A_n, \dots$ are a sequence of i.i.d. variables, and $N_{\text{arv}} + 1$ is the smallest n such that the partial sum $A_1 + \dots + A_n \geq a$, by Wald's equation,

$$\mathbf{E}\left[\sum_{n=1}^{N_{\text{arv}}(a)+1} A_n\right] = \mathbf{E}[N_{\text{arv}}(a) + 1]\mathbf{E}[A] = \frac{\mathbf{E}[N_{\text{arv}}(a) + 1]}{\lambda}. \tag{C.11}$$

Combining (C.8)–(C.11), we get $\mathbf{E}[N_{\text{arv}}(a) + 1]/\lambda \leq a + A_{\text{max}}$, implying (C.7). \square

Appendix D. Towards proving stability of the G/G/k/setup under work-conserving scheduling policies

In this section, we sketch some initial ideas that we believe could be used to prove the stability of the G/G/k/setup under Gittins and other complex scheduling policies. For simplicity, we focus on the G/G/k under stationary scheduling policies.

D.1. Partial proof sketch

Assume we are using some scheduling policy that makes the system a Markov process. Consider the embedded Markov process of our system at arrival instants. We construct a Lyapunov function and show its negative drift for this embedded Markov process when the total work in the system or the largest job size is outside a bounded region. The negative drift of the Lyapunov function should imply the stability of the system after applying a certain version of the Foster–Lyapunov theorem, although there are still some technical obstacles, as we discuss in the next section.

At any time $t \geq 0$, we let $W(t)$ be the total work in the system, let $S_{\text{big}}(t)$ be the size of the largest job in the system, and let $A_{\text{res}}(t)$ be the residual arrival time.

Fix large constants $\alpha > 0$ and $\beta > 0$. Define the Lyapunov function $V = W - \rho A_{\text{res}} + k(S_{\text{big}} - \alpha)^+$. Note that at each arrival point, $A_{\text{res}} = 0$ and $V \geq 0$. The continuous rate of change of V is

$$\begin{aligned} \dot{V} &= -(1 - \rho) + 1 - J - k\mathbb{1}\{S_{\text{big}} > \alpha\} \dot{S}_{\text{big}}(t) \\ &\leq -(1 - \rho) + 1 - J - \mathbb{1}\{S_{\text{big}} > \alpha, J < 1\} \\ &\leq -(1 - \rho) + (1 - J)\mathbb{1}\{S_{\text{big}} \leq \alpha\}, \end{aligned}$$

where in the first inequality, we use the fact that when $J < 1$, all jobs are in service, so the size of the largest job decreases at a rate $1/k$; and in the second inequality, we use the fact that $\mathbb{1}\{J < 1\} \geq 1 - J$. The expected jump of V at an arrival point is

$$\begin{aligned} \Delta V &= \mathbf{E}[S - \rho A + k \max((S_{\text{big}} - \alpha)^+, (S - \alpha)^+) - k(S_{\text{big}} - \alpha)^+] \\ &\leq k\mathbf{E}[(S - \alpha)^+], \end{aligned}$$

where we have used the fact that $\mathbf{E}[S] = \rho\mathbf{E}[A]$. Consider the change of V between two consecutive arrivals:

$$\begin{aligned} \mathbf{E}[V(A)] - V(0) &= \int_0^A \dot{V}(t) dt + \lambda \Delta V \\ &\leq \mathbf{E}\left[\int_0^A (-(1 - \rho) + (1 - J(t))\mathbb{1}\{S_{\text{big}}(t) \leq \alpha\}) dt\right] + \lambda k\mathbf{E}[(S - \alpha)^+] \\ &= -\lambda(1 - \rho) + \mathbf{E}\left[\int_0^A (1 - J(t))\mathbb{1}\{S_{\text{big}}(t) \leq \alpha\} dt\right] + \lambda k\mathbf{E}[(S - \alpha)^+]. \end{aligned}$$

Let $\mathcal{R} = [0, k(\alpha + \beta) + \beta] \times [0, \alpha + \beta]$, a compact region. We show that $\mathbf{E}[V(A)] - V(0) < 0$ as long as $(W(0), S_{\text{big}}(0)) \notin \mathcal{R}$. There are two cases to consider:

- If $S_{\text{big}}(0) > \alpha + \beta$, then $S_{\text{big}}(t) > \alpha + \beta - t$. As a result,

$$\mathbf{E}[V(A)] - V(0) \leq -\lambda(1 - \rho) + \mathbf{E}[(A - \beta)^+] + \lambda k\mathbf{E}[(S - \alpha)^+].$$

- If $S_{\text{big}}(0) \leq \alpha + \beta$ but $W(0) > k(\alpha + \beta) + \beta$, then we have $S_{\text{big}}(t) \leq \alpha + \beta$, and $W(t) \geq k(\alpha + \beta) + \beta - t$. For $t \leq \beta$, $N(t) \geq W(t)/S_{\text{big}} \geq k$, implying $J(t) = 1$. As a result,

$$\mathbf{E}[V(A)] - V(0) \leq -\lambda(1 - \rho) + \mathbf{E}[(A - \beta)^+] + \lambda k\mathbf{E}[(S - \alpha)^+].$$

Taking sufficiently large α and β , we get $\mathbf{E}[V(A)] - V(0) < 0$ in both cases.

D.2. Remaining obstacles


We have given a Lyapunov function that has negative drift for the embedded Markov process when (W, S_{big}) is outside a compact region \mathcal{R} . However, (W, S_{big}) is not a complete state description of the Markov process. Moreover, the set of states such that $(W, S_{\text{big}}) \in \mathcal{R}$ might not be compact. We believe one can show that this set of states is *petite* [74], even if it is not compact, which will suffice. However, this seems to require reasoning about the details of the job state space.

Of course, so far we have just been focusing on the existence of a stationary distribution for the embedded process at arrival instants. Once this is established, we need to use it to construct a stationary distribution for the continuous-time process, but this step is more standard.

References

- [1] Linus E. Schrage, A proof of the optimality of the shortest remaining processing time discipline, *Oper. Res.* 16 (3) (1968) 687–690, <http://dx.doi.org/10.1287/opre.16.3.687>.
- [2] Samuli Aalto, Urtzi Ayesta, Rhonda Righter, On the Gittins index in the $M/G/1$ queue, *Queueing Syst.* 63 (1–4) (2009) 437–458, <http://dx.doi.org/10.1007/s11134-009-9141-x>.
- [3] Samuli Aalto, Urtzi Ayesta, Rhonda Righter, Properties of the Gittins index with application to optimal scheduling, *Probab. Engrg. Inform. Sci.* 25 (3) (2011) 269–288, <http://dx.doi.org/10.1017/S0269964811000015>.
- [4] John C. Gittins, Kevin D. Glazebrook, Richard R. Weber, *Multi-Armed Bandit Allocation Indices*, second ed., Wiley, Chichester, UK, 2011.
- [5] Ziv Scully, Mor Harchol-Balter, The Gittins policy in the $M/G/1$ queue, in: 19th International Symposium on Modeling and Optimization in Mobile, Ad Hoc, and Wireless Networks (WiOpt 2021), IFIP, Philadelphia, PA, 2021, pp. 248–255, <http://dx.doi.org/10.23919/WiOpt52861.2021.9589051>.
- [6] Ziv Scully, Isaac Grosf, Mor Harchol-Balter, The Gittins policy is nearly optimal in the $M/G/k$ under extremely general conditions, *Proc. ACM Meas. Anal. Comput. Syst.* 4 (3) (2020) 43, <http://dx.doi.org/10.1145/3428328>, 29 pages.
- [7] Isaac Grosf, Ziv Scully, Mor Harchol-Balter, Alan Scheller-Wolf, Optimal scheduling in the multiserver-job model under heavy traffic, *Proc. ACM Meas. Anal. Comput. Syst.* 6 (3) (2022) 1–32, <http://dx.doi.org/10.1145/3570612>.
- [8] Ziv Scully, A New Toolbox for Scheduling Theory (Ph.D. Dissertation), Carnegie Mellon University, Pittsburgh, PA, 2022, <https://ziv.codes/pdf/scully-thesis.pdf>.
- [9] Percy H. Brill, Linda Green, Queues in which customers receive simultaneous service from a random number of servers: A system point approach, *Manage. Sci.* 30 (1) (1984) 51–68, <http://dx.doi.org/10.1287/mnsc.30.1.51>.
- [10] Mor Harchol-Balter, Open problems in queueing theory inspired by datacenter computing, *Queueing Syst.* 97 (1) (2021) 3–37, <http://dx.doi.org/10.1007/s11134-020-09684-6>.
- [11] Weina Wang, Qiaomin Xie, Mor Harchol-Balter, Zero queueing for multi-server jobs, *Proc. ACM Meas. Anal. Comput. Syst.* 5 (1) (2021) 7, <http://dx.doi.org/10.1145/3447385>, 25 pages.
- [12] Yige Hong, Weina Wang, Sharp waiting-time bounds for multiserver jobs, 2022, [arXiv:2109.05343](https://arxiv.org/abs/2109.05343) [cs, math] <http://arxiv.org/abs/2109.05343>.
- [13] Alexander Rummyantsev, Evsey Morozov, Stability criterion of a multiserver model with simultaneous service, *Ann. Oper. Res.* 252 (1) (2017) 29–39, <http://dx.doi.org/10.1007/s10479-015-1917-2>.
- [14] Alexander Rummyantsev, Robert Basmadjian, Sergey Astafiev, Alexander Golovin, Three-level modeling of a speed-scaling supercomputer, *Ann. Oper. Res.* (2022) <http://dx.doi.org/10.1007/s10479-022-04830-0>.
- [15] Isaac Grosf, Mor Harchol-Balter, Alan Scheller-Wolf, WCFS: A new framework for analyzing multiserver systems, *Queueing Syst.* 102 (1–2) (2022) 143–174, <http://dx.doi.org/10.1007/s11134-022-09848-6>.
- [16] Mohan L. Chaudhry, James G.C. Templeton, *A First Course in Bulk Queues*, Wiley, New York, NY, 1983.
- [17] S.W. Fuhrmann, A note on the $M/G/1$ queue with server vacations, *Oper. Res.* 32 (6) (1984) 1368–1373.
- [18] S.W. Fuhrmann, Robert B. Cooper, Stochastic decompositions in the $M/G/1$ queue with generalized vacations, *Oper. Res.* 33 (5) (1985) 1117–1129, <http://dx.doi.org/10.1287/opre.33.5.1117>.
- [19] Mor Harchol-Balter, *Performance Modeling and Design of Computer Systems: Queueing Theory in Action*, Cambridge University Press, Cambridge, UK, 2013.
- [20] Yuan Li, David A. Goldberg, Simple and explicit bounds for multi-server queues with universal $1/(1 - \rho)$ scaling, 2017, [arXiv:1706.04628](https://arxiv.org/abs/1706.04628) [math] <http://arxiv.org/abs/1706.04628>.
- [21] Jalani K. Williams, Mor Harchol-Balter, Weina Wang, The $M/M/k$ with deterministic setup times, *Proc. ACM Meas. Anal. Comput. Syst.* 6 (3) (2022) 1–45, <http://dx.doi.org/10.1145/3570617>.
- [22] Kevin D. Glazebrook, José Niño-Mora, Parallel scheduling of multiclass $M/M/m$ queues: Approximate and heavy-traffic optimization of achievable performance, *Oper. Res.* 49 (4) (2001) 609–623, <http://dx.doi.org/10.1287/opre.49.4.609.11225>.
- [23] Kevin D. Glazebrook, An analysis of Klimov’s problem with parallel servers, *Math. Methods Oper. Res.* 58 (1) (2003) 1–28, <http://dx.doi.org/10.1007/s001860300278>.
- [24] Rhonda Righter, J. George Shanthikumar, Genji Yamazaki, On extremal service disciplines in single-stage queueing systems, *J. Appl. Probab.* 27 (2) (1990) 409–416, <http://dx.doi.org/10.2307/3214660>.
- [25] Sayan Banerjee, Amarjit Budhiraja, Amber L. Puhua, Heavy traffic scaling limits for shortest remaining processing time queues with heavy tailed processing time distributions, *Ann. Appl. Probab.* 32 (4) (2022) 2587–2651, <http://dx.doi.org/10.1214/21-AAP1741>.
- [26] Minghong Lin, Adam Wierman, Bert Zwart, Heavy-traffic analysis of mean response time under shortest remaining processing time, *Perform. Eval.* 68 (10) (2011) 955–966, <http://dx.doi.org/10.1016/j.peva.2011.06.001>.
- [27] John C. Gittins, David M. Jones, A dynamic allocation index for the sequential design of experiments, in: Joseph M. Gani, Károly Sarkadi, István Vincze (Eds.), *Progress in Statistics*, in: *Colloquia Mathematica Societatis János Bolyai*, Number 9, North-Holland, Amsterdam, The Netherlands, 1974, pp. 241–266, <https://ci.nii.ac.jp/ncid/BA03522866>.
- [28] Rhonda Righter, J. George Shanthikumar, Scheduling multiclass single server queueing systems to stochastically maximize the number of successful departures, *Probab. Engrg. Inform. Sci.* 3 (3) (1989) 323–333, <http://dx.doi.org/10.1017/S0269964800001194>.
- [29] Ziv Scully, Mor Harchol-Balter, Alan Scheller-Wolf, SOAP: One clean analysis of all age-based scheduling policies, *Proc. ACM Meas. Anal. Comput. Syst.* 2 (1) (2018) 16, <http://dx.doi.org/10.1145/3179419>, 30 pages.
- [30] Ziv Scully, Mor Harchol-Balter, Alan Scheller-Wolf, Simple near-optimal scheduling for the $M/G/1$, *Proc. ACM Meas. Anal. Comput. Syst.* 4 (1) (2020) 11, <http://dx.doi.org/10.1145/3379477>, 29 pages.
- [31] Rocco van Vreumingen, *Queueing Systems with Non-Standard Service Policies and Server Vacations* (Master’s thesis), University of Amsterdam, Amsterdam, The Netherlands, 2019, <http://www.mt-support.nl/VanVreumingen/Rocco/wp-content/uploads/2020/01/Master-thesis-Queueing-systems-with-non-standard-service-policies-and-server-vacations.pdf>.
- [32] Ziv Scully, Isaac Grosf, Mor Harchol-Balter, Optimal multiserver scheduling with unknown job sizes in heavy traffic, *Perform. Eval.* 145 (2021) 102150, <http://dx.doi.org/10.1016/j.peva.2020.102150>, 31 pages.
- [33] Mor Harchol-Balter, Allen B. Downey, Exploiting process lifetime distributions for dynamic load balancing, *ACM Trans. Comput. Syst.* 15 (3) (1997) 253–285, <http://dx.doi.org/10.1145/263326.263344>.
- [34] Mor Harchol-Balter, *Network Analysis without Exponential Assumptions* (Ph.D. thesis), University of California, Berkeley, Berkeley, CA, 1996, <https://www.cs.cmu.edu/~harchol/Thesis/thesis.html>.
- [35] Mark E. Crovella, Azer Bestavros, Self-similarity in world wide web traffic: Evidence and possible causes, *IEEE/ACM Trans. Netw.* 5 (6) (1997) 835–846, <http://dx.doi.org/10.1109/90.650143>.
- [36] Kihong Park, Walter Willinger (Eds.), *Self-Similar Network Traffic and Performance Evaluation*, Wiley, New York, NY, 2000, <http://dx.doi.org/10.1002/047120644X>.
- [37] David L. Peterson, Data center I/O patterns and power laws, in: 22nd International Computer Measurement Group Conference, Computer Measurement Group, San Diego, CA, 1996, pp. 1034–1045, http://www.cmg.org/?s2member_file_download=/proceedings/1996/96INT631.pdf.

- [38] Ziv Scully, Lucas van Kreveld, Onno J. Boxma, Jan-Pieter Dorsman, Adam Wierman, Characterizing policies with optimal response time tails under heavy-tailed job sizes, *Proc. ACM Meas. Anal. Comput. Syst.* 4 (2) (2020) 30, <http://dx.doi.org/10.1145/3392148>, 33 pages.
- [39] Jesus R. Artalejo, Antonis Economou, Maria J. Lopez-Herrero, Analysis of a multiserver queue with setup times, *Queueing Syst.* 51 (1–2) (2005) 53–76, <http://dx.doi.org/10.1007/s11134-005-1740-6>.
- [40] Anshul Gandhi, Mor Harchol-Balter, $M/G/k$ with Exponential Setup, Technical Report CMU-CS-09-166, Carnegie Mellon University, Pittsburgh, PA, 2009, p. 44, <http://dx.doi.org/10.1184/R1/6607073>.
- [41] Anshul Gandhi, Mor Harchol-Balter, Ivo Adan, Server farms with setup costs, *Perform. Eval.* 67 (11) (2010) 1123–1138, <http://dx.doi.org/10.1016/j.peva.2010.07.004>.
- [42] Anshul Gandhi, Mor Harchol-Balter, $M/G/k$ with staggered setup, *Oper. Res. Lett.* 41 (4) (2013) 317–320, <http://dx.doi.org/10.1016/j.orl.2013.03.006>.
- [43] Anshul Gandhi, Sherwin Doroudi, Mor Harchol-Balter, Alan Scheller-Wolf, Exact analysis of the $M/M/k$ /setup class of Markov chains via recursive renewal reward, *Queueing Syst.* 77 (2) (2014) 177–209, <http://dx.doi.org/10.1007/s11134-014-9409-7>.
- [44] Jamol Pender, Tuan Phung-Duc, A law of large numbers for $M/M/c$ /Delayoff-Setup queues with nonstationary arrivals, in: Sabine Wittevrongel, Tuan Phung-Duc (Eds.), *Analytical and Stochastic Modelling Techniques and Applications*, Vol. 9845, Springer, Cham, Switzerland, 2016, pp. 253–268, http://dx.doi.org/10.1007/978-3-319-43904-4_18.
- [45] Peter D. Welch, On a generalized $M/G/1$ queueing process in which the first customer of each busy period receives exceptional service, *Oper. Res.* 12 (5) (1964) 736–752, <http://dx.doi.org/10.1287/opre.12.5.736>.
- [46] Bharat T. Doshi, A note on stochastic decomposition in a $GI/G/1$ queue with vacations or set-up times, *J. Appl. Probab.* 22 (2) (1985) 419–428, <http://dx.doi.org/10.2307/3213784>.
- [47] Wolfgang Bischof, Analysis of $M/G/1$ -queues with setup times and vacations under six different service disciplines, *Queueing Syst.* 39 (4) (2001) 265–301, <http://dx.doi.org/10.1023/A:1013992708103>.
- [48] Gautam Choudhury, On a batch arrival Poisson queue with a random setup time and vacation period, *Comput. Oper. Res.* 25 (12) (1998) 1013–1026, [http://dx.doi.org/10.1016/S0305-0548\(98\)00038-0](http://dx.doi.org/10.1016/S0305-0548(98)00038-0).
- [49] Qi-Ming He, Elizabeth Jewkes, Flow time in the $MAP/G/1$ queue with customer batching and setup times, *Commun. Statist. Stoch. Models* 11 (4) (1995) 691–711, <http://dx.doi.org/10.1080/15326349508807367>.
- [50] Huan Li, Yixin Zhu, A new approach to the $G/G/1$ queue with generalized setup time and exhaustive service, *J. Appl. Probab.* 31 (4) (1994) 1083–1097, <http://dx.doi.org/10.2307/3215331>.
- [51] Masakiyo Miyazawa, Decomposition formulas for single server queues with vacations : A unified approach by the rate conservation law, *Commun. Statist. Stoch. Models* 10 (2) (1994) 389–413, <http://dx.doi.org/10.1080/15326349408807301>.
- [52] Bharat T. Doshi, Queueing systems with vacations – a survey, *Queueing Syst.* 1 (1) (1986) 29–66, <http://dx.doi.org/10.1007/BF01149327>.
- [53] Shweta Upadhyaya, Queueing systems with vacation: An overview, *Int. J. Math. Oper. Res.* 9 (2) (2016) 167, <http://dx.doi.org/10.1504/IJMOR.2016.077997>.
- [54] Onno J. Boxma, Wim P. Groenendijk, Pseudo-conservation laws in cyclic-service systems, *J. Appl. Probab.* 24 (4) (1987) 949–964, <http://dx.doi.org/10.2307/3214218>.
- [55] Bharat T. Doshi, Generalizations of the stochastic decomposition results for single server queues with vacations, *Commun. Statist. Stoch. Models* 6 (2) (1990) 307–333, <http://dx.doi.org/10.1080/15326349908807149>.
- [56] John D.C. Little, Little's law as viewed on its 50th anniversary, *Oper. Res.* 59 (3) (2011) 536–549, <http://dx.doi.org/10.1287/opre.1110.0940>.
- [57] Robert M. Loynes, The stability of a queue with non-independent inter-arrival and service times, *Math. Proc. Camb. Phil. Soc.* 58 (3) (1962) 497–520, <http://dx.doi.org/10.1017/S0305004100036781>.
- [58] Jack Kiefer, Jacob Wolfowitz, On the theory of queues with many servers, *Trans. Amer. Math. Soc.* 78 (1) (1955) 1–18, <https://www.jstor.org/stable/1992945>.
- [59] Ward Whitt, Existence of limiting distributions in the $GI/GI/s$ queue, *Math. Oper. Res.* 7 (1) (1982) 88–94, <http://dx.doi.org/10.1287/moor.7.1.88>.
- [60] Karl Sigman, One-dependent regenerative processes and queues in continuous time, *Math. Oper. Res.* 15 (1) (1990) 175–189, <http://www.jstor.org/stable/3689937>.
- [61] Evsey Morozov, Bart Steyaert, *Stability Analysis of Regenerative Queueing Models: Mathematical Methods and Applications*, Springer, Cham, Switzerland, 2021, <http://dx.doi.org/10.1007/978-3-030-82438-9>.
- [62] Anton Braverman, J.G. Dai, Masakiyo Miyazawa, Heavy traffic approximation for the stationary distribution of a generalized Jackson network: The BAR approach, *Stoch. Syst.* 7 (1) (2017) 143–196, <http://dx.doi.org/10.1287/15-SSY199>.
- [63] Masakiyo Miyazawa, Diffusion approximation for stationary analysis of queues and their networks: A review, *J. Oper. Res. Soc. Japan* 58 (1) (2015) 104–148, <http://dx.doi.org/10.15807/jorsj.58.104>.
- [64] Ziv Scully, Isaac Grosf, Michael Mitzenmacher, Uniform bounds for scheduling with job size estimates, in: 13th Innovations in Theoretical Computer Science Conference (ITCS 2022), in: *Leibniz International Proceedings in Informatics (LIPIcs)*, Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik, Berkeley, CA, 2022, 41, <http://dx.doi.org/10.4230/LIPIcs.ITCS.2022.114>, 30 pages.
- [65] Masakiyo Miyazawa, Rate conservation laws: A survey, *Queueing Syst.* 15 (1) (1994) 1–58, <http://dx.doi.org/10.1007/BF01189231>.
- [66] François Baccelli, Pierre Brémaud, *Elements of Queueing Theory: Palm-Martingale Calculus and Stochastic Recurrences*, second ed., *Applications of Mathematics*, Number 26, Springer, Berlin, Germany, 2003.
- [67] Søren Asmussen, *Applied Probability and Queues*, second ed., in: *Stochastic Modelling and Applied Probability*, Number 51, Springer, New York, NY, 2003, <http://dx.doi.org/10.1007/b97236>.
- [68] Adam Wierman, Mor Harchol-Balter, Takayuki Osogami, Nearly insensitive bounds on SMART scheduling, *ACM SIGMETRICS Perform. Eval. Rev.* 33 (1) (2005) 205–216, <http://dx.doi.org/10.1145/1071690.1064236>.
- [69] Isaac Grosf, Mor Harchol-Balter, Alan Scheller-Wolf, Stability for two-class multiserver-job systems, 2020, [arXiv:2010.00631](https://arxiv.org/abs/2010.00631) [cs]. <http://arxiv.org/abs/2010.00631>.
- [70] Ronald W. Wolff, Chia-Li Wang, Idle period approximations and bounds for the $GI/G/1$ queue, *Adv. Appl. Probab.* 35 (3) (2003) 773–792, <http://dx.doi.org/10.1239/aap/1059486828>.
- [71] Jingwen Li, Jihong Ou, Characterizing the idle-period distribution of $GI/G/1$ queues, *J. Appl. Probab.* 32 (1) (1995) 247–255, <http://dx.doi.org/10.2307/3214933>.
- [72] Isaac Grosf, Ziv Scully, Mor Harchol-Balter, SRPT for multiserver systems, *Perform. Eval.* 127–128 (2018) 154–175, <http://dx.doi.org/10.1016/j.peva.2018.10.001>.
- [73] Isaac Grosf, Ziv Scully, Mor Harchol-Balter, Load balancing guardrails: Keeping your heavy traffic on the road to low response times, *Proc. ACM Meas. Anal. Comput. Syst.* 3 (2) (2019) 42, <http://dx.doi.org/10.1145/3341617.3326157>, 31 pages.
- [74] Sean P. Meyn, Richard L. Tweedie, *Markov Chains and Stochastic Stability*, second ed., in: *Cambridge Mathematical Library*, Cambridge University Press, Cambridge, UK, 2009, <http://dx.doi.org/10.1017/CBO9780511626630>.



Yige Hong is a PhD student at Carnegie Mellon University in the Computer Science Department, advised by Weina Wang. Yige works on the performance analysis and optimal control of complex stochastic systems like multiserver jobs, the $G/G/k$ model, stochastic bin-packing, and restless bandits. A particular emphasis of his work is understanding those systems at a large scale, as motivated by the needs of data centers and other large service systems. Yige uses stochastic analysis tools like Lyapunov functions, coupling, and Stein's method to design scalable control policies with provable guarantees.



Ziv Scully is an assistant professor at Cornell University in the School of Operations Research and Information Engineering. He completed his PhD in Computer Science at Carnegie Mellon University in 2022, advised by Mor Harchol-Balter and Guy Blelloch. Between graduating from CMU and joining Cornell, Ziv was a research fellow at the UC Berkeley Simons Institute for the Data-Driven Decision Processes program; and then an NSF FODSI postdoc at Harvard and MIT, mentored by Michael Mitzenmacher and Piotr Indyk. Ziv researches the theory of decision making under uncertainty, with a particular emphasis on scheduling and resource allocation. Ziv's work has been recognized by awards from INFORMS, ACM SIGMETRICS, and IFIP PERFORMANCE, including winning the 2022 George Nicholson Student Paper Competition and receiving the 2022 SIGMETRICS Doctoral Dissertation Award.