

Improving Multiresource Job Scheduling with Markovian Service Rate Policies

ZHONGRUI CHEN, University of North Carolina at Chapel Hill, USA

ISAAC GROSOFF, Northwestern University, USA

BENJAMIN BERG, University of North Carolina at Chapel Hill, USA

Modern cloud computing workloads are composed of *multiresource jobs* that require a variety of computational resources in order to run, such as CPU cores, memory, disk space, or hardware accelerators. A single cloud server can typically run many multiresource jobs in parallel, but only if the server has sufficient resources to satisfy the demands of every job. A scheduling policy must therefore select sets of multiresource jobs to run in parallel in order to minimize the *mean response time* across jobs – the average time from when a job arrives to the system until it is completed. Unfortunately, achieving low response times by selecting sets of jobs that fully utilize the available server resources has proven to be a difficult problem.

In this paper, we develop and analyze a new class of policies for scheduling multiresource jobs, called Markovian Service Rate (MSR) policies. While prior scheduling policies for multiresource jobs are either highly complex to analyze or hard to implement, our MSR policies are simple to implement and are amenable to response time analysis. We show that the class of MSR policies is *throughput-optimal* in that we can use an MSR policy to stabilize the system whenever it is possible to do so. We also derive bounds on the mean response time under an MSR algorithm that are tight up to an additive constant. These bounds can be applied to systems with different preemption behaviors, such as fully preemptive systems, non-preemptive systems, and systems that allow preemption with setup times. We show how our theoretical results can be used to select a good MSR policy as a function of the system arrival rates, job service requirements, the server's resource capacities, and the resource demands of the jobs.

CCS Concepts: • **Mathematics of computing** → **Markov processes**; • **Computer systems organization** → **Cloud computing**; • **Computing methodologies** → **Model development and analysis**.

Additional Key Words and Phrases: Scheduling, Queueing Theory, Cluster Scheduling, Multiresource Jobs, Multiserver Jobs, Markovian Service Rate

ACM Reference Format:

Zhongrui Chen, Isaac Groszof, and Benjamin Berg. 2025. Improving Multiresource Job Scheduling with Markovian Service Rate Policies. *Proc. ACM Meas. Anal. Comput. Syst.* 9, 2, Article 25 (June 2025), 36 pages. <https://doi.org/10.1145/3727117>

1 Introduction

Modern data center servers include a wide variety of computational resources, including CPU cores, network cards, local or disaggregated memory and storage, and specialized accelerators such as GPUs. Each job request in the data center requires some set of diverse resources in order to run, such as a few CPU cores, some amount of memory, and some amount of storage [25, 26, 36].

Authors' Contact Information: [Zhongrui Chen](mailto:zhongrui@cs.unc.edu), jcpwfloi@cs.unc.edu, University of North Carolina at Chapel Hill, Chapel Hill, North Carolina, USA; [Isaac Groszof](mailto:izzy.groszof@northwestern.edu), izzy.groszof@northwestern.edu, Northwestern University, Evanston, Illinois, USA; [Benjamin Berg](mailto:ben@cs.unc.edu), ben@cs.unc.edu, University of North Carolina at Chapel Hill, Chapel Hill, North Carolina, USA.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

© 2025 Copyright held by the owner/author(s). Publication rights licensed to ACM.

ACM 2476-1249/2025/6-ART25
<https://doi.org/10.1145/3727117>

Each job also has an associated *service requirement*, which describes how long the job must run once it has been allocated its requested resources. Because a single job requires several types of resources, we refer to these jobs as *multiresource jobs*. To serve multiresource jobs quickly, a single server generally aims to run many multiresource jobs in parallel. However, the server also has a limited capacity for each resource. Hence, a set of multiresource jobs can only be run in parallel if their aggregate demand for each resource is less than the server's capacity for that resource. For example, in order to run some set of jobs concurrently, the number of total CPU cores demanded by the jobs should not exceed the total number of CPU cores in the server, and the total memory required by these jobs should not exceed the DRAM capacity of the server. This raises the question of how a scheduling policy should choose sets of multiresource jobs to run in parallel on a data center server.

When scheduling multiresource jobs, the goal is often to reduce job *response times*, the time from when a job arrives to the system until it is completed [5, 36]. In particular, this paper aims to design scheduling policies to minimize *mean response time*, the long-run average response time across a stream of arriving multiresource jobs. Unfortunately, optimizing the response times of multiresource jobs has proven to be difficult [11, 18, 25, 34]. Much of the literature on multiresource jobs only proves system stability results [11, 26], while response time analysis has been largely restricted to simple policies such as *First-Come-First-Served* (FCFS) that perform poorly in practice [16]. Furthermore, many existing results on scheduling multiresource jobs suggest using a MaxWeight policy [26] that is too complex to run in a real system. As a result, modern data centers are generally greatly overprovisioned to achieve low response times [36], and rely on heuristic scheduling policies, such as BackFilling policies like First-Fit, that admit no theoretical analysis or guarantees [14].

This paper introduces a new class of policies for scheduling multiresource jobs called Markovian Service Rate (MSR) policies. We prove strong stability results for this class of policies and bounds on mean response time. Using these bounds, we derive an accurate approximation of the mean response time of an MSR policy. Best of all, our MSR policies require a single offline optimization step, after which the policy runs in constant time per scheduling decision.

1.1 The Problem: Scheduling to Reduce Wastage

The key challenge in scheduling multiresource jobs is that it is difficult to maximize the average utilization of all the server resources. Intuitively, finding sets of jobs with high utilization of the server resources requires solving a bin-packing problem. Hence, naïve scheduling policies can greatly underutilize system resources. We describe the average amount of idle resources under a policy as the *wastage* of the policy. In general, high wastage leads to higher mean response time, and possibly even instability.

To illustrate this problem, we consider the FCFS policy for multiresource jobs described in [16]. Whenever a job arrives or departs the system, FCFS repeatedly examines the job at the front of the queue and puts the job in service as long as the server has enough of resources to accommodate it. This process continues until either the queue is empty or the next job in the queue would exceed the server's capacity for one or more resources. Although it is a simple, intuitive policy, FCFS can greatly underutilize the available system resources. For example, consider the case where there are two *types* of jobs, with type-1 jobs requesting 4 CPU cores and type-2 jobs requesting 2 CPU cores. Assume that each job must run for one second, the server has 8 CPU cores in total, and arriving jobs alternate between type-1 and type-2 jobs. In this case, FCFS grabs a 4-core job and a 2-core job from the front of the queue. The next job in the queue requires four cores, but only two are available, so the remaining two cores in the server will be idle. Hence, in our example, FCFS always wastes at least two cores and completes at most two jobs per second. Alternatively, consider a no-wastage policy that alternates between running four 2-core jobs together or two 4-core jobs together. This

Table 1. Comparison of MSR policies to other multiresource job scheduling policies.

Policy	Throughput Optimal	Low Complexity	Mean Response Time Analysis	Preemption Model
FCFS [16]		x	x	non-preemptive
First-Fit [14]		x		non-preemptive
ServerFilling [14] ¹	Sometimes	x	Sometimes	preemptive
MaxWeight [26]	x		Heavy traffic ²	preemptive
Randomized-Timers [31]	x	x		non-preemptive
MSR	x	x	x	general

no-wastage policy can complete up to three jobs per second. As a result, the no-wastage policy may perform well in cases where the FCFS is unstable. We formalize this argument in Section 5.

FCFS suffers in this example because its high wastage limits the number of jobs that can be run in parallel. While it is easy to spot better packings of the jobs in this simple example, the complexity of finding low-wastage allocations explodes as the number of job types grows (e.g., some jobs request three cores) and as the number of resource types grows (e.g., jobs request CPUs, GPUs, and DRAM).

1.2 Our Goal: Improved Low-complexity Scheduling Policies

To avoid the pitfalls of an FCFS-style policy, the canonical work on scheduling multiresource jobs suggests finding low-wastage allocations by solving complex bin-packing problems repeatedly to continuously determine the next scheduler action [26, 34]. These are the so-called *MaxWeight* algorithms that have been shown to maximize the capacity region of the system (i.e., throughput optimality) and achieve empirically low mean response time. Unfortunately, running *MaxWeight* in real systems is impractical due to the complexity of the bin-packing problem that must be solved repeatedly. *MaxWeight* also preempts jobs frequently, which may not be feasible in all systems. It is a long-standing open problem to find and analyze policies that are both simple enough to run in practice, but still achieve low mean response time similar to a *MaxWeight* policy.

The closest the scheduling community has come to solving this problem is the *Randomized-Timers* approach of [11], which describes a non-preemptive policy for scheduling multiresource jobs that can compute each scheduling decision with low complexity (complexity is linear in the number of job types). *Randomized-Timers* is throughput-optimal, but the policy often results in high mean response time compared to *MaxWeight*. Additionally, the *Randomized-Timers* algorithm is a complicated randomized algorithm and there is no analysis of its mean response time.

Our goal in this paper is to develop policies with the best features of *MaxWeight* and *Randomized-Timers*. That is, we seek low-complexity policies that can achieve low mean response time under a variety of job preemption models. Our policies should stabilize the system whenever possible, and should be accompanied by strong theoretical guarantees about mean response time.

1.3 Our Solution: MSR Scheduling Policies

This paper analyzes the *Markovian Service Rate* (MSR) class of policies for scheduling multiresource jobs. MSR policies use a finite-state Continuous Time Markov Chain (CTMC) to determine which set of jobs to run at every moment in time. The states of this CTMC, known as the *candidate set* of

¹ServerFilling is throughput optimal and has a mean response time analysis in the restricted setting of single-dimensional jobs with power-of-two resource requirements, but is not throughput optimal in the full multiresource setting we consider.

²*MaxWeight* undergoes state space collapse, which can be used to give a heavy traffic analysis of mean response time [34].

possible scheduling actions, as well as the CTMC's transition rates, are determined offline via a one-time optimization step that considers the service requirement distributions and arrival rates for a given workload. We show how to select an MSR policy that stabilizes the system whenever it is possible to do so. By deriving additively-tight mean response time bounds for any MSR policy, we also show that a well-chosen MSR policy can run with low complexity and achieve a low mean response time. Specifically, we show how to construct an MSR policy that is constant-competitive with MaxWeight in a heavy-traffic limit. Furthermore, we find that MSR policies are comparable to MaxWeight at all loads in simulation.

Our analysis generalizes to a variety of job preemption models. We consider the cases where jobs are fully preemptible, where jobs are non-preemptible, and where job preemptions are accompanied by setup times. In each case, we show how to select an appropriate MSR policy and analyze the mean response time under this policy. Specifically, the contributions of the paper are as follows:

- First, Section 5 introduces the class of MSR scheduling policies, which are easy to implement with low computational complexity. We show that, under a variety of job preemption behaviors, there exists an MSR policy that can stabilize the system if it is possible to do so. Hence, we say that the class of MSR policies is throughput-optimal under each preemption behavior.
- Next, in Section 6, we analyze the mean response time of the multiresource job system under any MSR policy by decoupling the system into a set of $M/M/1$ Markovian Modulated Service Rate systems. By analyzing these decoupled problems, we prove additively tight bounds on the mean queue length of an MSR policy. We use our bounds to derive a mean response time approximation for MSR policies that is shown to be highly accurate in simulation.
- In Sections 7 and 8, we use our theoretical results to optimize the choice of MSR policy under a variety of job preemption behaviors. We validate our choices by comparing MSR policies to policies from the literature in both theory and in simulation using a trace from the Google Borg cluster scheduler. We find that our MSR policies are competitive with these other policies, while also being analytically tractable and simple to implement. Specifically, after performing an offline optimization step, the MSR policies we evaluate run in constant time.

Table 1 compares our new MSR policies to existing policies for scheduling multiresource jobs.

2 Prior Work

Scheduling in Real-world Systems. Multiresource jobs are ubiquitous in datacenters and super-computing centers where each job requires some amount of several different resources (e.g., CPUs, GPUs, memory, disk space, network bandwidth) in order to run. As a result, the systems community has extensively considered the problem of scheduling multiresource jobs [5, 19, 24, 33, 36, 37]. Unfortunately, almost all of this work depends on a complex set of heuristic scheduling policies that do not provide theoretical performance guarantees.

The Multiserver Job Model. Prior theoretical work has considered the multiserver job model [18] where each job demands multiple servers in order to run. Several papers have developed throughput-optimal policies for this multiserver job system, such as the MaxWeight Policy [25] and the Randomized-Timers policy [11]. However, neither of these policies admits a general analysis of mean response time. Other work [13, 14, 38] has analyzed the multiserver job system in a variety of special cases such as restricting jobs to belong to one of two classes [13], considering the system in a variety of scaling limits [38], and requiring job server demands to be powers-of-2 [14]. In these special cases, the analysis relies heavily on the fact that there is only one resource type, servers.

The Multiresource Job Model. The multiresource job model is a generalization of the multiserver job model. In this setting, jobs can now request several different types of resources. Several papers [3, 11, 12, 20, 25, 31] have considered the multiresource job model. In [25] the MaxWeight policy is

also shown to be throughput optimal in the multiresource job model. Similarly, the Randomized-Timers policy [11, 31] is throughput optimal in the multiresource job setting. Analyses of mean response time in the multiresource job setting are more limited. For example, [16] recently analyzed the mean response time of a simple FCFS policy in the multiresource job case. [3] proposes a throughput-optimal policy but the response time analysis assumes there are only 2 types of jobs.

Markov Modulated Service Rate Systems. Our plan is to analyze MSR policies in the multiresource job setting. These are policies whose modulating processes are CTMCs. As a result, the multiresource job systems that we analyze will look like a collection of single-server queueing systems with Markov modulated service rates. These Markov modulated systems have been studied extensively in the queueing literature [2, 4, 6, 7, 9, 16, 28, 32]. While [2, 9] first studied the heavy-traffic behavior of a system with Markov modulated *arrival rates*, [6] generalized their results to allow Markov modulated service rates. Matrix analytic methods [4, 28, 32] have also been used to analyze these systems. More recently, [7, 16] have obtained response time bounds in Markov modulated systems by using drift-based techniques. We will extend these drift-based techniques to help us analyze mean response time in the multiresource job model.

Scheduling in Networks. While we focus on analyzing multiresource job scheduling, some work on scheduling in networks considers related problems. Achievable throughput analysis in generalized switches is analogous to trying to find the stability region of multiresource job systems [27, 34, 35]. Other work considers scheduling in wireless networks, where only a subset of wireless channels can send messages at the same time, leading to bin-packing effects [23, 29]. Unfortunately, the policies discussed in this work are too complex for response time analysis [23, 29, 34], and/or are computationally expensive [34].

3 Model

Our goal is to model the performance of jobs running in a data center. When a job is submitted to the system, a load balancer will first select a server to run each job, and each individual server is then responsible for scheduling its assigned jobs [36]. This paper models the behavior of a single server in this setting, deferring the load balancing problem to future work.

We consider a server with R different types of resources. Let the vector $\mathbf{P} = (P_1, P_2, \dots, P_R) \in \mathbb{R}^R$ denote the resource capacity of the server, where P_j is the amount of resource j that the server has. The server is tasked with completing a stream of jobs that arrive to the system over time. All dispatched jobs are either in service, or are held in a central queue that can be served in any order.

3.1 Multiresource Jobs

We refer to the jobs in our model as *multiresource jobs*. A multiresource job is defined by an ordered pair (\mathbf{d}, x) , where $\mathbf{d} \in \mathbb{R}_+^R$ is a *resource demand vector* denoting the job's requirement for each resource type and $x \in \mathbb{R}_+$ is the job's *service requirement*, denoting the amount of time the job must be run on the server before it is completed. We assume a job's resource demands are constant throughout its lifetime. The server can therefore run multiple multiresource jobs in parallel if the sum of the jobs' resource demands does not exceed the server capacity for each resource.

We consider workloads composed of a finite number of *job types* representing the different kinds of applications. Intuitively, jobs of the same type represent different instances of the same program or different virtual machines of the same instance type. Hence, jobs of the same type have the same resource demand vector. Specifically, we assume there are K types of jobs, and that all type- i jobs share the same resource demand row vector \mathbf{D}_i . Let $\mathbf{D} = (\mathbf{D}_1, \mathbf{D}_2, \dots, \mathbf{D}_K)$ denote the matrix of these demand vectors. We assume the service requirements of type- i jobs are sampled i.i.d. from a common service distribution, $S_i \sim \exp(\mu_i)$. While the system knows the resource demands of each arriving job, we assume that service requirements are unknown to the system.

We define a *possible schedule* to be a set of jobs that can be run in parallel without violating any of the resource constraints of the server. We describe a possible schedule using a vector $\mathbf{u} = (u_1, u_2, \dots, u_K) \in \mathbb{Z}_+^K$, where u_i denotes how many type- i jobs we are putting into service in parallel. Because service requirements are exponentially distributed, this description suffices to describe the behavior of the system. The vector \mathbf{u} is a possible schedule if and only if $\sum_{i=1}^K u_i \mathbf{D}_i = \mathbf{u} \mathbf{D} \leq \mathbf{P}$. We call the set of all possible schedules the *schedulable set*, \mathcal{S} . Formally, $\mathcal{S} = \{\mathbf{u} \in \mathbb{Z}_+^K \mid \mathbf{u} \cdot \mathbf{D} \leq \mathbf{P}\}$.

3.2 Arrival Process

We consider the case where a stream of incoming jobs arrives at the system over time. Let $\{A_i(t), t \geq 0\}$ be the counting process tracking the number of type- i arrivals before time t . For convenience, we define $A_i(t_1, t_2) := A_i(t_2) - A_i(t_1)$ as the number of arrivals during the time interval $[t_1, t_2]$. We assume A_i evolves according to a Poisson process with rate λ_i . Let $\boldsymbol{\lambda} = (\lambda_1, \lambda_2, \dots, \lambda_K) \in \mathbb{R}_+^K$ denote the vector of job arrival rates and $\Lambda = \sum_{i=1}^K \lambda_i$ denote the total arrival rate. Finally, let $\mathbf{A}(t) = (A_1(t), A_2(t), \dots, A_K(t))$ denote the vector of cumulative arrivals by time t .

3.3 Scheduling Policies

A *scheduling policy* p in the multiresource job model chooses a possible schedule to use at every moment in time to serve some subset of the jobs in the system. We describe p via a *modulating process* $\{m^p(t), t \geq 0\}$ with N^p states, where $m : \mathbb{R}_+ \rightarrow \{1, 2, \dots, N^p\}$. Note that N^p is a property of the policy p , not an exponential function. Let m^p be the stationary distribution of the modulating process. Each state in the modulating process corresponds to a schedule in \mathcal{S} . Let $\mathbf{u}^p(t) = (u_1^p(t), u_2^p(t), \dots, u_K^p(t))$ be the schedule that corresponds to $m^p(t)$ for any $t \geq 0$. For notational convenience, we also define $u_{s,i}^p = u_i^p(t \mid m^p(t) = s)$ for all $s \in \{1, 2, \dots, N^p\}$ and $i \in \{1, 2, \dots, K\}$. Note that the schedules corresponding to $m^p(t)$ may, in general, depend on the system state (e.g., queue lengths).

We define a class of simple scheduling policies called *Markovian Service Rate* (MSR) policies in Definition 1.

Definition 1 (MSR Policies). *An MSR policy p is a scheduling policy whose modulating process $\{m^p(t), t \geq 0\}$ is an irreducible Continuous Time Markov Chain (CTMC) with $N^p < \infty$ states, where the updates of the CTMC are conditionally independent of the underlying queueing system given the system events (e.g. completions and setups). Let \mathbf{G}^p denote the infinitesimal generator of the CTMC. Let \mathbf{u}^p denote the limiting distribution of the schedule used by policy p , and let $\mathbb{E}[\mathbf{u}^p]$ be the steady-state average schedule under policy p .*

MSR policies are simple in the sense that they use a finite-state CTMC to choose their schedules at every moment in time. Furthermore, this CTMC is required to be mostly independent of the system state, although its state transitions can correspond to arrivals and departures in the system. Specifically, the MSR policy's modulating process is not allowed to be correlated with any queue length information. Our definition is compatible with the definition of an MMSR system in [16]. In this paper, we will show that the class of MSR policies is throughput-optimal and can achieve low mean response time while also being simpler to analyze and implement than existing approaches such as MaxWeight and Randomized-Timers.

3.4 Preemption

As described in Section 2, the choice of scheduling policy for multiresource jobs has historically depended on the preemption behavior of jobs. To capture this dynamic, we consider jobs with a variety of preemption behaviors. Each preemption behavior imposes a different set of constraints on what a policy's modulating process can do. For instance, when jobs are preemptible with no

overhead, a modulating process may transition arbitrarily between schedules by preempting jobs as necessary. However, when jobs are non-preemptible, a modulating process cannot perform transitions that would necessitate preemptions. For example, a non-preemptive policy cannot directly transition to a schedule with a much smaller number of type- i jobs in service than the current schedule. This paper considers the case where jobs are *fully preemptible* with no overhead, the case where jobs are *non-preemptible*, and the case where jobs can be preempted by waiting for a *setup time* [10].

The class of MSR policies can handle multiresource jobs under each of these preemption models by imposing different restrictions on its modulating process. For preemptible jobs, the modulating process can be any irreducible, finite-state CTMC. For non-preemptible jobs, the modulating process must avoid preemptions altogether. That is, jobs can be added to the schedule if there is sufficient capacity, otherwise jobs must be removed from the schedule one at a time at moments of job completions. When jobs are preemptible with setup times, the modulating process will spend time in some additional states that correspond to the required setup times. Details of each of these restrictions on state transitions are discussed formally in Section 5.1. We refer to MSR policies with *preemptible* jobs, *non-preemptible* jobs, and jobs with *setup times* as *pMSR*, *nMSR*, and *sMSR* policies, respectively.

3.5 Queuing Dynamics

We model the queue length as the CTMC $\{\mathbf{Q}(t) = (Q_1(t), Q_2(t), \dots, Q_K(t)), t \geq 0\}$, where $Q_i(t)$ tracks the number of type- i jobs in system at time t . The queueing dynamics of the system follow

$$\mathbf{Q}(t + \delta) = \mathbf{Q}(t) + \mathbf{A}(t, t + \delta) - \hat{\mathbf{C}}(t, t + \delta), \quad (1)$$

where $\hat{\mathbf{C}}(t, t + \delta)$ denotes the number of jobs that completed service during time interval $[t, t + \delta)$. We call $\hat{\mathbf{C}}(t, t + \delta)$ *actual completions* during time interval $[t, t + \delta)$.

Then, we will relate $\hat{\mathbf{C}}(t, t + \delta)$ to the schedule we pick at time t , $\mathbf{u}(t)$. Note that the server may not be serving $\mathbf{u}_i(t)$ type- i jobs at time t if there are not enough type- i jobs in the system. The actual number of type- i jobs in service is $\min(\mathbf{u}_i(t), Q_i(t))$. We let $\mathbf{C}(t, t + \delta)$ denote the number of *potential completions* during the interval $[t, t + \delta)$ by simulating the number of jobs that would have completed if exactly $\mathbf{u}(t)$ jobs were in service. Consequently, $\mathbf{C}(t, t + \delta) \geq \hat{\mathbf{C}}(t, t + \delta)$. We let $\mathbf{Z}(t, t + \delta)$ denote the *unused service* during time interval $[t, t + \delta)$ where $\hat{\mathbf{C}}(t, t + \delta) + \mathbf{Z}(t, t + \delta) = \mathbf{C}(t, t + \delta)$.

We can then rewrite (1) as $\mathbf{Q}(t + \delta) = \mathbf{Q}(t) + \mathbf{A}(t, t + \delta) - \mathbf{C}(t, t + \delta) + \mathbf{Z}(t, t + \delta)$, or

$$\mathbf{Q}(t + \delta) = (\mathbf{Q}(t) + \mathbf{A}(t, t + \delta) - \mathbf{C}(t, t + \delta))^+,$$

where the positive part of x , $(x)^+$, is defined as $(x)^+ = \max(x, 0)$. Note that for any job type, i , $\mathbf{Z}_i(t, t + \delta) > 0$ implies $\mathbf{Q}_i(t + \delta) = 0$ because the queue has to be empty before unused service occurs. Because job service requirements are exponentially distributed as $\exp(\mu_i)$, we have

$$E[\hat{\mathbf{C}}(t, t + \delta)] \leq E[\mathbf{C}(t, t + \delta)] = \boldsymbol{\mu} \otimes \int_t^{t+\delta} \mathbf{u}(x) dx \quad \forall \delta > 0,$$

where $\boldsymbol{\mu}$ is the *service requirement vector* defined as $\boldsymbol{\mu} = (\mu_1, \mu_2, \dots, \mu_K)$, and \otimes denotes the *Hadamard product*.

3.6 Performance Metrics

Our goal is to analyze the steady-state behavior of the server as $t \rightarrow \infty$. We define

$$\mathbf{Q} \sim \lim_{t \rightarrow \infty} \mathbf{Q}(t).$$

We are interested in analyzing the *steady-state mean queue length* of the server, $\mathbb{E}[Q]$, where

$$\mathbb{E}[Q] = \frac{1}{\Lambda} \langle \lambda, \mathbb{E}[Q] \rangle = \frac{1}{\Lambda} \langle \lambda, \lim_{t \rightarrow \infty} \mathbb{E}[Q(t)] \rangle.$$

We can then apply Little's Law to determine the mean response time across jobs, $\mathbb{E}[T]$, as

$$\mathbb{E}[T] = \frac{\mathbb{E}[Q]}{\Lambda},$$

where T is the *steady-state response time of a job in the system*. Hence, to analyze $\mathbb{E}[T]$, we will focus on analyzing $\mathbb{E}[Q]$.

While there has been significant work on analyzing the fairness of a scheduling policy by looking at higher moments of response time [21] or *mean slowdown* [39], this work focuses on optimizing the mean response time across jobs. As we will discuss in Section 7.1, we can design policies that both achieve good mean response time and follow mild fairness properties.

Note that the above steady-state limits may not exist in all cases. We say the system is *stable* under a particular scheduling policy if $Q(t)$ is ergodic and the steady-state queue length distribution exists [17]. When the system is stable, the steady-state distribution of Q exists and $\mathbb{E}[Q] < \infty$. The *capacity region*, C^p , of a scheduling policy p is defined as the set of arrival rates such that the system is stable under p . The capacity region of a *system*, C , is the set of all arrival rates such that *some* scheduling policy exists that stabilizes the system. If a policy's capacity region equals the capacity region of the system, then we say this scheduling policy is *throughput-optimal* [26].

The results of [26, 31] show that the capacity region of any multiresource job system is

$$C = \{ \lambda \mid (1 + \epsilon)\lambda \oslash \mu \in \text{Conv}(\mathcal{S}), \epsilon > 0 \} \quad (2)$$

where $\text{Conv}(\mathcal{S})$ is the convex set of points that can be written as a weighted average of the points in \mathcal{S} and $\lambda \oslash \mu$ denotes Hadamard division between the vectors. Intuitively, there exists a scheduling policy that can stabilize the system if and only if there exists a convex combination of possible schedules such that the average service rate of type- i jobs exceeds the average arrival rate of type- i jobs for all $i \in [1, K]$.

Here, we define the load on the system in terms of utilization of the server when the arrival rate vector is λ . Specifically, we define the *system load*, $\rho > 0$, to be the smallest number such that

$$\frac{\lambda \oslash \mu}{\rho} \in \text{Conv}(\mathcal{S}). \quad (3)$$

Note that if $\lambda \in C$, then $\rho < 1$ according to (2).

Then, [26] defines the *MaxWeight* policy $\mathbf{u}^{\text{MaxWeight}}(t)$ as the solution to the optimization problem

$$\mathbf{u}^{\text{MaxWeight}}(t) = \arg \max_{\mathbf{u} \in \mathcal{S}} \langle \mathbf{u}, \mathbf{Q}(t) \rangle.$$

This policy is known to be throughput-optimal. Unfortunately, MaxWeight also requires frequent preemption and repeatedly optimizing over the set of all possible schedules. Solving this NP-hard optimization problem can be prohibitively expensive. We aim to devise and analyze throughput-optimal policies that avoid this complexity.

3.7 Summary of Modeling Assumptions

To maintain tractability, our model includes a few assumptions that warrant discussion.

Job Sizes and Arrivals. We assume that job sizes are exponentially distributed, while real-world job sizes generally follow heavily-tailed distributions. Fortunately, the analytic techniques in Section 6 can potentially be extended to handle phase-type job sizes, providing a promising direction for future work. We describe these extensions in Appendix G. We have also assumed a stationary

Poisson arrival process that is less variable than arrivals in a real system. We address both of these issues by evaluating our results using real-world data from the Google Borg scheduler in Section 8. **Steady State Analysis.** Our analysis develops queueing-theoretic results about the steady-state behavior of the system. In a system where the workload changes rapidly and unpredictably, one might be concerned that the system never reaches stationarity. The results of Section 8 show that our scheduling policies still perform well over the course of a thirty-day trace of real-world datacenter traffic. Our analysis should be viewed as complementary to the worst-case and competitive analysis approaches that assume adversarial workloads and result in loose performance bounds [22].

Resource Requirements. We assume that job resource requirements are constant throughout a job's lifetime. While jobs in some systems can change resource demands in real time [1], systems like Borg let users reserve resources statically. While jobs may not fully use their demanded resources, scheduling decisions in reservation-based systems are still largely based on these reservations [36].

We also assume that the jobs' resource requirements are known beforehand. This assumption is practical in some settings, such as the case where a cloud provider sells virtual machines of different instance types. In other cases, the demands may have to be predicted by looking at recent runs.

Heterogeneous Resources. While the fleet of servers in a modern datacenter is heterogeneous, this work focuses on scheduling within a single server. Within a server, we assume that resources of the same type are homogeneous. For some resources, such as GPUs, the performance difference between GPU models can be so huge that they should be handled as separate resource types. For CPU cores, differences between cores are more easily abstracted away from the scheduler. For example, Google Borg normalizes CPU requests to *Google Compute Units* so that 1 unit provides the same performance regardless of the underlying hardware [36]. While resources such as memory and disk space can have variable performance characteristics, there are many mechanisms in software and hardware designed to make these resources appear homogeneous as a program runs.

4 Overview of Results

The rest of this paper is devoted to analyzing various MSR policies. Broadly, our theoretical results can be grouped into two categories. First, in Section 5, we prove results about the stability of the system under various MSR policies, and show how to select an MSR policy that stabilizes the system under various preemption models. Second, in Sections 6 and 7, we analyze the mean response time under an MSR policy and use this analysis to further guide the selection of a stable MSR policy with low mean response time. We now provide an overview of these results.

Stability Results. First, we prove the throughput-optimality of the class of MSR policies when jobs are fully preemptible. In Theorem 2, we show that if it is possible for any policy to stabilize the system, there exists a pMSR policy that will stabilize the system.

Theorem 2. *When scheduling preemptible jobs with no preemption overhead, the class of MSR policies is throughput-optimal. That is, if there exists a scheduling policy that can stabilize a multiresource job system with K job types and arrival rates λ , then there exists an MSR policy, p , with $N^p \leq K$ candidate schedules that stabilizes the system. Specifically, an MSR policy can stabilize a system with preemptible jobs for any $\lambda \in C$.*

This result has a nice geometric intuition. We prove that choosing a pMSR policy to stabilize the system is equivalent to finding a set of candidate schedules that lie on $\text{Conv}(\mathcal{S})$ whose completion rate vectors span the arrival vector in the K -dimensional space. We then argue that, by Carathéodory theorem, there exists a set of at most K schedules that satisfy this condition, meaning there exists an MSR policy, p , with $N^p \leq K$.

We note that our proof of Theorem 2 is not constructive — it does not provide a method for selecting a pMSR policy. Hence, in Corollary 3, we construct a pMSR policy to stabilize the system by solving a constrained optimization problem over the space of possible modulating processes. We state this formally in Corollary 3.

Corollary 3. *For a pMSR policy p , let \mathbf{W}^p denote the matrix of candidate schedules,*

$$\mathbf{W}^p = \begin{pmatrix} u_{1,1}^p & u_{1,2}^p & \cdots & u_{1,K}^p \\ \cdots & \cdots & \cdots & \cdots \\ u_{NP,1}^p & u_{NP,2}^p & \cdots & u_{NP,K}^p \end{pmatrix} \in \mathbb{Z}_{NP \times K}^+.$$

Let π^p be the limiting distribution of the modulating process $m^p(t)$, $\pi^p = (\pi_1^p, \pi_2^p, \dots, \pi_{NP}^p) \in \mathbb{R}^{NP}$. To construct a stable pMSR policy, it is necessary to find values \mathbf{W}^p and π^p such that

$$\begin{cases} \mathbf{W}^p \mathbf{D} \leq \mathbf{P} & \text{(Capacity constraint)} \\ \pi_i^p \geq 0, \forall i & \text{(Positive fractions)} \\ \pi_1^p + \pi_2^p + \cdots + \pi_{NP}^p \leq 1 & \text{(Fractions sum to 1)} \\ \pi^p \mathbf{W}^p > \lambda \odot \boldsymbol{\mu} & \text{(Stability condition from Lemma 1)} \\ \mathbf{W}^p \text{ integers} \end{cases}.$$

Because Theorem 2 guarantees that there exists an MSR policy, p , with $NP \leq K$ whenever it is possible to stabilize the system, we can safely assume that $NP = K$ when solving this optimization problem.

Given the solution to this optimization problem, we can construct a modulating process with candidate schedules \mathbf{W}^p and limiting distribution is π^p by solving a system of $NP + 1$ linear equations.

Both Theorem 2 and Corollary 3 make the strong assumption that all jobs can be preempted with no overhead at any time. Fortunately, we can generalize both results to the case where jobs are non-preemptible and the case where preemptions incur setup times. These generalizations are proven as Theorems 4 and 5, respectively. Both theorems show how a pMSR policy constructed using Corollary 3 can be adjusted to construct a stable policy under different preemption models, and thus both of these proofs are constructive.

Response Time Analysis Results. While Corollary 3 constructs a stable MSR policy, there may be many MSR policies that can stabilize the system. Our goal is therefore to analyze the mean response time of a given MSR policy, and use this analysis to choose an MSR policy with low mean response time.

Our analysis exploits the fact that the MSR policy's modulating process is independent of the system state. This allows us to decouple the system and individually analyze the number of type- i jobs in the system for each of the K job types. Specifically, we analyze the expected number of type- i jobs under an MSR policy p , $\mathbb{E}[Q_i^p]$, by comparing it to $\mathbb{E}[Q_i^{p-1}]$, the number of jobs in a single-server queue with Markov-modulated service rates. We refer to the single-server system as the MSR-1 system.

We first note that the difference between the mean queue lengths of these systems is at most the maximum number of type- i jobs served in parallel. We prove this formally in Theorem 6 via a coupling argument.

Theorem 6. *For any job type, i , consider an MSR policy p and its corresponding MSR-1 policy $p-1$. Let β_i^p denote the maximum number of type- i jobs served in parallel under p . Then,*

$$\mathbb{E}[Q_i^{p-1}] \leq \mathbb{E}[Q_i^p] \leq \mathbb{E}[Q_i^{p-1}] + \beta_i^p.$$

Next, we leverage the results of [16] to bound $\mathbb{E}[Q_i^{p-1}]$. Our bound comprises two terms. First, there is a primary term that scales as $\Theta\left(\frac{1}{1-\rho_i}\right)$, where ρ_i is the load on the decoupled type- i system. Second, there is a term corresponding to unused service — times when the type- i system has unused idle capacity. We combine this bound with Theorem 6 to prove a bound on $\mathbb{E}[Q_i^p]$ in Theorem 8 (see Section 6.3 for a full theorem statement along with the required notation). In addition to our upper and lower bounds on mean queue length, we develop a simple heuristic to approximate $\mathbb{E}[Q_i^p]$ that always falls between our bounds.

We use this analysis to guide our choice of MSR policies under various preemption models in Section 7. Ideally, we could select a policy by directly optimizing our response time bounds or approximations subject to the quadratic stability constraints proven in Corollary 3. However, our bounds and approximations include terms that must be computed numerically once the modulating process is fixed, so it is intractable to optimize over these expressions directly. We instead suggest using a simple heuristic to choose an MSR policy — to minimize the maximum value of ρ_i for any job type, i . This heuristic provides several benefits:

- The heuristic is linear. Hence, we can formulate a Mixed Integer Quadratic Constrained Programming (MIQCP) problem that optimizes the heuristic subject to the quadratic constraints that guarantee system stability. A solution to the MIQCP, as defined in equation (7) of Section 7.1, can be used to construct a stable pMSR policy, p .
- The pMSR policy we select performs well under heavy load. This holds because minimizing our heuristic also minimizes system load, ρ , and the primary term in our bound on $\mathbb{E}[Q^p]$ scales as $\Theta\left(\frac{1}{1-\rho}\right)$. Notably, we prove in Theorem 10 that in the heavy traffic limit as $\rho \rightarrow 1$, p is constant-competitive with MaxWeight.
- The heuristic is minimized by balancing the load experienced by each job type. Hence, while fairness is not a central concern of this paper, we show that an MSR policy can achieve low overall mean response time without arbitrarily hurting one type of jobs.

How to use our results in practice. Running a stable MSR policy with low mean response time consists of both an offline and an online step.

First, we initialize the policy by solving an MICQP, (7), to determine the candidate set and limiting distribution of the modulating process. Next, we solve a system of equations to construct a modulating process that meets this specification. While there are many possible structures we can use for the modulating process, a simple heuristic is to arrange the candidate schedules in a loop structure. The structure of the modulating process is discussed in Section 7. While this step is computationally expensive, it happens offline.

Second, the online portion of the MSR policy runs with complexity at most $O(K)$ by maintaining at most K exponential timers to track transitions of the modulating process. When using a modulating process with a loop structure, as suggested above, only 1 exponential timer is needed and the online portion of the MSR policy runs with complexity $O(1)$. Hence, MSR policies move the vast majority of their complexity offline, resulting in efficient online execution.

5 Markovian Service Rate (MSR) Policies

In this section, we will prove the class of MSR policies is *throughput-optimal* for scheduling multiresource jobs. In other words, if there exists a scheduling policy that can stabilize a given system, we can select an MSR policy that stabilizes the system. The results of this section only provide stability conditions for various MSR policies. In order to construct performant MSRs in practice, we will need to use the more detailed response time analysis developed in Section 6. We will then show in Section 7 how to use our stability conditions and response time analysis to construct performant MSR policies under a variety of preemption models.

5.1 Throughput-Optimality of MSR policies

Our high-level goal in this section is to show that the class of MSR policies is throughput-optimal. However, we will see that proving this claim will depend on the preemption behavior of jobs. Hence, we will prove this style of result for a variety of preemption behaviors in Theorems 2, 4 and 5.

We begin by first establishing necessary and sufficient stability criteria for any MSR policy. Intuitively, the system should be stable whenever the long-run average completion rate of an MSR policy exceeds the arrival rate for every job type. We prove this stability condition as Lemma 1.

Lemma 1 (Stability Condition). *Given a multiresource system with K job types, arrival vector λ and service requirement vector μ , the system is stable under an MSR policy p if and only if the long-run average completion rate for each job type exceeds the long run arrival rate. That is, the system is stable if and only if $\mu \otimes \mathbb{E}[u^p] > \lambda$, where \otimes denotes the Hadamard product between vectors.*

PROOF. This proof uses the Foster-Lyapunov criteria [34] to establish stability conditions for the system. We use the L2-norm of the queue length vector $Q(t)$ as our Lyapunov function and show that the required drift conditions are satisfied if and only if the average completion rate of jobs is greater than the average arrival rate. See Appendix A for details. \square

Lemma 1 applies to any MSR policy regardless of the preemption model being considered. Hence, we can use these stability criteria to reason about throughput-optimality under a variety of preemption behaviors in the following sections.

Fully Preemptible Jobs. Next, we show that Lemma 1 suffices to prove throughput-optimality when scheduling fully preemptible jobs. In this case, we assume that jobs can be preempted at arbitrary times, arbitrarily frequently, with no associated overhead. As such, we place no constraints on the modulating process. We call an MSR policy for scheduling fully preemptible multiresource jobs a *pMSR policy*. We prove the throughput-optimality of pMSR policies in Theorem 2.

Theorem 2. *When scheduling preemptible jobs with no preemption overhead, the class of MSR policies is throughput-optimal. That is, if there exists a scheduling policy that can stabilize a multiresource job system with K job types and arrival rates λ , then there exists an MSR policy, p , with $N^p \leq K$ candidate schedules that stabilizes the system. Specifically, an MSR policy can stabilize a system with preemptible jobs for any $\lambda \in C$.*

PROOF. As described in Section 3.6, the capacity region $\lambda \in C$ is known and is given by (2). Equation (2) tells us that $\lambda \otimes \mu$ must lie strictly on the interior of the convex hull of \mathcal{S} . Hence, for some $\epsilon' > 0$, $(1 + \epsilon')\lambda \otimes \mu$ lies on the border of the convex hull of \mathcal{S} , and this point can be written as a convex combination of schedules on the border of the convex hull of \mathcal{S} . By Carathéodory's Theorem [8], there exists at least one such convex combination consisting of at most K schedules. Our goal is to set the transitions of the modulating process so that the stationary probabilities of being in each state match the weights of the convex combination. Finding appropriate transition rates then amounts to solving a system of $K - 1$ linearly independent equations and $K(K - 1)$ unknowns. This guarantees that we can construct a pMSR policy, p , with $N^p \leq K$ such that $\mathbb{E}[u^p] = (1 + \epsilon')\lambda \otimes \mu$ for some $\epsilon' > 0$. By Lemma 1, $\{Q^p(t), t \geq 0\}$ is positive recurrent and the system is stable. \square

Crucially, Theorem 2 says that a stable MSR policy only needs K states in its modulating process, confirming that pMSR policies provide a low-complexity solution for scheduling multiresource jobs. At any moment in time, the only information the policy needs to maintain in order to stabilize the system are K states comprising the candidate set, the $K(K - 1)$ transition rates between these states, and the time and direction of the next transition. By contrast, policies like MaxWeight may consider an exponential number of schedules at every scheduling decision. The modulating process

for an MSR policy can also be chosen offline during a one-time optimization step that considers the arrival rates, service rates, and resource demands of each job type.

While Theorem 2 tells us that some stable pMSR policy exists whenever some stable policy exists, it does not tell us how to construct a stable pMSR policy. We provide steps to construct a stable pMSR policy, p , in Corollary 3.

Corollary 3. *For a pMSR policy p , let \mathbf{W}^p denote the matrix of candidate schedules,*

$$\mathbf{W}^p = \begin{pmatrix} u_{1,1}^p & u_{1,2}^p & \cdots & u_{1,K}^p \\ \cdots & \cdots & \cdots & \cdots \\ u_{NP,1}^p & u_{NP,2}^p & \cdots & u_{NP,K}^p \end{pmatrix} \in \mathbb{Z}_{NP \times K}^+.$$

Let $\boldsymbol{\pi}^p$ be the limiting distribution of the modulating process $m^p(t)$, $\boldsymbol{\pi}^p = (\pi_1^p, \pi_2^p, \dots, \pi_{NP}^p) \in \mathbb{R}^{NP}$. To construct a stable pMSR policy, it is necessary to find values \mathbf{W}^p and $\boldsymbol{\pi}^p$ such that

$$\begin{cases} \mathbf{W}^p \mathbf{D} \leq \mathbf{P} & \text{(Capacity constraint)} \\ \pi_i^p \geq 0, \forall i & \text{(Positive fractions)} \\ \pi_1^p + \pi_2^p + \cdots + \pi_{NP}^p \leq 1 & \text{(Fractions sum to 1)} \\ \boldsymbol{\pi}^p \mathbf{W}^p > \boldsymbol{\lambda} \odot \boldsymbol{\mu} & \text{(Stability condition from Lemma 1)} \\ \mathbf{W}^p \text{ integers} \end{cases}.$$

Because Theorem 2 guarantees that there exists an MSR policy, p , with $NP \leq K$ whenever it is possible to stabilize the system, we can safely assume that $NP = K$ when solving this optimization problem.

Given the solution to this optimization problem, we can construct a modulating process with candidate schedules \mathbf{W}^p and limiting distribution is $\boldsymbol{\pi}^p$ by solving a system of $NP + 1$ linear equations.

PROOF. The capacity constraint $\mathbf{W}^p \mathbf{D} \leq \mathbf{P}$ suggests all candidate schedules are in the schedulable set \mathcal{S} . Lemma 1 implies that if $\boldsymbol{\pi}^p \mathbf{W}^p > \boldsymbol{\lambda} \odot \boldsymbol{\mu}$, then the system is stable. By Theorem 2, the optimization problem is feasible. \square

While we provide constraints for the optimization problem in Corollary 3, we have no objective function here because the results of Lemma 1 did not tell us how to compare two stable pMSR policies. In Section 6, we develop analytical tools to optimize over the space of stable MSR policies.

Non-preemptible Jobs. We refer to an MSR policy for scheduling non-preemptible jobs as an *nMSR policy*. Unlike pMSRs, an nMSR policy's modulating process can no longer change states arbitrarily. Specifically, we must ensure that no preemptions are required for any state transitions.

Solely increasing the number of type- i jobs scheduled does not require preemptions, because this action simply reserves or reclaims some free resources. Decreasing the number of type- i jobs scheduled, however, generally requires waiting for some running type- i jobs to complete. We therefore limit an nMSR policy to decrease the number of scheduled jobs of any type by at most one each time the modulating process changes states. Because moving to a schedule with one fewer type- i job in service may require completing a type- i job, an nMSR policy q 's modulating process can transition with rate at most $\mu_i u_i^q(t)$ in this case. While these added restrictions clearly break our proof of Theorem 2, we will argue that the class of nMSR policies remains throughput-optimal.

The key to our argument is to design modulating processes with two types of states: *working states* and *switching states*. Working states correspond to schedules that actually help stabilize the system. Switching states exist solely to facilitate switching between working states without preemptions. It is straightforward to see that, for any pair of working states, there exists a finite sequence of switching states that permit transitions from one working state to the other without preemptions. For example, a sequence of switching states could first reduce the number of type-1

jobs schedules one at a time, then reduce the number of type-2 jobs, etc. We define a *switching route* to be a particular sequence of switching states connecting two working states.

To prove the throughput optimality of the class of nMSR policies, we argue that any pMSR policy, p , that stabilizes the system can be augmented with some switching routes to create an nMSR policy that stabilizes the system. We prove this claim in Theorem 4.

Theorem 4. *When scheduling non-preemptible jobs, the class of nMSR policies is throughput-optimal. That is, if there exists a policy that stabilizes a multiresource job system with K job types and arrival rates λ , then there exists an nMSR policy, q , with $N^q \leq K$ working states that stabilizes the system. Specifically, an nMSR policy can stabilize a system with non-preemptible jobs when $\lambda \in C$.*

PROOF. Our proof first invokes Theorem 2 to claim the existence of a pMSR policy, p , with $N^p \leq K$ that stabilizes the system. When scheduling non-preemptive jobs, we would like our stationary distribution to be close to that of p . To accomplish this, we let the working states of q equal the states of m^p and assume arbitrary switching routes. We then fix the probability that q is in each working state *given* that the modulating process in some working state to equal the stationary distribution of m^p . By scaling all the transition rates out of the working states proportionally to some switching rate, α , we control the amount of time q spends in the working states without changing our chosen conditional probabilities. We then use a Renewal-Reward argument to show that setting α sufficiently low makes the fraction of time spent in any of the working states arbitrarily close to 1. This implies $\mathbb{E}[u^q] \approx \mathbb{E}[u^p]$. The full proof of this claim is given in Appendix E. \square

Our proof of Theorem 4 demonstrates the value of preemption. In order to stabilize the system, our nMSR policies must move through switching states whose corresponding schedules may be highly inefficient. To avoid spending too much time in these inefficient switching states, an nMSR policy's modulating process must change between working states on a much slower timescale than a pMSR policy, which can directly move between efficient working states. We will measure the impact of this slower modulation both in theory and in simulation in Sections 6 and 7, respectively.

Preemption with Setup Times. Our model of MSR policies is even general enough to capture the case where jobs can be preempted, but where preemptions incur a *setup time*. Here, a setup time represents a period of time during preemptions where server resources are occupied but not actively used by any job. Setup times generally occur in real systems when the job that is being preempted must be checkpointed and its resources must be reclaimed before the preempting job can start running. To model this process, we associate every preemption with a i.i.d. drawn random setup time of length $\Gamma \sim \exp(\gamma)$. We refer to γ as the *setup rate* of the system. When a job is preempted, it holds its demanded resources for Γ seconds before the preempting job is put in service. Neither the preempted job nor the preempting job can complete during this setup time.

We call an MSR policy for scheduling jobs with setup times an *sMSR policy*. The class of sMSR policies shares the challenges with nMSR policies — the modulating process must be constrained to model preemption overheads. Hence, we can again divide an sMSR policy's modulating process into working and switching states. In sMSRs, the transition rates between switching states will depend on the rate γ rather than service rates. Furthermore, the schedules corresponding to each switching state must be altered for jobs that are experiencing setup times but not actually running. While we will delve more deeply into the performance of sMSR policies in Section 7, we will begin here by proving the throughput optimality of sMSR policies in Theorem 5.

Theorem 5. *When scheduling jobs with setup costs, the class of sMSR policies is throughput-optimal. That is, if there exists a scheduling policy that can stabilize a multiresource job system with K job types and arrival rates λ , then there exists an sMSR policy, r , with $N^r \leq K$ working states that stabilizes the system. Specifically, an sMSR policy can stabilize a system with preemptible jobs when $\lambda \in C$.*

PROOF. This proof is nearly identical to the proof of Theorem 4. While the exact structures of the switching routes between working states differ between the nMSR and sMSR cases, in both cases we choose working states and transitions in order to mirror the behavior of some pMSR policy, p , that stabilizes the system. We use the same argument with the switching rate, α , to show that we can construct an sMSR policy based on p that also stabilizes the system. \square

Taken together, the throughput optimality theorems of this section show that there is a large design space of MSR policies that can stabilize a multiresource job system. However, our current tools for analyzing these MSR policies can only tell us whether the system is stable. To choose a particular MSR policy that achieves low mean response time, we need to develop a stronger analysis of the performance of an MSR policy. We develop this analytic framework in Section 6.

6 Response Time Bounds for MSR Algorithms

Although we have shown that MSR policies are useful for stabilizing the system, it is still unclear how to select an MSR policy that achieves low mean response time. Hence, in this section, we prove bounds on the mean response of an arbitrary MSR policy p . We will also show how these bounds can be used to provide a closed-form prediction of mean response time under an MSR policy that is highly accurate in practice. Our prediction formula can then be used to guide the selection of a performant MSR policy. Specifically, we will use our response time analysis to select a stable MSR policy with low mean response time under a variety of job preemption models in Section 7.

Our analysis begins by decoupling the original system with K job types into K separate systems. We will analyze the behavior of these decoupled systems one at a time by comparing the multiresource job system to a simpler single-job-at-a-time system with modulated service rates. This comparison allows us to prove bounds on the mean queue length for each type of jobs in Theorem 8. We then use our bounds to derive an accurate mean queue length *approximation* in Section 6.4.

6.1 Decoupling the Multiresource Job System

We begin by noting that there are two equivalent perspectives from which one can view the multiresource job system. Perspective 1, as described in Section 3, is that there is a central queue holding multiresource jobs of all types. Perspective 2 is that we can imagine K separate queues, one for each type of jobs. In general, the service rate of the type- i queue can depend on the numbers of jobs in all other queues, so the queues in perspective 2 do not evolve independently, limiting the usefulness of this perspective.

Under an MSR policy, however, Definition 1 tells us that the number of type- i jobs scheduled by an MSR policy p depends only on the state of the modulating process. As a result, when considering MSR policies, it is simpler to analyze the expected length of each queue in perspective 2 than it is to analyze the central queue in perspective 1. Specifically, consider the type- i queue from perspective 2 under an MSR policy p . This queue has Poisson arrivals with rate λ_i and jobs whose service requirements are distributed as $\exp(\mu_i)$. The number of jobs allowed in service at time t evolves according to the CTMC $\{u_i^p(t)\}$, which is assumed to be independent of the overall state of the system. Hence, we call this the *decoupled system* for type- i jobs. In this section, we will bound the mean queue length in each decoupled system, $\mathbb{E}[Q_i]$. See Figure 1 for an example comparing perspective 1 and perspective 2 under an MSR policy.

To bound $\mathbb{E}[Q_i]$, we will compare the decoupled system to a similar, simpler system, the $M/M/1$ Markovian Service Rate system (MSR-1) system. A *corresponding MSR-1 system* for type- i jobs is a single server system that serves jobs one at a time and where the service rate of the server varies according to a finite-state CTMC. Hence, we can construct a corresponding MSR-1 system that looks highly similar to our decoupled system, except that the decoupled system runs multiple type- i

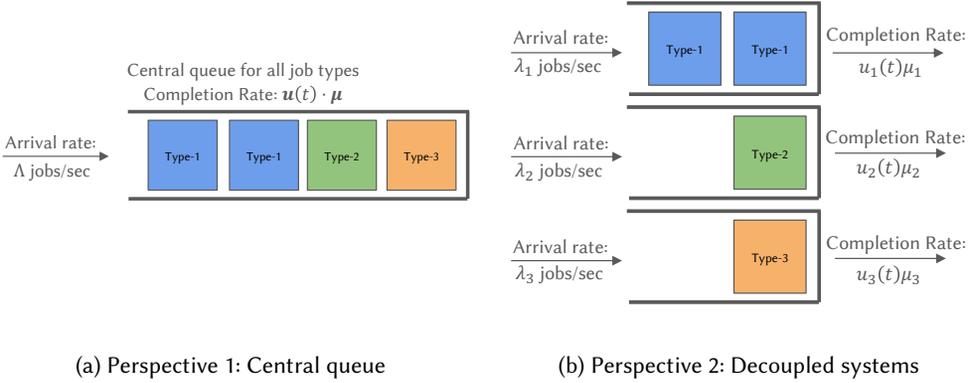


Fig. 1. Two views of the multiresource job system under an MSR policy. While arriving jobs are stored in a central queue, MSR policies allow us to analyze jobs of each type separately.

jobs in parallel. We can then analyze the MSR-1 system, and use our analysis to derive bounds on the decoupled system for each job type. Formally, we define the MSR-1 system as follows.

Definition 2. A corresponding MSR-1 system for type- i jobs is an $M/M/1$ queueing system with Markov-modulated service rates. Let λ_i be the arrival rate of the system. The corresponding MSR-1 system shares the same modulating process, $\{m^p(t), t \geq 0\}$ as the original MSR system. We denote this policy derived from p for the MSR-1 system as $p-1$. Hence, at time t , the corresponding MSR-1 system completes jobs at a rate of $u_i^{p-1}(t) = \mu_i u_i^p(t)$. Following the definition, $u_{s,i}^{p-1} = \mu_i u_{s,i}^p$.

Let $\{Q_i^{p-1}(t), t \geq 0\}$ denote the queue length process of the corresponding MSR-1 system under $p-1$. Our goal is to analyze the mean queue length under $p-1$ and then use this analysis to provide queue length bounds for the original system under MSR policy p .

The intuition behind our argument is that, when there are sufficiently many jobs in the system, the MSR and MSR-1 systems complete type- i jobs at the same rate. When there are not many jobs in the MSR system, the MSR system may complete jobs at a slower rate. However, we know the queue length of the MSR system is small in this case. We formalize this argument in Theorem 6.

Theorem 6. For any job type, i , consider an MSR policy p and its corresponding MSR-1 policy $p-1$. Let β_i^p denote the maximum number of type- i jobs served in parallel under p . Then,

$$\mathbb{E}[Q_i^{p-1}] \leq \mathbb{E}[Q_i^p] \leq \mathbb{E}[Q_i^{p-1}] + \beta_i^p.$$

PROOF. We prove this claim via a coupling argument in Appendix B. □

Theorem 6 tells us that the MSR and MSR-1 systems perform similarly (within an additive constant) at all system loads. It is worth noting, however, that these bounds become asymptotically tight as the arrival rates in both systems increase and the mean queue length of the MSR-1 system goes to infinity. This makes intuitive sense — our proof argues that the only differences between the MSR and MSR-1 systems happen when there are few jobs in one of the two systems, which becomes increasingly rare as the arrival rates increase. We now proceed to leverage our bounds by analyzing the MSR-1 system in order to bound mean queue length in the MSR system.

6.2 Analysis of the MSR-1 System

We now analyze an MSR-1 system running under policy $p-1$. Recently, [16] bounded the mean queue length of an MSR-1 system by using the concept of *relative completions*. The concept was

an independent rediscovery of a technique which had previously been introduced by [9] for the Markovian arrival settings, and extended to the Markovian service setting by [6], proving equivalent results for mean queue length. A recent tech report, [15], provides a helpful expanded presentation of the technique of [16], and a more thorough review of the state of the literature. For any state s in the modulating process of $p-1$, the relative completions of the state is defined as the long-run difference in completions between the MSR-1 system that starts from state s and the MSR-1 system that starts from a random state chosen according to the stationary distribution of $\{m^p(t)\}$. Formally, let $C_i^{p-1}(t)$ denote the cumulative type- i completions under $p-1$ by time t . We define $\Delta(s)$ to be the relative completions of state s in the modulating process where

$$\Delta(s) = \lim_{t \rightarrow \infty} C_i^{p-1}(t \mid m^p(0) = s) - \mathbb{E}[u_i^{p-1}]t.$$

We also define v_i^{p-1} to be the completion-rate-weighted stationary random variable for the state in $\{v_i^{p-1}(t)\}$, which is defined such that

$$P\{v_i^{p-1} = s\} = P\{m^p = s\} \frac{u_{s,i}^{p-1}}{\mathbb{E}[u_i^{p-1}]}.$$

A key result from [15, 16] can then be stated as follows.

Remark 7. In [15, 16], it is shown that

$$\mathbb{E}[Q_i^{p-1}] = \frac{\rho_i + \mathbb{E}[\Delta(v_i^{p-1})]}{1 - \rho_i} - \mathbb{E}_U[\Delta(v_i^{p-1})],$$

where $\rho_i = \frac{\lambda_i}{\mathbb{E}[u_i^{p-1}]}$ and $\mathbb{E}_U[\Delta(v_i^{p-1})]$ denotes the expectation of $\Delta(v_i^{p-1})$ taken over moments when the system experiences unused service.

6.3 Queue Length Bounds for the MSR System

Combining Remark 7 with Theorem 6 yields the following bound for the MSR system.

Theorem 8. For any multiresource job system under an MSR policy p , the mean queue length of type- i jobs is bounded in terms of the corresponding MSR-1 system under policy $p-1$ as

$$\mathbb{E}[Q_i^p] = \frac{\rho_i + \mathbb{E}[\Delta(v_i^{p-1})]}{1 - \rho_i} + B,$$

where

$$- \max_s \Delta(s) \leq B \leq - \min_s \Delta(s) + \beta_i^p.$$

and β_i^p is the maximum number of type- i jobs served in parallel under p .

PROOF. First, we invoke Theorem 6 to relate the mean type- i queue length in the MSR system to a corresponding MSR-1 system, giving

$$\mathbb{E}[Q_i^{p-1}] \leq \mathbb{E}[Q_i^p] \leq \mathbb{E}[Q_i^{p-1}] + \beta_i^p.$$

We then apply Remark 7 to this bound to obtain

$$\frac{\rho_i + \mathbb{E}[\Delta(v_i^{p-1})]}{1 - \rho_i} - \mathbb{E}_U[\Delta(v_i^{p-1})] \leq \mathbb{E}[Q_i^p] \leq \frac{\rho_i + \mathbb{E}[\Delta(v_i^{p-1})]}{1 - \rho_i} - \mathbb{E}_U[\Delta(v_i^{p-1})] + \beta_i^p$$

Finally, we note that

$$\min_s \Delta(s) \leq \mathbb{E}_U[\Delta(v_i^{p-1})] \leq \max_s \Delta(s),$$

which gives the desired bounds. \square

To compute the bounds from Theorem 8 for a given system, one must compute both the distribution of v_i^{p-1} and the extrema of Δ . Although these quantities do not have a general closed-form, computing the quantities for a given system is straightforward.

First, computing the distribution of v_i^{p-1} amounts to computing the stationary distribution of the modulating process $m^p(t)$ and weighting each term by its completion rate of type- i jobs. From Theorem 2, we know that a modulating process needs only K states in order to stabilize the system, so this stationary distribution can generally be computed quickly using standard techniques.

Next, computing the minimum and maximum of the Δ function requires computing $\Delta(s)$ for each state s . Let $\Delta = (\Delta(1), \Delta(2), \dots, \Delta(N^p))^T$. We compute Δ as follows.

Lemma 9. *In an MSR-1 system under policy $p-1$ whose modulating process has the infinitesimal generator \mathbf{G}^{p-1} , Δ can be computed as the solution to*

$$\begin{cases} \mathbf{G}^p \cdot \Delta = \mathbb{E}[u_i^{p-1}] \mathbf{1}_K - (u_{1,i}^{p-1}, u_{2,i}^{p-1}, \dots, u_{N^{p-1},i}^{p-1}) \\ \mathbb{E}[\Delta(m^p)] = 0 \end{cases}.$$

PROOF. We begin with relating the value of $\Delta(s)$ to the Δ -value of its neighboring states:

$$\Delta(s) = \frac{u_{s,i}^{p-1} - \mathbb{E}[u_i^{p-1}]}{\sum_{s' \neq s} \mathbf{G}^p_{s,s'}} + \frac{\sum_{s' \neq s} \mathbf{G}^p_{s,s'} \Delta(s')}{\sum_{s' \neq s} \mathbf{G}^p_{s,s'}}. \quad (4)$$

Equation 4 is composed of two terms: the relative completions that accrue before the first transition out of state s , and the relative completions that accrue after the process leaves s .

Multiplying $\sum_{s' \neq s} \mathbf{G}^p_{s,s'}$ on both sides, we get:

$$\begin{aligned} \Delta(s) \sum_{s' \neq s} \mathbf{G}^p_{s,s'} &= u_{s,i}^{p-1} - \mathbb{E}[u_i^{p-1}] + \sum_{s' \neq s} \mathbf{G}^p_{s,s'} \Delta(s') \\ \sum_{s' \neq s} \mathbf{G}^p_{s,s'} \Delta(s') - \Delta(s) \sum_{s' \neq s} \mathbf{G}^p_{s,s'} &= \mathbb{E}[u_i^{p-1}] - u_{s,i}^{p-1} \\ \sum_{s'} \mathbf{G}^p_{s,s'} \Delta(s') &= \mathbb{E}[u_i^{p-1}] - u_{s,i}^{p-1} \end{aligned} \quad (5)$$

$$\mathbf{G}^p \cdot \Delta = \mathbb{E}[u_i^{p-1}] \mathbf{1}_{N^p} - \mu_i (u_{1,i}^{p-1}, u_{2,i}^{p-1}, \dots, u_{N^{p-1},i}^{p-1}) \quad (6)$$

Step (6) holds because (5) holds for all $s \in [1, N^p]$. See [16] for proof that $\mathbb{E}[\Delta(m^p)] = 0$. \square

Lemma 9 tells us that computing all Δ terms for an MSR-1 policy $p-1$ amounts to inverting the policy's generator matrix \mathbf{G}^p . Again, note that this will generally be a computationally inexpensive operation due to the limited number of candidate schedules needed by our MSR policies.

6.4 A Queue Length Approximation for the MSR System

Intuitively, the upper bound from Theorem 8 will be looser at lower loads, since Theorem 6 is most accurate at high loads. Specifically, to obtain a bound in Theorem 8, we assumed that the MSR system rarely experiences unused service. To better model the effects of unused service on the MSR system, we propose the following queue length approximation.

$$\mathbb{E}[Q_i^p] \approx P_Q^{M/M/k_i^*} \frac{\rho_i + \rho_i \mathbb{E}[\Delta(v_i^{p-1})]}{1 - \rho_i} + \rho_i \cdot k_i^*.$$

Here, $k_i^* = \mathbb{E}[u_i^p]$ and $P_Q^{M/M/k_i^*}$ denotes the queueing probability in an $M/M/k_i^*$ system. We derive this approximation by estimating the relationship between the MSR and MSR-1 systems and

applying this estimate instead of Theorem 6 in the proof of Theorem 8. The details of this argument are provided in Appendix C.

7 Choosing a Stable MSR Policy with Low Response Time

Section 6 established results on how to stabilize the system using an MSR policy. However, these results do not describe how to choose between the many possible MSR policies that might stabilize the system. In this section, we use our response time analysis from Section 6 to optimize the design of an MSR policy. We will consider how to select a good pMSR policy (Section 7.1), how to select a good nMSR policy (Section 7.2), and how to select a good sMSR policy (Section 7.3).

To illustrate the effectiveness of our optimization techniques, we will compare the MSR policies we design to the following three state-of-the-art scheduling policies:

MaxWeight is a preemptive policy proposed in [25] that chooses a schedule by solving a bin-packing problem on every arrival or departure. MaxWeight tries to maximize the weighted number of jobs in service, where the weight of type- i jobs equals to the queue length of type- i jobs.

Randomized-Timers is a non-preemptive policy proposed in [11] that uses moments of when jobs complete to change schedules. The policy uses a complex randomized algorithm to discover good solutions to the weighted bin-packing problem used by MaxWeight while the policy runs.

First-Fit is a simple, non-preemptive policy similar to FCFS. Sometimes known as BackFilling [14], First-Fit examines the queued jobs in FCFS order upon each job arrival or completion. If a job is found that can be added to the set of running jobs, First-Fit serves this job. Running jobs are never preempted. First-Fit is generally thought to be a better version of the FCFS policy from Section 1.

We will compare our MSR policies using an example of a server with multiresource jobs. Specifically, we consider a server with three resources — 20 cores, 15 GB of DRAM and 50 Gbps of network bandwidth. We consider a workload with three job types, whose demand vectors are given as

$$\mathbf{D}_1 = (3, 7, 1) \quad \mathbf{D}_2 = (4, 1, 1) \quad \mathbf{D}_3 = (10, 1, 5).$$

For this example, we vary ρ by linearly scaling the arrival rate vector, $\boldsymbol{\lambda}$. See Appendix F for further details. While this example is illustrative, we evaluate a more realistic workload in Section 8.

7.1 Choosing a pMSR Policy (Preemptive MSR)

Corollary 3 provides an Mixed Integer Quadratic Constrained Program (MIQCP) whose solutions correspond to the set of stable pMSR policies. Our goal is to augment this MIQCP with an objective function to help optimize our choice of pMSR policy. One approach would be to use either the bounds of Theorem 8 or our queue length approximation functions directly as the objective functions. Unfortunately, these functions lack a general closed form, and would make the optimization problem intractable. We instead propose a simpler objective function to heuristically minimize our queue length upper bounds.

We select a pMSR policy using the following two-step approach. First, we solve the following MIQCP for the candidate schedules \mathbf{W}^p and the limiting distribution $\boldsymbol{\pi}^p$:

$$\left\{ \begin{array}{l} \min \quad \max_i \rho_i \quad (\text{where } \rho_i = \lambda_i / (\boldsymbol{\pi}^p \mathbf{W}^p)_i) \\ \text{s.t.} \quad \mathbf{W}^p \mathbf{D} \leq \mathbf{P} \quad (\text{Capacity constraint}) \\ \quad \pi_i^p \geq 0, \forall i \quad (\text{Positive fractions}) \\ \quad \pi_1^p + \pi_2^p + \dots + \pi_{N^p}^p \leq 1 \quad (\text{Fractions sum to 1}) \\ \quad \boldsymbol{\pi}^p \mathbf{W}^p > \boldsymbol{\lambda} \oslash \boldsymbol{\mu} \quad (\text{Stability condition from Lemma 1}) \\ \quad \mathbf{W}^p \text{ integers} \end{array} \right. \quad (7)$$

Appendix J gives a numerical example of equation (7).

Second, we compute a modulating process with a simple loop structure that matches the parameters computed in (7). Here, we sort the set of candidate schedules lexicographically and let each state transition only to the next neighboring state in this ordering until all states have been visited. Given this loop structure and the desired values of π^p , we can compute the necessary transition rates by solving a system of N^p equations with $N^p + 1$ unknowns. Solving this system defines the ratios of the transition rates relative to one another, but still allows all transition rates to be scaled by a factor $\alpha > 0$. We refer to α as the *switching rate* of the system because it controls the overall rate of preemptions in the pMSR system. The loop-structured modulating process with switching rate α for our example system is shown in the Appendix in Figure 6a.

Our bounds suggest that the mean response time of a pMSR policy is minimized by setting α to be arbitrarily high. Intuitively, the system in this case begins to behave like a non-modulated system using a constant schedule of $\mathbb{E}[\mathbf{u}^p]$. We can actually prove this property using our mean queue length bounds for small pMSR instances with two job types (see Appendix D). Figure 2a confirms this result by showing how the mean queue length under a pMSR policy improves as α grows. Notably, although our theoretical results show that taking $\alpha \rightarrow \infty$ is beneficial in this example, Figure 2a shows that we capture most of the performance benefit as long as $\alpha \geq 2$. Hence, we will assume that α is some constant chosen to be sufficiently large.

The above process for selecting a pMSR policy has both intuitive and theoretical justifications. Intuitively, the ρ_i terms in our bounds denote the load on the system experienced by type- i jobs. Our bounds show that $\mathbb{E}[Q_i]$ scales as $\Theta(\frac{1}{1-\rho_i})$. Hence, it is intuitive to use an objective function in (7) that balances load across job types. Formally, we can prove that the policy we select compares favorably to MaxWeight in heavy traffic. Specifically, we show that the pMSR policy we select is constant-competitive with MaxWeight in the limit as $\rho \rightarrow 1$. Our proof relies on the fact that both the pMSR and MaxWeight queue lengths scale as $\Theta(\frac{1}{1-\rho})$ in heavy traffic. Hence, while the second step of our selection process uses a simple loop structure for the modulating process, further optimization of the modulating process only gives constant factor improvement in mean queue length relative to MaxWeight in heavy traffic.

We formally define our notion of heavy-traffic scaling in Definition 3 in Appendix H. Using this definition, we prove Theorem 10 as follows.

Theorem 10. *Under the heavy-traffic scaling regime in Definition 3, the solution to (7) is the same at all loads. The policy, p^* , that we select using (7) is constant-competitive with MaxWeight in heavy traffic. That is, there exists a constant $C > 0$ such that*

$$\lim_{\rho \rightarrow 1} \frac{\mathbb{E}[Q^p]}{\mathbb{E}[Q^{MaxWeight}]} \leq C$$

PROOF. See Appendix H for a full proof. □

7.2 Choosing an nMSR Policy (Nonpreemptive MSR)

Given the choice of pMSR policy in Section 7.1, the proof of Theorem 4 yields a related stable nMSR policy. Specifically, our proof suggests setting the working states to be the states of $m^p(t)$ and then using arbitrary switching routes between these states. To stabilize the system, we can scale the transition rates out of the working states by a very small switching rate $\alpha > 0$. This method for choosing an nMSR policy has two main issues. First, switching routes are chosen arbitrarily instead of being chosen to optimize mean response time. Second, while α must be small to guarantee stability, the optimal switching rate α^* is not obvious here.

The space of possible switching routes between working states is massive. Because our queue length bounds must be recomputed for each choice of switching routes, it is hard to optimize this

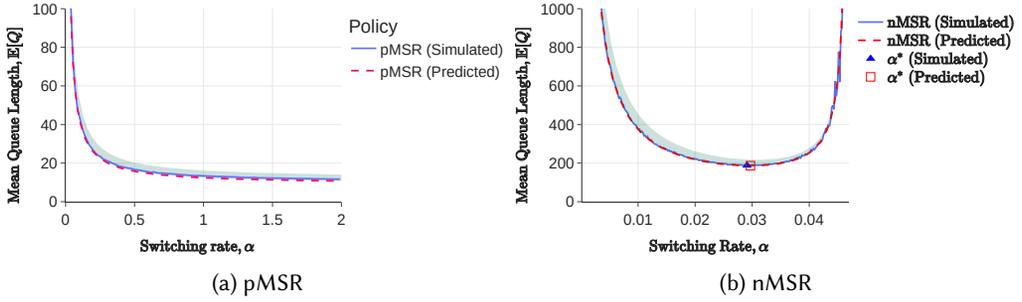


Fig. 2. The effect of switching rate on MSR policies under various preemption behaviors in the example from Section 7. The shaded regions depict our queue length bounds. Our pMSR policies benefit from a high switching rate when system load $\rho = 0.9$. However, nMSR policies suffer when α is too high or too low. Our queue length prediction is accurate and can be used to select α^* in both cases.

choice using our results. However, we sampled many possible switching routes for the example system in this section and found that the choice of switching routes had a minimal impact on mean queue length. Hence, we defer a full analysis of optimizing switching routes to future work. An example of an nMSR policy with simple switching routes is shown in Figure 6b of Appendix F.

Our queue length bounds *are* useful for optimizing the choice of α . Given the pMSR policy, p , selected in Section 7.1, let $q(\alpha)$ be the nMSR policy based on p whose transitions out of working states are scaled by $\alpha > 0$. Because it is non-preemptive, $q(\alpha)$ can suffer when α is high and the policy spends too much time in switching states. Conversely, if α is too small, the relative completions terms in our queue length analysis can become large. We therefore use our bounds to select the optimal switching rate, α^* , that balances the cost of spending time in switching states with the cost of switching too slowly between working states.

Figure 2b shows that our mean queue length predictions allow us to accurately predict α^* , and that choosing the correct value of α dramatically reduces mean queue length. Specifically, even though we chose α to minimize our queue length prediction, the value that minimizes this prediction (α^* **Predicted**) is within 2.1% of the α^* derived in simulation (α^* **Simulated**). The result of using α^* **Predicted** is a mean queue length that is within 1.5% of optimal.

Furthermore, we can bound the potential error from minimizing our prediction function instead of an exact formula for $\mathbb{E}[Q]$. The $\mathbb{E}[Q]$ from using α^* **Predicted** is at most the minimum value of our upper bound over all α 's. Similarly, the $\mathbb{E}[Q]$ under α^* **Simulated** is at least the minimum value of our lower bound over all α 's. In our example, the difference between the minima of our upper and lower bounds is 16.3%. Hence, while our prediction function is more accurate in practice, we can *guarantee* that using α^* **Predicted** is within 16.3% of the nMSR policy with the optimal α . By contrast, Figure 2b shows that choosing the wrong α can easily increase $\mathbb{E}[Q]$ by an order of magnitude or even result in instability. We can use our $\mathbb{E}[Q]$ bounds similarly to provide error bounds whenever optimizing parameters of an MSR policy using our prediction function.

7.3 Choosing an sMSR Policy (MSR with Setup)

To select an sMSR policy, r , we closely follow our selection process for the nMSR policy, q . The central difference between these policies is the structure of their switching routes, which differ in two key ways. First, the rates that the policies move through switching states will differ because q switches by waiting for completions with rate μ_i and r switches by waiting for setup times with setup rate γ . Second, the nMSR policy completes jobs during its switching states, but the time r

spends in its switching states is wasted because it is unable to complete the jobs being preempted. Nonetheless, again consider a range of sMSR policies, $r(\alpha)$, whose working state transitions are scaled by α , using our analysis to select an optimal switching rate α^* .

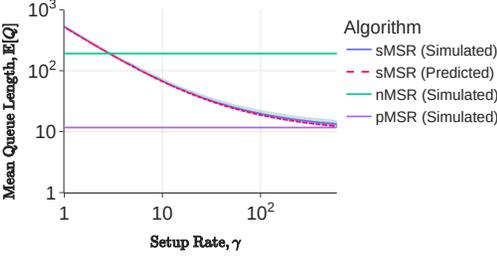


Fig. 3. Mean queue length of the sMSR policy $r(\alpha^*)$ as a function of setup rate, γ when $\rho = 0.9$. The shaded regions depict our queue length bounds for $r(\alpha^*)$. We compare the performance of $r(\alpha^*)$ to the nMSR policy, $q(\alpha^*)$, and the pMSR policy, p .

Because one can always use an nMSR policy and avoid costly preemptions altogether, it is natural to compare our sMSR policy, $r(\alpha^*)$, to the nMSR policy $q(\alpha^*)$. Specifically, one must understand how fast the setup times must be in order for $r(\alpha^*)$ to outperform $q(\alpha^*)$.

Figure 3 shows that our mean queue length bounds accurately predict when it is better to use an sMSR policy than an nMSR policy. The $r(\alpha^*)$ policy performs better than $q(\alpha^*)$ when γ is significantly higher than μ_i for all i , as expected. As γ grows and preemption overhead decreases, we see the performance of $r(\alpha^*)$ approach the performance of p . In these experiments, we used our mean queue length bounds to select a different α^* for each policy and for each value of γ .

7.4 The Example System

Figure 4 summarizes our results from this section by measuring the performance of our optimized MSR policies and the competitor policies as a function of system load.

Figure 4a shows our pMSR policy, p , is close to the performance of MaxWeight. While MaxWeight beats p at moderate loads, they appear to converge as $\rho \rightarrow 1$. First-Fit performs well at $\rho < .9$ but becomes unstable around $\rho = .9$. The non-preemptive Randomized-Timers lags far behind p .

We also evaluate the nMSR policy, $q(\alpha^*)$, in Figures 4b and 4c. MaxWeight greatly outperforms $q(\alpha^*)$ at all loads because it is preemptive. Furthermore, $q(\alpha^*)$ is not the best of the non-preemptive policies at low loads. However, $q(\alpha^*)$ becomes the best non-preemptive policy under high load. Here, First-Fit becomes unstable, while Randomized-Timers suffers from changing between schedules too slowly. In fact, Randomized-Timers performed so poorly under high load that many of our simulations did not converge (these points were omitted).

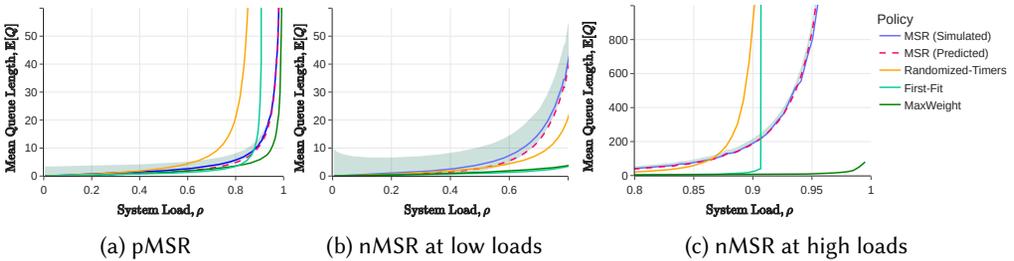


Fig. 4. Mean queue length under pMSR and nMSR policies in the example from Section 7. The pMSR policy uses $\alpha = 2$. We use α^* for the nMSR policy at each load. The shaded regions depict the queue length bounds for our MSR policies. The dashed lines represent our queue length approximations.

8 Evaluation

Section 7 shows how to optimize the performance of MSR policies under our theoretical model. This raises the question of how our policies perform under a real-world workload that violates several of our modeling assumptions (see Section 3.7). To address this question, we performed simulations based on publicly available data from the Google Borg system [36]¹.

8.1 Experimental Setup

Using the data from [36], we constructed a trace of 69 million requests that were submitted to Borg over a 30-day period in 2019. The trace contains ten job types, and jobs demand both CPU cores and memory. Job sizes and interarrival times follow highly variable distributions that are far from our theoretical model. For details on the construction and features of this trace, see Appendix I.

To simulate different arrival rates, we downsampled the trace by various downsampling factors. Following the notation of Section 3, we refer to our downsampling factor as *system load*, ρ , and normalize this factor so a value of $\rho < 1$ implies that at least one policy we tested was stable.

Our goal is to compare the mean response time, $\mathbb{E}[T]$, under different scheduling policies as a function of ρ . The competitor policies we consider follow the policy definitions of Section 7. Specifically, we once again consider the MaxWeight, Randomized-Timers, and First-Fit policies. For the MSR policy variants, we derive the set of working states by solving the optimization problem in (7). We again determine transition rates by assuming a loop-structured modulating process with an overall switching rate of α . For the pMSR policies we consider, we set $\alpha = 0.1$. For the nMSR policies, we set $\alpha = 3.6 \times 10^{-6}$, a value chosen to ensure stability when $\rho = .9$.

8.2 A Practical Optimization: MSR With BackFilling

All of our MSR policy variants suffer during moments of unused service. Here, resources earmarked by the modulating process to serve type- i jobs are left idle when there are not enough type- i jobs in the system. This wasted service has two impacts on our response time bounds. First, there are the additive terms in our bounds that correspond directly to the unused service for each job type. Reducing the amount of unused service will reduce these terms. Second, unused service also imposes an opportunity cost on the system — the wasted capacity for class i jobs could have been used to serve some other type of jobs, j , reducing the $\Theta\left(\frac{1}{1-\rho_j}\right)$ term in our bound. Hence, particularly as system load increases, reducing unused service even slightly can have an outside benefit on mean response time by reducing the load experienced by other job types.

We therefore simulate variants of our MSR policies that include a practical optimization for filling unused service capacity in the system. If an MSR policy tries to serve a class i job but does not find one available, we will try to fill the excess capacity by serving additional jobs from the central queue according to the First-Fit policy. We refer to these MSR variants as *MSR with BackFilling* policies. When a pMSR with BackFilling policy changes working states, it can simply preempt any additional jobs that were put into service by BackFilling, thus guaranteeing that all jobs will finish no later under pMSR with BackFilling as compared to the original pMSR policy. Therefore, the throughput-optimality results and response time upper bounds for pMSR policies extend to the analogous pMSR with BackFilling policies. When an nMSR or sMSR with BackFilling policy changes working states, however, additional jobs may need to be preempted or completed in order to change states, making it hard to guarantee a performance benefit from BackFilling. Nonetheless, our simulations show that even with non-preemptible jobs the benefits of BackFilling outweigh the

¹Code is available online at <https://github.com/jcpwflor/msr-borg>

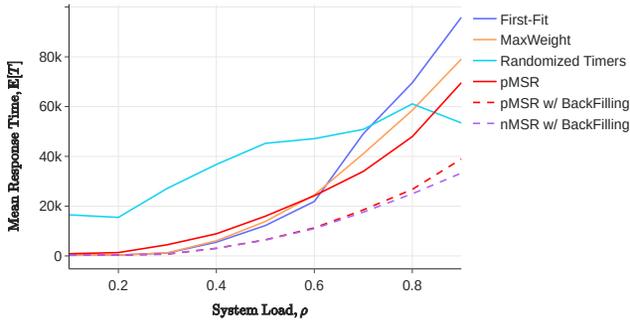


Fig. 5. Evaluation of MSR policies using a Google Borg trace. The nMSR policy without BackFilling is not pictured, but performed roughly 10x worse than the policies shown due to excessive unused service. The pMSR and nMSR policies with BackFilling provide low mean response times at all loads.

extra time the system spends in switching states. A more rigorous analysis of MSR policies with BackFilling will be addressed in future work.

8.3 Experimental Results

Our simulation results comparing scheduling policies from the literature to a variety of MSR policies using the Google Borg trace are shown in Figure 5. Notably, the nMSR policy we tested is not pictured because it performs roughly 10x worse than the policies in Figure 5 due to its high degree of unused service. While First-Fit and MaxWeight perform well at low loads, they are beaten by the pMSR policy under moderate loads. The pMSR policy outperforms First-Fit and MaxWeight by 40% and 15% respectively when $\rho = .8$. Notably, the MSR with BackFilling variants of the pMSR and nMSR policies provide excellent mean response times, outperforming both First-Fit and MaxWeight by a factor of 2 when load is high. While Randomized-Timers lags behind the MSR with BackFilling policies, it begins to be competitive at higher loads. The non-monotonicity of the Randomized-Timers data is caused by the difficulty of tuning various policy parameters.

It is worth noting that the optimization code used to compute the candidate set for the MSR policies completed in just 15.28 seconds on a modest 28-core server using the Gurobi solver. Unlike MaxWeight, this optimization code is not on the critical path of the MSR scheduling policies, and thus could be continuously re-run in the background using modest computational resources.

9 Conclusion

This paper devises a class of low-complexity scheduling policies for multiresource jobs called MSR policies. An MSR policy selects a set of candidate schedules to switch between using a CTMC. The class of MSR policies is throughput-optimal despite its simplicity. We show how to select an MSR policy that achieves low mean response time in a variety of practical cases.

While our analysis of MSR policies yields accurate predictions of mean queue length, MSR policies have some limitations. Specifically, MSR policies are outperformed by MaxWeight and First-Fit at moderate loads. These competitor policies have complex, infinite state modulating processes that depend on which jobs are in the queue. Our current analysis does not apply to these complex cases, but a natural direction for future work is to design policies that incorporate some state-dependent behavior while remaining tractable. Specifically, we are encouraged by our simulations of MSR with BackFilling, and will aim to generalize our analysis to handle these policies.

Acknowledgments

We thank our shepherd, Dr. Amr Rizk, and the anonymous reviewers for their insightful feedback that helped improve our work. This work is supported by a Northwestern IEMS Startup Grant, AFOSR Grant FA9550-24-1-0002, a UNC Chapel Hill Startup Grant, and National Science Foundation grants NSF-CCF-2403195 and NSF-IIS-2322974.

References

- [1] B. Berg, J.P. Dorsman, and M. Harchol-Balter. 2018. Towards Optimality in Parallel Scheduling. *ACM POMACS* 1, 2 (2018).
- [2] David Y Burman and Donald R Smith. 1986. An asymptotic analysis of a queueing system with Markov-modulated arrivals. *Operations Research* 34, 1 (1986), 105–119.
- [3] Zhongrui Chen, Isaac Grosf, and Benjamin Berg. 2024. Simple Policies for Multiresource Job Scheduling. *ACM SIGMETRICS Performance Evaluation Review* 52, 2 (2024), 27–29.
- [4] Florin Ciucu and Felix Poloczek. 2018. Two extensions of Kingman’s GI/G/1 bound. *Proceedings of the ACM on Measurement and Analysis of Computing Systems* 2, 3 (2018), 1–33.
- [5] Christina Delimitrou and Christos Kozyrakis. 2014. Quasar: Resource-efficient and qos-aware cluster management. *ACM SIGPLAN Notices* 49, 4 (2014), 127–144.
- [6] Mitko Dimitrov. 2011. Single-server queueing system with Markov-modulated arrivals and service times. *Pliska Stud. Math. Bulg* 20 (2011), 53–62.
- [7] Atilla Eryilmaz and Rayadurgam Srikant. 2012. Asymptotically tight steady-state queue length bounds implied by drift conditions. *Queueing Systems* 72 (2012), 311–359.
- [8] Ruy Fabila-Monroy and Clemens Huemer. 2017. Caratheodory’s theorem in depth. *Discrete & Computational Geometry* 58 (2017), 51–66.
- [9] Gennadi Falin and Anatoli Falin. 1999. Heavy traffic analysis of M/G/1 type queueing systems with Markov-modulated arrivals. *Top* 7, 2 (1999), 279–291.
- [10] Anshul Gandhi, Sherwin Doroudi, Mor Harchol-Balter, and Alan Scheller-Wolf. 2013. Exact analysis of the M/M/k/setup class of Markov chains via recursive renewal reward. In *Proceedings of the ACM SIGMETRICS/international conference on Measurement and modeling of computer systems*. 153–166.
- [11] Javad Ghaderi. 2016. Randomized algorithms for scheduling VMs in the cloud. *Proceedings - IEEE INFOCOM 2016-July* (7 2016). <https://doi.org/10.1109/INFOCOM.2016.7524536>
- [12] Javad Ghaderi. 2016. Simple high-performance algorithms for scheduling jobs in the cloud. *2015 53rd Annual Allerton Conference on Communication, Control, and Computing, Allerton 2015* (4 2016), 345–352. <https://doi.org/10.1109/ALLERTON.2015.7447025>
- [13] Isaac Grosf, Mor Harchol-Balter, and Alan Scheller-Wolf. 2020. Stability for Two-class Multiserver-job Systems. (10 2020). <https://arxiv.org/abs/2010.00631v1>
- [14] Isaac Grosf, Mor Harchol-Balter, and Alan Scheller-Wolf. 2022. WCFS: a new framework for analyzing multiserver systems. *Queueing Systems* 102, 1-2 (10 2022), 143–174. <https://doi.org/10.1007/S11134-022-09848-6/METRICS>
- [15] Isaac Grosf, Yige Hong, and Mor Harchol-Balter. 2024. Analysis of Markovian Arrivals and Service with Applications to Intermittent Overload. *arXiv preprint arXiv:2405.04102* (2024).
- [16] Isaac Grosf, Yige Hong, Mor Harchol-Balter, and Alan Scheller-Wolf. 2023. The RESET and MARC techniques, with application to multiserver-job analysis. *Performance Evaluation* 162 (2023), 102378.
- [17] M. Harchol-Balter. 2013. *Performance Modeling and Design of Computer Systems: Queueing Theory in Action*. Cambridge University Press.
- [18] Mor Harchol-Balter. 2022. The multiserver job queueing model. *Queueing Systems* 100, 3 (2022), 201–203.
- [19] Benjamin Hindman, Andy Konwinski, Matei Zaharia, Ali Ghodsi, Anthony D Joseph, Randy H Katz, Scott Shenker, and Ion Stoica. 2011. Mesos: A platform for fine-grained resource sharing in the data center. In *NSDI*, Vol. 11. 22–22.
- [20] Sungjin Im, Nathaniel Kell, Janardhan Kulkarni, and Debmalya Panigrahi. 2015. Tight Bounds for Online Vector Scheduling. *Proceedings - Annual IEEE Symposium on Foundations of Computer Science, FOCS 2015-December* (12 2015), 525–544. <https://doi.org/10.1109/FOCS.2015.39>
- [21] Sungjin Im, Benjamin Moseley, Kamesh Munagala, and Kirk Pruhs. 2020. Dynamic weighted fairness with minimal disruptions. *Proceedings of the ACM on Measurement and Analysis of Computing Systems* 4, 1 (2020), 1–18.
- [22] Stefano Leonardi and Danny Raz. 2007. Approximating total flow time on parallel machines. *J. Comput. System Sci.* 73, 6 (2007), 875–891.
- [23] Qiao Li and Rohit Negi. 2011. Distributed throughput-optimal scheduling in ad hoc wireless networks. In *2011 IEEE International Conference on Communications (ICC)*. IEEE, 1–5.
- [24] Don Lipari. 2012. The SLURM Scheduler Design. *SLURM User Group*. http://slurm.schedmd.com/slurm_ug_2012/SUG-2012-Scheduling.pdf (2012).
- [25] Siva Theja Maguluri and Rayadurgam Srikant. 2013. Scheduling jobs with unknown duration in clouds. *IEEE/ACM Transactions On Networking* 22, 6 (2013), 1938–1951.
- [26] Siva Theja Maguluri, R. Srikant, and Lei Ying. 2014. Heavy traffic optimal resource allocation algorithms for cloud computing clusters. *Performance Evaluation* 81 (2014), 20–39. <https://doi.org/10.1016/j.peva.2014.08.002>
- [27] Mihalis G Markakis, Eytan Modiano, and John N Tsitsiklis. 2013. Max-weight scheduling in queueing networks with heavy-tailed traffic. *IEEE/ACM Transactions on networking* 22, 1 (2013), 257–270.

- [28] Marcel F Neuts. 1978. The M/M/1 queue with randomly varying arrival and service rates. *Opsearch* 15 (1978), 139–157.
- [29] Jian Ni, Bo Tan, and R. Srikant. 2012. Q-CSMA: Queue-length-based CSMA/CA algorithms for achieving maximum throughput and low delay in wireless networks. *IEEE/ACM Transactions on Networking* 20, 3 (6 2012), 825–836. <https://doi.org/10.1109/TNET.2011.2177101>
- [30] Konstantinos Psounis, Pablo Molinero-Fernández, Balaji Prabhakar, and Fragkiskos Papadopoulos. 2005. Systems with multiple servers under heavy-tailed workloads. *Performance Evaluation* 62, 1-4 (2005), 456–474.
- [31] Konstantinos Psychas and Javad Ghaderi. 2018. Randomized algorithms for scheduling multi-resource jobs in the cloud. *IEEE/ACM Transactions on Networking* 26, 5 (10 2018), 2202–2215. <https://doi.org/10.1109/TNET.2018.2863647>
- [32] GJK Regterschot and JHA De Smit. 1986. The queue M//G/1 with Markov modulated arrivals and services. *Mathematics of operations research* 11, 3 (1986), 465–483.
- [33] SchedMD. 2021. SLURM Workload Manager. (2021). https://slurm.schedmd.com/heterogeneous_jobs.html.
- [34] Alexander L. Stolyar. 2004. MaxWeight scheduling in a generalized switch: State space collapse and workload minimization in heavy traffic. *https://doi.org/10.1214/aoap/1075828046* 14, 1 (2 2004), 1–53. <https://doi.org/10.1214/AOAP/1075828046>
- [35] Leandros Tassioulas and Anthony Ephremides. 1990. Stability properties of constrained queueing systems and scheduling policies for maximum throughput in multihop radio networks. In *29th IEEE Conference on Decision and Control*. IEEE, 2130–2132.
- [36] Muhammad Tirmazi, Adam Barker, Nan Deng, Md E Haque, Zhijing Gene Qin, Steven Hand, Mor Harchol-Balter, and John Wilkes. 2020. Borg: the next generation. In *Proceedings of the Fifteenth European Conference on Computer Systems*. 1–14.
- [37] Abhishek Verma, Luis Pedrosa, Madhukar Korupolu, David Oppenheimer, Eric Tune, and John Wilkes. 2015. Large-scale cluster management at Google with Borg. In *Proceedings of the Tenth European Conference on Computer Systems (Bordeaux, France) (EuroSys '15)*. Association for Computing Machinery, New York, NY, USA, Article 18, 17 pages. <https://doi.org/10.1145/2741948.2741964>
- [38] Qiaomin Xie Weina Wang and Mor Harchol-Balter. 2021. Zero Queueing for Multi-Server Jobs. *Proceedings of the ACM on Measurement and Analysis of Computing Systems* 5, 1 (2 2021), 1–25. <https://doi.org/10.1145/3447385>
- [39] Adam Wierman and Mor Harchol-Balter. 2003. Classifying scheduling policies with respect to unfairness in an M/GI/1. In *Proceedings of the 2003 ACM SIGMETRICS international conference on Measurement and modeling of computer systems*. 238–249.

Appendix

A Proof of Lemma 1

Lemma 1 (Stability Condition). *Given a multiresource system with K job types, arrival vector λ and service requirement vector μ , the system is stable under an MSR policy p if and only if the long-run average completion rate for each job type exceeds the long run arrival rate. That is, the system is stable if and only if $\mu \otimes \mathbb{E}[\mathbf{u}^p] > \lambda$, where \otimes denotes the Hadamard product between vectors.*

PROOF. These criteria are clearly necessary for stability, since it is easy to see that when the condition is violated, $\mathbb{E}[Q_i^p]$ is infinite for at least one job type, i .

We must then show that our condition is sufficient for stability. We first define a test function $V(\mathbf{Q}(t))$ to capture the system state as a number,

$$V(\mathbf{Q}(t)) = \|\mathbf{Q}\|_2^2.$$

For convenience of notation, we set $V(t) = V(\mathbf{Q}(t))$. Intuitively, this test function is a measure of queue length. The larger the test function is, the longer the overall queue is. Then, we invoke Foster-Lyapunov Theorem [34] and show that $\{\mathbf{Q}(t), t \geq 0\}$ is positive-recurrent. Let $\Delta V(\mathbf{q})$ denote the *drift* at state \mathbf{q} , where drift denotes the instantaneous change rate of the test function when the current state is \mathbf{q} . Formally,

$$\Delta V(\mathbf{q}) = \lim_{t \rightarrow 0} \frac{1}{t} \mathbb{E}[V(t) - V(0) \mid \mathbf{Q}(0) = \mathbf{q}].$$

Foster-Lyapunov Theorem states that when the drift function is negative for all but a finite set of exception states, the Markov chain $\{\mathbf{Q}(t), t \geq 0\}$ is positive-recurrent.

Remark 11 (Foster-Lyapunov Theorem). *Let \mathcal{M} be a finite subset of \mathbb{Z}_+^K . If there exists constants $\epsilon, b > 0$, such that*

$$\forall \mathbf{q} \in \mathbb{Z}_+^K, \lim_{t \rightarrow 0} \Delta V(\mathbf{q}) \leq -\epsilon + b \mathbf{1}_{\mathbf{q} \in \mathcal{M}},$$

then $\{\mathbf{Q}(t), t \geq 0\}$ is positive recurrent.

We now bound the drift of our test function V . For some constant $c_1 > 0$,

$$\begin{aligned} \Delta V(\mathbf{q}) &= \lim_{t \rightarrow 0} \frac{1}{t} \mathbb{E}[V(t) - V(0) \mid \mathbf{Q}(0) = \mathbf{q}] \\ &= \lim_{t \rightarrow 0} \frac{1}{t} \mathbb{E}[\|\mathbf{q} + \mathbf{A}(t) - \hat{\mathbf{C}}(t)\|_2^2 - \|\mathbf{q}\|_2^2 \mid \mathbf{Q}(0) = \mathbf{q}] \\ &\leq \lim_{t \rightarrow 0} \frac{1}{t} \mathbb{E}[\|\mathbf{q} + \mathbf{A}(t) - \mathbf{C}(t)\|_2^2 - \|\mathbf{q}\|_2^2 \mid \mathbf{Q}(0) = \mathbf{q}] \quad (8) \\ &= \lim_{t \rightarrow 0} \frac{1}{t} \mathbb{E}[\|\mathbf{A}(t) - \mathbf{C}(t)\|_2^2 + 2\langle \mathbf{q}, \mathbf{A}(t) - \mathbf{C}(t) \rangle \mid \mathbf{Q}(0) = \mathbf{q}] \\ &= \lim_{t \rightarrow 0} \frac{1}{t} (\mathbb{E}[\|\mathbf{A}(t) - \mathbf{C}(t)\|_2^2] + 2\langle \mathbf{q}, \mathbb{E}[\mathbf{A}(t) - \mathbf{C}(t) \mid \mathbf{Q}(0) = \mathbf{q}] \rangle) \\ &\leq \lim_{t \rightarrow 0} c_1^2 K + \frac{2}{t} \langle \mathbf{q}, \mathbb{E}[\mathbf{A}(t) - \mathbf{C}(t) \mid \mathbf{Q}(0) = \mathbf{q}] \rangle. \quad (9) \end{aligned}$$

(8) follows from the fact that when $\hat{\mathbf{C}}_i(t) < C_i(t)$, i.e., when unused service occurs, $\mathbf{q}_i + \mathbf{A}_i(t) - \hat{\mathbf{C}}_i(t) = 0, \forall i$, therefore setting $\hat{\mathbf{C}}(t)$ to $\mathbf{C}(t)$ won't decrease the norm. (9) follows from the fact that $\mathbb{E}\left[\frac{\mathbf{A}(t)}{t}\right] = \boldsymbol{\lambda}$ and $\mathbb{E}\left[\frac{\mathbf{C}(t)}{t}\right] \leq \mathbb{E}\left[\frac{\mathbf{A}(t)}{t}\right]$. Next, we simplify (9):

$$\begin{aligned} \lim_{t \rightarrow 0} \frac{2}{t} \langle \mathbf{q}, \mathbb{E}[\mathbf{A}(t) - \mathbf{C}(t) \mid \mathbf{Q}(0) = \mathbf{q}] \rangle &= \lim_{t \rightarrow 0} \frac{2}{t} \langle \mathbf{q}, \mathbb{E}[\mathbf{A}(t) \mid \mathbf{Q}(0) = \mathbf{q}] - \mathbb{E}[\mathbf{C}(t) \mid \mathbf{Q}(0) = \mathbf{q}] \rangle \\ &= \lim_{t \rightarrow 0} \frac{2}{t} \langle \mathbf{q}, \mathbb{E}[\mathbf{A}(t)] - \mathbb{E}[\mathbf{C}(t)] \rangle \quad (10) \end{aligned}$$

$$\begin{aligned} &= \lim_{t \rightarrow 0} \frac{2}{t} \langle \mathbf{q}, \boldsymbol{\lambda} t - \boldsymbol{\mu} t \otimes \mathbb{E}[\mathbf{u}^p] \rangle \quad (11) \\ &= 2\langle \mathbf{q}, \boldsymbol{\lambda} - \boldsymbol{\mu} \mathbb{E}[\mathbf{u}^p] \rangle \end{aligned}$$

(10) follows from the assumption that $\{\mathbf{A}(t)\}$ and $\{\mathbf{C}(t)\}$ are independent of queue length. (11) follows from the fact that the expected service rate is $\boldsymbol{\mu} \otimes \mathbb{E}[\mathbf{u}^p]$. Therefore, if $\boldsymbol{\lambda} < \boldsymbol{\mu} \otimes \mathbb{E}[\mathbf{u}^p]$, then there exists an $\epsilon, b > 0$ and some \mathcal{M} such that

$$\forall \mathbf{q} \in \mathbb{Z}_+^K, \Delta V(\mathbf{q}) \leq -\epsilon + b \mathbf{1}_{\mathbf{q} \in \mathcal{M}}.$$

Hence, by Remark 11, $\{\mathbf{Q}(t), t \geq 0\}$ is a positive recurrent CTMC if $\boldsymbol{\lambda} < \boldsymbol{\mu} \otimes \mathbb{E}[\mathbf{u}^p]$. \square

B Proof of Theorem 6

Theorem 6. *For any job type, i , consider an MSR policy p and its corresponding MSR-1 policy $p-1$. Let β_i^p denote the maximum number of type- i jobs served in parallel under p . Then,*

$$\mathbb{E}[Q_i^{p-1}] \leq \mathbb{E}[Q_i^p] \leq \mathbb{E}[Q_i^{p-1}] + \beta_i^p.$$

PROOF. We prove this claim via a coupling argument. First, observe that the arrival processes and modulating processes are identical between the MSR and MSR-1 systems. Hence, to compare the two systems on a given sample path, we will assume that arrival times are the same in both systems and that both systems' modulating processes are in corresponding states at any time t .

Next, we couple the departure processes of the two systems as follows. At any time, t , the MSR system attempts to serve $u_i^p(t)$ type- i jobs in parallel. Similarly, the service rate of type- i jobs in the MSR-1 system is $\mu_i u_i^p(t)$. To couple the systems, we set $u_i^p(t)$ exponential timers at any time t corresponding to an event (arrival, departure, or change to the modulating process) in either system. If one of these timers expires before the next event, it triggers a type- i completion in the MSR-1 system as long as there is at least one job to complete. In the MSR system, we map the type- i jobs in service at time t to one timer each. Note that this means some timers may not have an associated job. If a timer expires that is associated with a job in the MSR system, this associated job completes.

We will show that, for any sample path under our coupling,

$$Q_i^{p-1}(t) \leq Q_i^p(t) \leq Q_i^{p-1}(t) + \beta_i^p \quad \forall t \geq 0 \quad (12)$$

We prove this claim by induction on t . Without loss of generality, we assume that the claim holds at time 0 where $Q_i^{p-1}(0) = Q_i^p(0) = 0$. Consider any time t where (12) is assumed to hold for induction. We will consider two cases.

First, consider the case where $Q_i^p(t) < \beta_i^p$. Let t' be the next time after time t where $Q_i^p(t) = \beta_i^p$. In the interval $[t, t']$, any completion in the MSR system corresponds to a completion in the MSR-1 system, unless the MSR-1 system has 0 type- i jobs. Hence, $Q_i^{p-1}(t) \leq Q_i^p(t)$ implies that at each successive completion time $c \in [t, t']$, $Q_i^{p-1}(c) \leq Q_i^p(c)$. Because arrivals affect both systems equally, we have that $Q_i^{p-1}(\tau) \leq Q_i^p(\tau)$ for each arrival time $\tau \in [t, t']$. Finally, because $Q_i^p(\tau) \leq \beta_i^p$ for $\tau \in [t, t']$, it is trivial to see that $Q_i^p(\tau) \leq Q_i^{p-1}(\tau) + \beta_i^p$ for $\tau \in [t, t']$.

Second, consider the case where $Q_i^p(t) \geq \beta_i^p$. Let t' be the next time after t where $Q_i^p(t) < \beta_i^p$. In the interval $[t, t']$, any completion in the MSR system corresponds to a completion in the MSR-1 system, unless the MSR-1 system has 0 type- i jobs. This once again implies that the first part of the inequality holds on the interval $[t, t']$. To prove the second part of the inequality, note that now any completion in the MSR-1 system necessarily corresponds to a completion in the MSR system, since the MSR system has a job associated with every timer for the entire interval. Hence, $Q_i^p(t) \leq Q_i^{p-1}(t) + \beta_i^p$ implies that at each successive completion time $c \in [t, t']$, $Q_i^p(c) \leq Q_i^{p-1}(c) + \beta_i^p$ because every decrease in Q_i^{p-1} corresponds to a simultaneous decrease in Q_i^p . Thus, $Q_i^p(\tau) \leq Q_i^{p-1}(\tau) + \beta_i^p$ for $\tau \in [t, t']$.

Hence, given that (12) holds at time t , it also holds on the entire interval $[t, t']$, where t' is at least one event time later than t . This completes the proof of (12) by induction.

Given that (12) holds for any sample path under our coupling, we have

$$Q_i^{p-1} \leq_{st} Q_i^p \leq_{st} Q_i^{p-1} + \beta_i^p$$

and thus

$$\mathbb{E}[Q_i^{p-1}] \leq \mathbb{E}[Q_i^p] \leq \mathbb{E}[Q_i^{p-1}] + \beta_i^p$$

as desired. \square

C A Queue Length Approximation for the MSR System

The slack in the bounds of Theorem 8 results from the complexity of comparing an $M/M/k$ system with modulated service rates (the MSR system) to an analogous $M/M/1$ system with modulated service rates (the MSR-1 system). However, when service rates are not modulated, this comparison follows the well-known form

$$\mathbb{E}[N^{M/M/k}] = P_Q^{M/M/k} \cdot \mathbb{E}[N^{M/M/1}] + k\rho^{M/M/1} \quad (13)$$

where $P_Q^{M/M/k}$ is the probability an arriving job queues in the $M/M/k$ system, $\rho^{M/M/1}$ is the system load of the $M/M/1$ system, and hence the second term in the sum represents the average number of jobs in service in the $M/M/k$ system. When service rates are generally distributed, previous works [30] also suggest that this difference is difficult to exactly model. To derive Theorem 6, we noticed that the MSR and MSR-1 systems perform similarly *during periods where both systems have sufficiently many jobs in their queues*. As a result, the upper bound from Theorem 6 follows a similar form as (13), but pessimistically bounds P_Q by 1 and bounds the expected number of jobs in service by β_i^p to relate the MSR and MSR-1 systems.

Although we cannot provide a tighter *bound* on these two quantities, we can provide *approximations* of these quantities, resulting in a queue length approximation that performs well at all loads. To approximate P_Q , we treat the MSR system as if it always tries to serve $k_i^* = \mathbb{E}[u_i^p]$ jobs, setting $P_Q = P_Q^{M/M/k_i^*}$. We approximate the average number of type- i jobs in service as $\rho_i \cdot k_i^*$. By applying these approximations instead of applying Theorem 6 in the proof of Theorem 8, we get

$$\mathbb{E}[Q_i^p] \approx P_Q^{M/M/k_i^*} \frac{\rho_i + \mathbb{E}[\Delta(v_i^{p-1})]}{1 - \rho_i} - \mathbb{E}_U[\Delta(v_i^{p-1})] + \rho_i \cdot k_i^*.$$

Finally, as suggested in [16], we approximate $\mathbb{E}_U[\Delta(v_i^{p-1})]$ as $\mathbb{E}[\Delta(v_i^{p-1})]$, yielding the following approximation of mean queue length for an MSR-1 system

$$\mathbb{E}[Q_i^p] \approx P_Q^{M/M/k_i^*} \frac{\rho_i + \rho_i \mathbb{E}[\Delta(v_i^{p-1})]}{1 - \rho_i} + \rho_i \cdot k_i^*.$$

If our prediction lies below or above our lower or upper bounds, respectively, we report the value of the closest bound. Our prediction and bounds are compared to simulations in Section 7.

D Switching Rates for Preemptive MSR Policies

Theorem 12. *For a two-state pMSR policy p , where $N^p = 2$, the mean queue length upper bound is minimized if we are switching schedules infinitely often. That is, $\alpha = \infty$ minimizes the mean length upper bound across all $p(\alpha)$'s.*

PROOF. Let $\alpha G^p = \begin{pmatrix} -\alpha x_1 & \alpha x_1 \\ \alpha x_2 & -\alpha x_2 \end{pmatrix}$ for some $x_1, x_2 > 0$. WLOG, we assume $u_{1,1}^{p-1} \geq u_{2,1}^{p-1}$. Let w_1, w_2 be two working states of the system. By Lemma 9, we get

$$\begin{cases} \begin{pmatrix} -\alpha x_1 & \alpha x_1 \\ \alpha x_2 & -\alpha x_2 \end{pmatrix} (\Delta_1(1), \Delta_1(2)) = \mathbb{E}[u_1^{p-1}](1, 1) - (u_{1,1}^{p-1}, u_{2,1}^{p-1}) \\ \Delta_1(1) \mathbf{P}\{u_1^{p-1} = w_1\} + \Delta_1(2) \mathbf{P}\{u_1^{p-1} = w_2\} = 0 \end{cases}$$

Also, by the definition of CTMC, we have

$$\begin{cases} \begin{pmatrix} -\alpha x_1 & \alpha x_1 \\ \alpha x_2 & -\alpha x_2 \end{pmatrix} \begin{pmatrix} \mathbf{P}\{u_1^{p-1} = w_1\} & \mathbf{P}\{u_1^{p-1} = w_2\} \end{pmatrix} = 0 \end{cases} .$$

Hence,

$$\begin{cases} \mathbf{P}\{u_1^{p-1} = w_1\} = \frac{x_2}{x_1 + x_2} \\ \mathbf{P}\{u_1^{p-1} = w_2\} = \frac{x_1}{x_1 + x_2} \end{cases} .$$

By definition of v_1 , we have

$$\begin{cases} \mathbf{P}\{v_1^{p-1} = w_1\} = \frac{u_{1,1}^{p-1}}{\mathbb{E}[u_1^{p-1}]} \frac{x_2}{x_1+x_2} \\ \mathbf{P}\{v_1^{p-1} = w_2\} = \frac{u_{2,1}^{p-1}}{\mathbb{E}[u_1^{p-1}]} \frac{x_1}{x_1+x_2} \end{cases}.$$

Solving this system, we get

$$\begin{cases} \Delta_1(1) = \alpha x_1 \frac{u_{1,1}^{p-1} - u_{2,1}^{p-1}}{(x_1+x_2)^2} \\ \Delta_1(2) = \alpha x_2 \frac{u_{2,1}^{p-1} - u_{1,1}^{p-1}}{(x_1+x_2)^2} \end{cases}.$$

Plugging into Theorem 8, we get

$$\mathbb{E}[Q_1^{p(\alpha)}] = \frac{\rho_1}{1 - \rho_1} + \frac{x_1 x_2 (u_{1,1}^{p-1} - u_{2,1}^{p-1})^2}{\alpha \mathbb{E}[u_1^{p-1}] (1 - \rho_1) (x_1 + x_2)^3} + B,$$

where $x_2 \frac{u_{2,1}^{p-1} - u_{1,1}^{p-1}}{\alpha (x_1 + x_2)^2} \leq B \leq x_1 \frac{u_{1,1}^{p-1} - u_{2,1}^{p-1}}{\alpha (x_1 + x_2)^2}$. Thus,

$$\mathbb{E}[Q_1^{p(\alpha)}] \leq \frac{\rho_1}{1 - \rho_1} + \frac{x_1 x_2 (u_{1,1}^{p-1} - u_{2,1}^{p-1})^2}{\alpha \mathbb{E}[u_1^{p-1}] (1 - \rho_1) (x_1 + x_2)^3} + x_1 \frac{u_{1,1}^{p-1} - u_{2,1}^{p-1}}{\alpha (x_1 + x_2)^2}.$$

Let $z = \frac{x_1 x_2 (u_{1,1}^{p-1} - u_{2,1}^{p-1})^2}{\alpha \mathbb{E}[u_1^{p-1}] (1 - \rho_1) (x_1 + x_2)^3} + x_1 \frac{u_{1,1}^{p-1} - u_{2,1}^{p-1}}{\alpha (x_1 + x_2)^2}$. We can see that $\lim_{\alpha \rightarrow \infty} z = 0$ and that $z \geq 0$. Analysis on $\mathbb{E}[Q_2^p]$ upper bound is analogous.

Hence, we have shown that setting $\alpha = \infty$ minimizes this queue length upper bound. \square

E Proof of Theorem 4

Theorem 4. *When scheduling non-preemptible jobs, the class of nMSR policies is throughput-optimal. That is, if there exists a policy that stabilizes a multiresource job system with K job types and arrival rates λ , then there exists an nMSR policy, q , with $N^q \leq K$ working states that stabilizes the system. Specifically, an nMSR policy can stabilize a system with non-preemptible jobs when $\lambda \in C$.*

PROOF. From Theorem 2, we know that there exists an MSR policy p that can stabilize the system. Hence, given some MSR policy p that stabilizes the system, we will show how to construct a non-preemptive policy that also stabilizes the system.

Obviously, if p is already non-preemptive, we have proven our claim. Otherwise, we can use p to construct an nMSR policy that stabilizes the system.

To do this, we first define $p(\alpha)$ to be the MSR policy that with infinitesimal generator $\mathbf{G}^{p(\alpha)} = \alpha \mathbf{G}^p$. That is, $p(\alpha)$ is an MSR policy whose modulating process is a scaled version of p 's modulating process. We will create a non-preemptive policy, $q(\alpha)$, that is based on $p(\alpha)$. To make $q(\alpha)$ non-preemptive, we must add some switching states between the states of $p(\alpha)$'s modulating process. The original states of $p(\alpha)$'s modulating process will form the working states of $q(\alpha)$. We assume that the added switching states form arbitrary (finite length) switching routes between these working states. Hence, $q(\alpha)$ is a valid nMSR policy. We will now show that, for sufficiently small α , $q(\alpha)$ stabilizes the system.

Let w_1, w_2, \dots, w_{N^p} denote the working states of $q(\alpha)$, and let $s_1, s_2, \dots, s_{N^q - N^p}$ denote all the switching states of $q(\alpha)$. Let $\pi^{q(\alpha)}$ denote the stationary distribution of states in the modulating

process of $q(\alpha)$. Let $\pi_W^{q(\alpha)}$ be the probability that we are in any working state defined as

$$\pi_W^{q(\alpha)} = \sum_{i=1}^{N^p} \pi_{w_i}^{q(\alpha)}.$$

By Lemma 1, we have $\boldsymbol{\mu} \otimes \mathbb{E}[\mathbf{u}^p] > \boldsymbol{\lambda}$. Therefore, there exists $\epsilon_1 > 0$ such that $\boldsymbol{\mu} \otimes \mathbb{E}[\mathbf{u}^p] - \epsilon_1 = \boldsymbol{\lambda}$. Next, we condition $\mathbb{E}[\mathbf{u}^{q(\alpha)}]$ on whether it is in working state.

$$\begin{aligned} \mathbb{E}[\mathbf{u}^{q(\alpha)}] &= \mathbb{E}[\mathbf{u}^{q(\alpha)} \mid \text{working state}] \pi_W^{q(\alpha)} + \mathbb{E}[\mathbf{u}^{q(\alpha)} \mid \text{not working state}] (1 - \pi_W^{q(\alpha)}) \\ &\geq \mathbb{E}[\mathbf{u}^{q(\alpha)} \mid \text{working state}] \pi_W^{q(\alpha)} \\ &= \mathbb{E}[\mathbf{u}^p] \pi_W^{q(\alpha)} \\ &= \mathbb{E}[\mathbf{u}^p] \pi_W^{q(\alpha)} \end{aligned}$$

Because Lemma 1 tells us that $\mathbb{E}[\mathbf{u}^p] - \epsilon_1 = \boldsymbol{\lambda}$ for some $\epsilon_1 > 0$, it will suffice to show that $\pi_W^{q(\alpha)}$ can be arbitrarily close to 1 when α is small. We prove this in Lemma 13.

Lemma 13. *For any $\epsilon > 0$, there exists an α such that*

$$\pi_W^{q(\alpha)} > 1 - \epsilon$$

PROOF. We will prove the Lemma via the renewal-reward theorem [17]. We let the renewal cycle be the time between visits to state w_1 by $\mathbf{u}^{q(\alpha)}(t)$. Consider a path $\mathbf{u}^{q(\alpha)}(t)$ takes, ω , during a renewal cycle, we use the random variable X to define the cycle length and let the reward function be $R^{q(\alpha)}(t) = \mathbf{1}[\mathbf{u}^{q(\alpha)}(t) \text{ is in a working state}]$. Let $R^{q(\alpha)}$ be the random variable denoting the reward accrued during one renewal cycle. Let $S^{q(\alpha)}$ be the random variable denoting the time $\mathbf{u}^{q(\alpha)}$ spends in switching states during a renewal cycle.

Then, by the renewal-reward theorem [17], we have

$$\begin{aligned} \pi_W^{q(\alpha)} &= \frac{\mathbb{E}[R^{q(\alpha)}]}{\mathbb{E}[X]} \\ &= \frac{\mathbb{E}[R^{q(\alpha)}]}{\mathbb{E}[R^{q(\alpha)}] + \mathbb{E}[S^{q(\alpha)}]} \\ &= \frac{\frac{1}{\alpha} \mathbb{E}[R^{q(1)}]}{\frac{1}{\alpha} \mathbb{E}[R^{q(1)}] + \mathbb{E}[S^{q(1)}]} \\ &= \frac{\mathbb{E}[R^{q(1)}]}{\mathbb{E}[R^{q(1)}] + \alpha \mathbb{E}[S^{q(1)}]}. \end{aligned} \tag{14}$$

Step (14) follows from

$$\begin{aligned} \mathbb{E}[R^{q(\alpha)}] &= \sum_{\omega} \mathbb{E}[R^{q(\alpha)} \mid \omega] \mathbf{P}\{\omega\} \\ &= \sum_{\omega} \frac{1}{\alpha} \mathbb{E}[R^{q(1)} \mid \omega] \mathbf{P}\{\omega\} \\ &= \frac{1}{\alpha} \mathbb{E}[R^{q(1)}] \end{aligned}$$

because the transition rates out of the working states are being multiplied by α in $q(\alpha)$ compared to $q(1)$. Therefore, we only need to set the α to be

$$0 < \alpha < \frac{\epsilon \mathbb{E}[R^{q(1)}]}{(1 - \epsilon) \mathbb{E}[S^{q(1)}]}$$

such that $\pi_W^{q(\alpha)} > 1 - \epsilon$. \square

From Lemma 13, we know that there exists some α such that $\pi_W^{q(\alpha)} = 1 - \epsilon_2$ holds for a given $\epsilon_2 > 0$. Hence, if we are scaling the transitions out of working states by α , we have,

$$\begin{aligned} \boldsymbol{\mu} \otimes \mathbb{E}[\mathbf{u}^q] &= \boldsymbol{\mu} \otimes \mathbb{E}[\mathbf{u}^p] \pi_W(\alpha) \\ &> \boldsymbol{\mu} \otimes \mathbb{E}[\mathbf{u}^p] (1 - \epsilon_2) \\ &= (\boldsymbol{\lambda} + \epsilon_1) (1 - \epsilon_2) \\ &= \boldsymbol{\lambda} + \epsilon_1 - \epsilon_2 - \epsilon_1 \epsilon_2. \end{aligned}$$

In this case, we can set $\epsilon_3 = \epsilon_1 - \epsilon_2 - \epsilon_1 \epsilon_2$. We need $\epsilon_3 = \epsilon_1 (1 - \epsilon_2) - \epsilon_2 > 0$. Therefore, given ϵ_1 , we can set ϵ_2 to be $\frac{\epsilon_2}{1 - \epsilon_2} < \epsilon_1$.

Lastly, by invoking Lemma 1, we have shown that the constructed nMSR policy with the α associated with ϵ_2 can stabilize the system. \square

F Example System from Section 7

To evaluate our MSR policies in Section 7, we consider a server with three resource types (WLOG, we call the resources CPU cores, DRAM, and network bandwidth). Our example server has 20 cores, 15 GB of DRAM and 50 Gbps of network bandwidth. We assume there are three job types with demands

$$\mathbf{D}_1 = (3, 7, 1) \quad \mathbf{D}_2 = (4, 1, 1) \quad \mathbf{D}_3 = (10, 1, 5).$$

When the system load, ρ , equals 1, we let

$$\boldsymbol{\lambda} = (.5, 2, 1) \quad \boldsymbol{\mu} = (1, 1, 1).$$

When we vary the system load in our experiments, we keep $\boldsymbol{\mu}$ fixed, and scale $\boldsymbol{\lambda}$ linearly such that

$$\boldsymbol{\lambda} = \rho \cdot (.5, 2, 1).$$

Note that the arrival rate vector was chosen such that $\rho < 1$ implies $\boldsymbol{\lambda} \in \mathcal{C}$.

This example was chosen to be non-trivial in two key ways. First, some possible schedules in this example are Pareto optimal, but do not lie on the boundary of the convex hull of \mathcal{S} . For example, the schedule (1, 1, 1) cannot fit an additional job of any type without removing some jobs from the schedule, but it does not lie on the convex boundary of \mathcal{S} . Hence, although the schedule (1, 1, 1) seems to pack the server efficiently, choosing this schedule will not help stabilize the system when system load is sufficiently high.

Second, the possible schedules on the convex boundary of \mathcal{S} are not trivially close together. Hence, in cases where preemption is limited, nMSR and sMSR policies will need non-trivial switching routes to move between working states. The modulating processes for the pMSR, nMSR, and sMSR policies we construct based on this example are shown in Figure 6.

G Phase-Type Job Sizes

In the modulating processes of pMSRs and sMSRs, the state transitions do not coincide with job completions. Therefore, given jobs whose sizes follow a different distribution, the modulating process will not change. As described in [15], the RESET and MARC techniques used in the proof of Theorem 8 can be expanded to handle the case where job sizes follow phase-type distribution. Hence, our queue length bounds from this theorem will continue to hold in this case. Specifically, when dealing with phase-type distributions, the analysis of the MSR-1 system requires computing relative completions terms for each phase of type- i jobs. Hence, while the computation of our bounds is still straightforward in this case, there is some added complexity to having these additional job phases.

The story is more complicated in the nMSR case, where the modulating process relies on job completions to transition between working states. Here, many additional states need to be added to the modulating process to capture all the possible switching routes that the nMSR policy may be forced to follow between working states. Hence, the size of the modulating process might grow exponentially with the number of job phases. Finding a compact representation of the modulating process for nMSR policies with phase-type jobs remains an open problem that we defer to future work.

H Proof of Theorem 10

Our goal is to compare an MSR policy to MaxWeight in a heavy traffic scaling regime where $\rho \rightarrow 1$. Due to the definition of ρ , there are many possible ways to achieve this scaling. In Definition 3, we define an intuitive heavy-traffic scaling regime where the arrival rate vector is scaled linearly to increase system load.

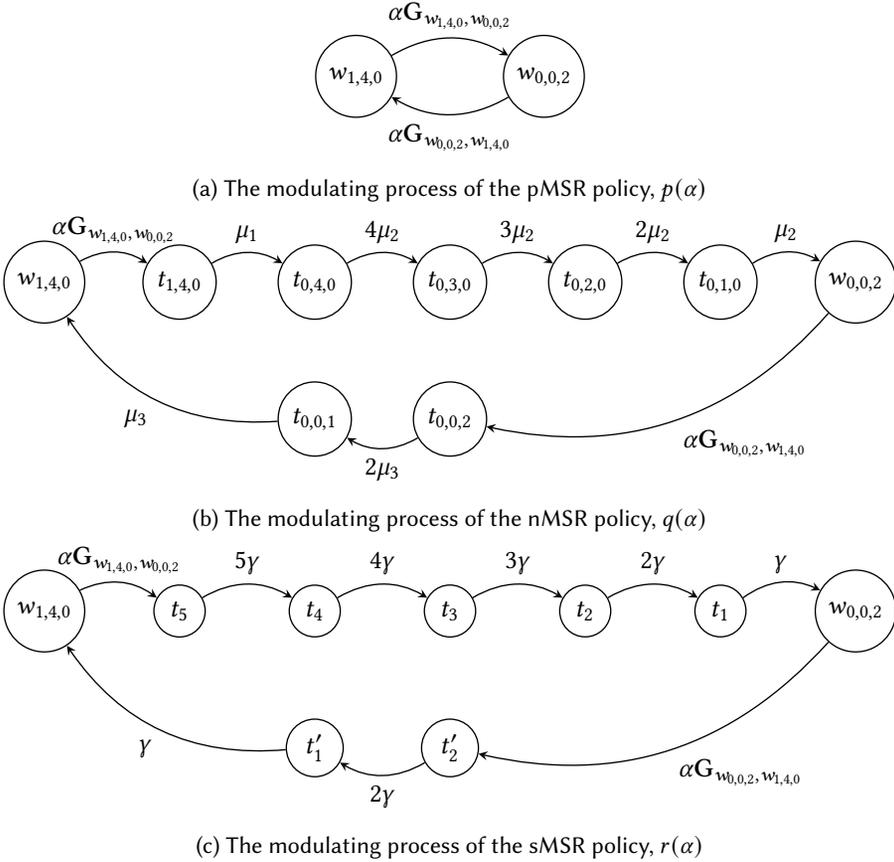


Fig. 6. Here we depict the structure of the modulating processes for $p(\alpha)$, $q(\alpha)$ and $r(\alpha)$. These policies are used for the examples described in Section 7. States marked with w are working states, and states marked with t are switching states. In (a) and (b), the subscripts denote the state's corresponding schedule. In (c), the subscripts for switching states denote the number of jobs experiencing a setup time due to being preempted.

Definition 3. Consider a system with initial arrival rate vector λ and initial system load ρ_0 . Let $\tilde{\lambda} = \frac{\lambda}{\rho_0}$. Here, $\tilde{\lambda}$ represents the supremum of the set of scaled arrival rate vectors under which the system may be stable. We consider the system in the limit as the arrival rate vector scales linearly to $\tilde{\lambda}$. Formally, for any value of system load, ρ , let the scaled arrival rate, $\lambda(\rho)$, be defined as

$$\lambda(\rho) = \tilde{\lambda} \cdot \rho.$$

We consider the heavy-traffic limit as $\rho \rightarrow 1$.

Using this definition of heavy traffic, we prove the following theorem.

Theorem 10. Under the heavy-traffic scaling regime in Definition 3, the solution to (7) is the same at all loads. The policy, p^* , that we select using (7) is constant-competitive with MaxWeight in heavy traffic. That is, there exists a constant $C > 0$ such that

$$\lim_{\rho \rightarrow 1} \frac{\mathbb{E}[Q^p]}{\mathbb{E}[Q^{\text{MaxWeight}}]} \leq C$$

PROOF. First, we prove that the solution to (7) is the same at all loads. This follows directly from Definition 3, as (7) is effectively finding a policy p^* whose average completion rate is $\tilde{\lambda}$. Because $\tilde{\lambda}$ is the same at all loads, the policy p^* does not depend on ρ .

To prove that p^* is constant competitive, we note that Lemma 9 tells us $\mathbb{E}[\Delta(v_i^{p^*-1})]$ is only dependent on \mathbf{G} , the infinitesimal generator of the modulating process. Since the policy p^* does not change with ρ , $\mathbb{E}[\Delta(v_i^{p^*-1})]$ is a constant. Therefore, $\mathbb{E}[Q_i^{p^*-1}]$ scales as $\Theta(1/(1 - \rho_i))$. Our optimization objective in (7) ensures that $\rho = \rho_i$ for all i , therefore $\mathbb{E}[Q^{p^*}]$ scales as $\Theta(1/(1 - \rho)) = \Theta(1/\epsilon)$. Hence, to prove that p^* is constant-competitive with MaxWeight, it suffices to show that $\mathbb{E}[Q^{\text{MaxWeight}}]$ also scales as $\Theta(1/(1 - \rho))$.

It is shown in [7] that $\mathbb{E}[Q^{\text{MaxWeight}}]$ scales as $\Theta(1/\epsilon)$ in the limit as $\epsilon \rightarrow 0$, where ϵ is the minimum distance between λ and the faces of $\text{Conv}(\mathcal{S})$. When ϵ is sufficiently small, the face of $\text{Conv}(\mathcal{S})$ closest to λ does not change as $\rho \rightarrow 1$. Hence, for sufficiently large ρ , let θ be the angle between the vector λ and the normal vector of the closest face on $\text{Conv}(\mathcal{S})$. Here, $\epsilon = (1 - \rho) \cos(\theta)$. Because ϵ and $1 - \rho$ differ by at most a constant factor, $\mathbb{E}[Q^{\text{MaxWeight}}]$ scales as $\Theta(1/(1 - \rho))$ as $\rho \rightarrow 1$. Hence, p^* is constant-competitive with MaxWeight. \square

I Description of the Google Borg Trace

Our trace is based on data collected from Google Borg over 30 days in 2019 [36]. The Borg dataset is divided into 8 cells that represent collections of machines in 8 different time zones. The trace describes job arrival times, sizes, and the normalized CPU and memory demands of each job.

We considered data from one cell. We first group the jobs in the trace into distinct job types, grouping together jobs whose CPU and memory demands are within 0.1%. While this resulted in hundreds of thousands of job types, the ten most frequent job types accounted for 51% of all requests. Hence, to construct a tractable scenario, we consider these ten most popular types. We further filtered out jobs that were killed prematurely or did not finish during the observation window of the trace.

The trace of the top ten job types contains 69 million requests. The average job size across all types was 310 seconds. The job size distribution is far from exponential – the distribution’s failure rate is initially increasing and then decreases quickly after 50 seconds. The arrival process in the trace is highly variable and often includes large batches of arriving jobs.

J An Example on the Initialization Step of MSR

Consider a server with 128 cores of CPU, 256 GiB of memory, 1TiB of disk space, and 100Gbps of network. Table 2 lists the 4 types of virtual machines some cloud company wants to sell. For simplicity, we assume the mean job sizes are all 1 second for all job types. That is, $\mu = (1, 1, 1, 1)$. The measurements at the company tell us we get 4, 5, 2, 1.5 virtual machine requests per second for type-1, type-2, type-3, type-4 virtual machines, respectively. We can therefore plug these numbers into equation (7) to solve for the candidate set and the fraction of time we need to spend in each candidate. In this case, the capacity of the system is $\mathbf{P} = (128, 256, 1024, 100)$, the demand matrix is

$$\mathbf{D} = \begin{pmatrix} 1 & 4 & 50 & 1 \\ 4 & 1 & 10 & 10 \\ 2 & 2 & 100 & 5 \\ 8 & 4 & 10 & 1 \end{pmatrix},$$

and the arrival rate vector being $\lambda = (4, 5, 2, 1.5)$.

It takes the Gurobi solver under 1ms to solve for this MIQCP. The results are:

$$\mathbf{W} = \begin{pmatrix} 0 & 7 & 4 & 10 \\ 0 & 5 & 9 & 5 \\ 0 & 10 & 0 & 0 \\ 10 & 9 & 0 & 0 \end{pmatrix}$$

and

$$\pi = (0.07633588 \quad 0.30534351 \quad 0.00763359 \quad 0.61068702).$$

Type	CPU	Memory	Disk	Network	Demand
Type-1	1 core	4GiB	50GiB	1Gbps	4 requests/sec
Type-2	4 cores	1GiB	10GiB	10Gbps	5 requests/sec
Type-3	2 cores	2GiB	100GiB	5Gbps	2 requests/sec
Type-4	8 cores	4GiB	10GiB	1Gbps	1.5 requests/sec

Table 2. The resource demands and arrival rates of the 4 job types in the example system.

Received January 2025; revised April 2025; accepted April 2025