

Empirical Gittins: M/G/1 Scheduling from Job Size Samples

Shefali Ramakrishna
Cornell University
sr899@cornell.edu

Amit Harlev
Cornell University
ah843@cornell.edu

Ziv Scully
Cornell University
zivscully@cornell.edu

1. INTRODUCTION

We consider the classic problem of minimizing mean response time in the M/G/1 queue. The optimal scheduling policy depends on the amount of information the scheduler has about job sizes (service times).

- (a) When sizes are known, Shortest Remaining Processing Time (SRPT) is optimal.
- (b) When sizes are unknown but *drawn from a known distribution*, the Gittins policy is optimal [3, 7].
- (c) When no information about sizes or their distribution is known, Randomized Multi-Level Feedback (RMLF) achieves the best known competitive ratio of $O(\log \frac{1}{1-\rho})$, where $\rho \in (0, 1)$ is the load [1, 2, 4]

In systems with unknown job sizes, the setting that often arises is one between (b) and (c). For example, we might not know job sizes or their distribution, but we may have information on past jobs' sizes, or be able to learn the distribution from future jobs' sizes. Motivated by this setting, we study the setting in which the scheduler has access only to a finite number of i.i.d. samples from the job size distribution. We ask: *can the scheduler construct a near-optimal policy from finitely many samples?*

1.1 Our Work: Empirical Gittins Policy

The Gittins policy uses the job size distribution to construct a priority rule, then uses that rule to schedule jobs. Specifically, this priority rule is a *rank function* [8], which maps a job's *age* (attained service) to its *rank* (lower rank is better priority). One can thus view the Gittins policy as a map

$$(\text{job size distribution, age}) \mapsto \text{rank.}$$

If we only have access to finitely many job size samples instead of the true job size distribution, a natural approach is to use the same map above, but plugging in the *empirical distribution* instead of the true distribution. We call this the *empirical Gittins* policy. This policy is simple and intuitive, but its performance relative to the optimum is unclear.

We present the first analysis of M/G/1 scheduling when the scheduler has only sample access to the job size distribution. In the finite-support setting, we prove that the empirical Gittins policy achieves $(1 + \varepsilon)$ -approximate mean response time with probability at least $1 - \delta$ when given $\Omega(\frac{1}{\varepsilon^2} \log \frac{1}{\delta})$ samples (Section 2.2). The main technical step of our proof is a generic result about the Gittins policy with a misspecified

This work was supported by the NSF under grant no. CMMI-2307008. Amit Harlev was supported by the DoD through the NDSEG Fellowship Program.

Copyright is held by author/owner(s).

job size distribution. Roughly speaking, we show that if the misspecified density function is within a γ factor of the true density function, then the resulting misspecified Gittins policy is within a γ^2 factor of optimal (Section 2.1).

We conclude by giving empirical evidence suggesting that empirical Gittins may perform well even with continuous job size distributions, though analyzing this case theoretically remains an open problem (Section 3).

1.2 Related Work: Learning and Scheduling

Recent work at the intersection of learning and queueing has focused on developing scheduling policies that adapt to unknown or changing environments [5, 9, 10]. These approaches are often motivated by the idea that real-world systems evolve over time, requiring policies that can learn from experience and adjust accordingly.

Our work is driven by the same high-level goal, namely designing scheduling policies that adapt well to changing environments, but differs in two key ways. First, prior work [5, 9, 10] assumes Markovian service (i.e. exponential or geometric job sizes) with multiple job classes, which means the learning problem reduces to estimating a vector of values. In contrast, we work with *general job size distributions* with no parametric assumptions. Second, prior work formulates an online learning problem, using transient analyses to obtain regret bounds. In contrast, we study a “one-shot” problem: we use a finite number of samples to construct a scheduling policy, then evaluate that policy's steady-state performance. This formulation is certainly simpler than online learning, but it allows us to use modern queueing tools [6, 7] to study the more complicated setting of general job size distributions.

2. RESULTS FOR FINITE-SUPPORT JOB SIZE DISTRIBUTIONS

We analyze the M/G/1 queue with job sizes drawn i.i.d. from distribution S . We study *nonclairvoyant* scheduling policies, meaning policies that do not use information about any individual job's size. We assume a standard overhead-free preempt-resume model. We assume load $\rho < 1$ for stability.

All of the policies we consider are *SOAP* policies [6, 8]. These are policies that schedule jobs using a *rank function* $r : \mathbb{R}_+ \rightarrow \mathbb{R}$ which maps each job's *age* a (a.k.a. attained service), namely the amount of time the job has been served so far, to a *rank* $r(a)$, which is a number representing a priority (lower is better). All SOAP policies follow the same general scheduling rule: *at every moment in time, serve the job with lowest rank*. In this section, ties for lowest rank can be broken arbitrarily (but we revisit this in Section 3).

The *Gittins rank* is defined in terms of an *efficiency func-*

tion $J(a, b)$, which captures how advantageous it is to serve a job of age a until it reaches age $b > a$. Formally, the Gittins rank is given by:

$$r_{\text{Gittins}}(a) = \inf_{b>a} J(a, b), \quad (2.1)$$

$$\text{where } J(a, b) = \frac{\mathbb{E}[\min(S, b) - a \mid S > a]}{\mathbb{P}[S \leq b \mid S > a]}. \quad (2.2)$$

To prove guarantees for empirical Gittins, we use a structural result from Scully [6]: any SOAP policy whose rank function approximates the true Gittins rank pointwise achieves near-optimal performance.

Theorem 2.1 (corollary of [6, Theorem 16.5]). *Consider an M/G/1 queue in which the scheduler does not observe individual job sizes. Let r_{Gittins} denote the Gittins rank function derived from the true job size distribution. Suppose an age-based priority policy π has rank function r_π satisfying*

$$\frac{1}{\gamma} r_{\text{Gittins}}(a) \leq r_\pi(a) \leq \gamma r_{\text{Gittins}}(a) \text{ for all ages } a$$

has mean response time within a factor of γ^2 of optimal:

$$\mathbb{E}[T_\pi] \leq \gamma^2 \mathbb{E}[T_{\text{Gittins}}].$$

We will show that given enough samples, the empirical Gittins policy has a high probability of having a rank function within a multiplicative factor of the true Gittins rank. By Theorem 2.1, this will imply a bound on the response time of the empirical Gittins policy. We do this in two steps:

- (Section 2.1) We show that an ‘‘approximately correct’’ job size distribution results in an ‘‘approximately correct’’ Gittins rank function.
- (Section 2.2) We show that with enough samples, the empirical distribution is ‘‘approximately correct’’ with high probability.

2.1 Gittins with Misspecified Size Distribution

Consider two finite-support PMFs p_1 and p_2 . The intuition is that p_1 is the true PMF of S and p_2 is a misspecified PMF, but most of our intermediate results will treat them symmetrically. Throughout, let r_i and J_i be the Gittins rank functions (2.1) and efficiency functions (2.2), respectively, when S is the distribution with PMF p_i .

Lemma 2.2. *Suppose*

$$\beta \leq \frac{p_2(s)}{p_1(s)} \leq \alpha$$

for some constants $0 < \beta \leq \alpha$ and for all $s \in \text{supp}(p_1)$. Then for all $0 < a < b$,

$$\frac{\beta}{\alpha} \leq \frac{J_2(a, b)}{J_1(a, b)} \leq \frac{\alpha}{\beta}.$$

Proof. By assumption, for all s in the support of p_1 , we have

$$\begin{aligned} J_1(a, b) &= \frac{\sum_{s>a} p_1(s) \cdot (\min(s, b) - a)}{\sum_{a<s\leq b} p_1(s)} \\ &\geq \beta \cdot \frac{\sum_{s>a} p_1(s) \cdot (\min(s, b) - a)}{\sum_{a<s\leq b} p_2(s)} \\ &\geq \frac{\beta}{\alpha} \cdot \frac{\sum_{s>a} p_2(s) \cdot (\min(s, b) - a)}{\sum_{a<s\leq b} p_2(s)} = \frac{\beta}{\alpha} J_2(a, b), \end{aligned}$$

and symmetrically, $J_2(a, b) \geq \frac{\beta}{\alpha} J_1(a, b)$. Combining yields the desired bounds. \square

Since the Gittins rank is defined as $r(a) = \inf_{b>a} J(a, b)$, the efficiency function bounds from Lemma 2.2 translate directly into bounds on the rank functions r_1 and r_2 . We now formalize this:

Lemma 2.3. *If the efficiency functions J_1 and J_2 have*

$$\frac{\beta}{\alpha} \leq \frac{J_2(a, b)}{J_1(a, b)} \leq \frac{\alpha}{\beta}$$

for all ordered pairs of ages $a < b$, then their corresponding Gittins ranks satisfy

$$\frac{\beta}{\alpha} \leq \frac{r_2(a)}{r_1(a)} \leq \frac{\alpha}{\beta}$$

for all ages a .

Proof. Let $b_1^*(a)$ and $b_2^*(a)$ be the minimizers of $J_1(a, b)$ and $J_2(a, b)$ respectively. Then

$$\begin{aligned} r_1(a) &= J_1(a, b_1^*(a)) \leq J_1(a, b_2^*(a)) \\ &\leq \frac{\alpha}{\beta} J_2(a, b_2^*(a)) = \frac{\alpha}{\beta} r_2(a), \end{aligned}$$

and, symmetrically, $r_2(a) \leq \frac{\alpha}{\beta} r_1(a)$. \square

Together, Lemma 2.2 and Lemma 2.3 immediately imply our first main result below.

Theorem 2.4. *Let p_1, p_2, r_1, r_2 be as defined at the start of this section, and let $\mathbb{E}[T_{i,j}]$ be the mean response time in an M/G/1 whose job size distribution has PMF p_i and is using the SOAP policy with rank function r_j (the Gittins rank function from PMF p_j). Suppose that there exist constants $0 < \beta \leq \alpha$ such that for all $s \in \text{supp}(p_1)$,*

$$\beta \leq \frac{p_2(s)}{p_1(s)} \leq \alpha.$$

Then

$$\mathbb{E}[T_{1,2}] \leq \left(\frac{\alpha}{\beta}\right)^2 \mathbb{E}[T_{1,1}].$$

Remark 2.5. All of the results in this section extend beyond the finite-support case. To generalize the result, view p_1 and p_2 as density functions of two job size distributions with respect to a common measure μ (e.g. the Lebesgue measure). Then all of the statements and proofs generalize by replacing sums with integrals with respect to μ . In fact, p_1 and p_2 just need to be densities of finite measures, not necessarily probability measures.

2.2 Performance Bound for Empirical Gittins

We now apply Theorem 2.4 to analyze Empirical Gittins. Let p_{true} be the PMF of the true size distribution S , and let P_{emp} be the *random* PMF of the empirical distribution from n samples from S . This means that for all s ,

$$n P_{\text{emp}}(s) \sim \text{Binomial}(n, p_{\text{true}}(s)). \quad (2.3)$$

We write $\mathbb{E}[T_{\text{true}}]$ for the mean response time of the true Gittins policy and $\mathbb{E}[T_{\text{emp}}]$ for that of empirical Gittins.

Theorem 2.6 (Performance guarantee for empirical Gittins). *Let $\delta, \varepsilon \in (0, 1)$. The empirical Gittins policy achieves*

$$\mathbb{P}\left[\mathbb{E}[T_{\text{emp}} \mid P_{\text{emp}}] \leq \left(\frac{1+\varepsilon}{1-\varepsilon}\right)^2 \mathbb{E}[T_{\text{true}}]\right] \geq 1 - \delta$$

as long as the number of samples satisfies

$$n \geq \frac{2 + \varepsilon}{q\varepsilon^2} \log\left(\frac{2k}{\delta}\right),$$

where $k = |\text{supp}(p_{\text{true}})|$ and $q = \min_{s \in \text{supp}(p_{\text{true}})} p_{\text{true}}(s)$.

Because $(\frac{1+\varepsilon}{1-\varepsilon})^2 \approx 1 + 4\varepsilon$ for small ε , this implies empirical Gittins with $\Omega(\frac{1}{\varepsilon^2} \log \frac{1}{\delta})$ samples is a $(1 + \varepsilon)$ -approximation for mean response time with probability at least $1 - \delta$.

Proof of Theorem 2.6. Let $\text{supp} = \text{supp}(p_{\text{true}})$ for brevity. By (2.3) and the Chernoff bound, for all $s \in \text{supp}$,

$$\mathbb{P}\left[\left|\frac{P_{\text{emp}}(s)}{p_{\text{true}}(s)} - 1\right| > \varepsilon\right] \leq 2 \exp\left(-\frac{\varepsilon^2 n p_{\text{true}}(s)}{2 + \varepsilon}\right).$$

Applying the union bound and recalling the definitions of q and k from the theorem statement yields

$$\mathbb{P}\left[\exists s \in \text{supp} : \left|\frac{P_{\text{emp}}(s)}{p_{\text{true}}(s)} - 1\right| > \varepsilon\right] \leq 2k \exp\left(-\frac{\varepsilon^2 n q}{2 + \varepsilon}\right).$$

Therefore, if $n \geq \frac{2+\varepsilon}{q\varepsilon^2} \log\left(\frac{2k}{\delta}\right)$, then

$$\mathbb{P}\left[\forall s \in \text{supp} : \frac{P_{\text{emp}}(s)}{p_{\text{true}}(s)} \in [1 - \varepsilon, 1 + \varepsilon]\right] \geq 1 - \delta,$$

from which the result follows by Theorem 2.4. \square

3. OPEN QUESTION: CONTINUOUS JOB SIZE DISTRIBUTIONS

The results in Section 2.2 depend upon properties such the number of support points k and the minimum point mass probability q . The bounds thus do not make sense for continuous distributions, and naive discretization arguments seem unlikely to work well, as they would lead to larger k and smaller q as the discretization gets finer, worsening the bound in Theorem 2.6.

Nevertheless, we observe in simulation that with enough samples, empirical Gittins performs well for continuous distributions, as Figure 3.1 shows for two examples: one (truncated) light-tailed in (a), and one (truncated) heavy-tailed in (b). Each histogram in Figure 3.1 shows the mean response time $\mathbb{E}[T_{\text{emp}} | P_{\text{emp}}]$ for 100 different randomly sampled empirical distributions P_{emp} .

In all cases, we simulate the empirical Gittins policy with FCFS tie-breaking. However, the empirical Gittins rank function is undefined for ages greater than the largest observed sample. We therefore use a “preemptive last-come-first-served (PLCFS) fallback” [?]: jobs with age greater than the largest observed sample have priority lower than all other jobs, and they are prioritized from latest to earliest arrival time.

References

- [1] Nikhil Bansal, Bart Kamphorst, and Bert Zwart. 2018. Achievable Performance of Blind Policies in Heavy Traffic. *Mathematics of Operations Research* 43, 3, 949–964.
- [2] Luca Becchetti and Stefano Leonardi. 2004. Nonclairvoyant Scheduling to Minimize the Total Flow Time on Single and Parallel Machines. *J. ACM* 51, 4, 517–539.
- [3] John C. Gittins, Kevin D. Glazebrook, and Richard R. Weber. 2011. *Multi-Armed Bandit Allocation Indices* (2 ed.). Wiley, Chichester, UK.
- [4] Bala Kalyanasundaram and Kirk R. Pruhs. 2003. Minimizing Flow Time Nonclairvoyantly. *J. ACM* 50, 4, 551–567.

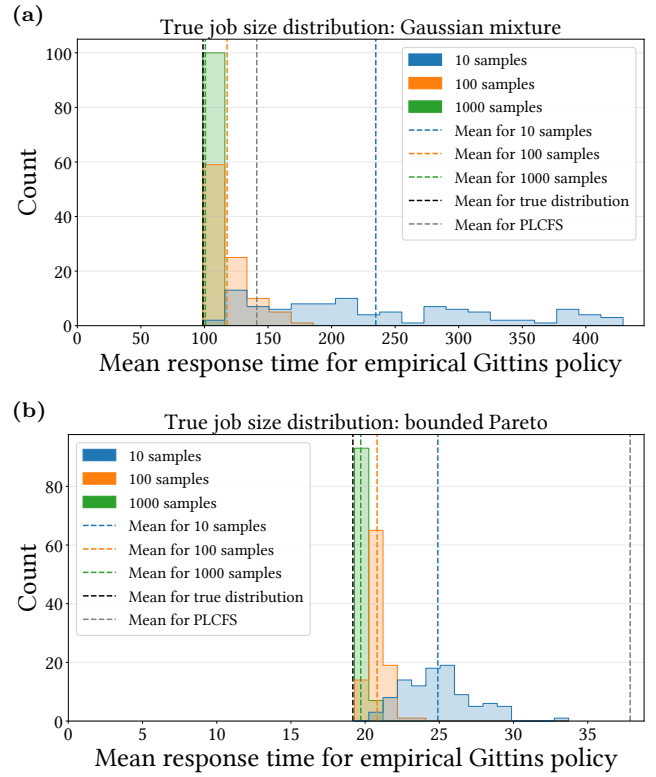


Figure 3.1. Simulated samples of $\mathbb{E}[T_{\text{emp}} | P_{\text{emp}}]$ for empirical Gittins policies with 10 (blue), 100 (orange), and 1000 (green) samples, each with 100 trials. Dotted lines indicate the average over the 100 P_{emp} PMFs across all trials. Also shown are mean response times for the true Gittins policy (black dotted line, left) and PLCFS (gray dotted line). Job size distributions are (a) a “1-6-14 Gaussian mixture”, which is a Gaussian mixture with means 1, 6, 14 (weight 1/3 each), standard deviations all 1/2, and support truncated to $[0, 16]$; and (b) a bounded Pareto distribution with shape parameter 2.1 and support $[1, 50]$. All empirical mean response times were computed via simulation for 80,000 busy periods (all standard errors $< 4\%$).

- [5] Subhashini Krishnasamy, Rajat Sen, Ramesh Johari, and Sanjay Shakkottai. 2020. Learning Unknown Service Rates in Queues: A Multiarmed Bandit Approach. *Operations Research* 69.
- [6] Ziv Scully. 2022. *A New Toolbox for Scheduling Theory*. Ph. D. Dissertation. Carnegie Mellon University, Pittsburgh, PA.
- [7] Ziv Scully and Mor Harchol-Balter. 2021. The Gittins Policy in the M/G/1 Queue. In *19th International Symposium on Modeling and Optimization in Mobile, Ad Hoc, and Wireless Networks (WiOpt 2021)*. IFIP, Philadelphia, PA, 248–255.
- [8] Ziv Scully, Mor Harchol-Balter, and Alan Scheller-Wolf. 2018. SOAP: One Clean Analysis of All Age-Based Scheduling Policies. *Proceedings of the ACM on Measurement and Analysis of Computing Systems* 2, 1, Article 16, 30 pages.
- [9] Zixian Yang, R. Srikant, and Lei Ying. 2023. Learning While Scheduling in Multi-Server Systems With Unknown Statistics: MaxWeight with Discounted UCB. In *Proceedings of The 26th International Conference on Artificial Intelligence and Statistics (Proceedings of Machine Learning Research, Vol. 206)*. PMLR, 4275–4312.
- [10] Yueyang Zhong, John R. Birge, and Amy R. Ward. 0. Learning to Schedule in Multiclass Many-Server Queues with Abandonment. *Operations Research* 0, 0, null.