Scheduling jobs with unknown sizes: Beyond SRPT

Nicolas Gast[1], Bruno Gaujal[1] and Adrien Obrecht[2],
1. Ghost, Inria Grenoble
2. ENS of Lyon

Frejus, March 2026

“Un coup de dés jamais n’abolira le hasard,
Quand bien même lancé dans des circonstances éternelles”

Motivation

We consider the problem of scheduling jobs on a single processor system to minimize their completion times.

This problem may look academic and idealized. However, it has attracted considerable interest, and the proposed solutions are both elegant and illuminating in identifying the features that are crucial for achieving low completion times.

Leverage information about job sizes to decide which job to schedule at time t.

The solutions

In the *clairvoyant case* (jobs announce their size when they arrive), the optimal scheduling policy is SRPT (Smallest Remaining Processing Time) (see Schrage. 1968).

In the *non-clairvoyant case*, (the size of the job is only revealed at completion time), the optimal scheduling policy is the Gittins index policy (Harchol-Balter, Scully 2021) when the distribution of the job sizes is known.

In the *non-clairoyant no-information case*: The size of the job is only known at completion time and the distribution of the job sizes is also unknown: Empirical Gittins Index.

Generalized Response Time

Jobs arrive in batches. The state of job i is denoted x_i and only evolves when being served (in a Markovian way).

A special state is done reached when the job is completed.

The holding cost incurred in state x is $\text{hold}(x)$, with $\text{hold}(\text{ done }) = 0$.

The goal is to minimize $\mathbb{E} \sum_{i=1}^{N} \text{hold}(X_i)$, N is the number of jobs and X_i is the state of job i in stationary regime.

Generalized Response Time

Jobs arrive in batches. The state of job i is denoted x_i and only evolves when being served (in a Markovian way).

A special state is done reached when the job is completed.

The holding cost incurred in state x is $\mathrm{hold}(x)$, with $\mathrm{hold}(\text{ done }) = 0$.

The goal is to minimize $\mathbb{E} \sum_{i=1}^{N} \mathrm{hold}(X_i)$, N is the number of jobs and X_i is the state of job i in stationary regime.

The classical case of response time minimization is:

$x_i =$ amount of execution of job i.

$\mathrm{hold}(x) = 1$ and $\mathrm{hold}(\text{ done }) = 0$.

Gittins Index

Intuition: reduce the holding cost as fast as possible.

The cost reduction rate from state x to state y is

$$\gamma(x, y) := \frac{\mathrm{hold}(x) - \mathrm{hold}(y)}{\tau(x \to y)}.$$

The average cost reduction rate from x to subset $\mathcal{Y}$ is

$$\Gamma(x, \mathcal{Y}) := \frac{\mathbb{E}\left[\mathrm{hold}(x) - \mathrm{hold}(y), y \in \mathcal{Y} \cup \{done\}\right]}{\mathbb{E}\left[\tau(x \to \mathcal{Y} \cup \{\text{ done }\})\right]}$$

The Gittins index of a job in state x is the average cost reduction rate considering all possible futures:

$$\nu(x) := \sup_{\mathcal{Y}} \Gamma(x, \mathcal{Y})$$

Gittins Index policy

Definition (Gittins index policy)

Serve the job that maximizes $\nu_i(x_i)$.

Theorem (Harchol-Balter, Scully 2021)

The Gittins index policy is optimal (minimizes the total expected cost).

The proof is technical. It first considers one job in isolation and uses a free penalization parameter to model the rest of the world and compare with the job (this is called the Gittins game). Then, it exploits the PASTA property of Poisson arrivals to conclude the proof.

Gittins Index policy

Definition (Gittins index policy)

Serve the job that maximizes $\nu_i(x_i)$.

Theorem (Harchol-Balter, Scully 2021)

The Gittins index policy is optimal (minimizes the total expected cost).

The proof is technical. It first considers one job in isolation and uses a free penalization parameter to model the rest of the world and compare with the job (this is called the Gittins game). Then, it exploits the PASTA property of Poisson arrivals to conclude the proof.

The GI Policy is quite interesting because its decision is based on indices computed for each job in isolation (scales linearly with the number of present jobs).

Also quite general: covers many cases (preemptive, non-preemptive, priorities, variable speed...)

Gittins Index for response time

Let φ be the db of S, with density f and RDF $F^c(x) = \mathbb{P}(S \geqslant x)$.

Here, the state is X_i is the amount of execution of job i. $\text{done} = S_i$, the size of job i. $\text{hold}(x) = 1$, for all x and $\text{hold}(\text{ done }) = 0$.

$$\begin{aligned}
\nu(x) &:= \sup_{y>x} \frac{\mathbb{E}(\text{hold}(x) - \text{hold}(y))}{\mathbb{E}\tau(x \to y \cup \{\text{ done }\})} \\
&= \sup_{y>x} \frac{\mathbb{P}(S \leqslant y | S > x)}{\mathbb{E}(\min(S, y) - x | S > x)} \\
&= \sup_{y>x} \frac{\int_x^y f(u)du}{\int_x^y F^c(u)du}.
\end{aligned}$$

Can be computed numerically when the distribution is known.

Remark In the deterministic case ($\mathbb{P}(S = s) = 1$), $\nu(x) = 1/(s - x)$: GI policy $\equiv$ SRPT).

What if φ is unknown?

Estimations of Gittins indices based on samples

Our approach: "*explore then commit*" .

We observe n samples of job sizes and infer empirical Gittins indices that hopefully converge as n goes to infinity.

Main challenges:

- The empirical distribution induced by n samples is made of n atoms all with equal probabilites. This is quite distinct from the actual distribution of job size.

- The Gittins index policy is not continuous w.r.t. index values.

The shift to the right

We assume φ has bounded support, its density is L-Lipschitz and lower lounded by ℓ.

We first define a shifted discretization of the measure φ: we denote by φ_Δ the shifted version of φ, where the distribution support only contains multiples of $1/\Delta$.

$$\varphi_\Delta(i/\Delta) = \varphi\left(\left(\frac{i-1}{\Delta}, \frac{i}{\Delta}\right]\right)$$
$$\varphi_\Delta(t) = 0 \text{ otherwise.}$$

Shifted empirical index: Working from samples.

The empirical distribution from n samples of the job size, $S_1, \ldots, S_n$ is $\widehat{\varphi_n} = \sum_{i=1}^{n} \frac{1}{n} \delta_{S_i}$.

- ν is the Gittins index function of φ, the distribution of the job size.
- ν_Δ is the Gittins index function of φ_Δ, the shifted discrete distribution.
- $\widehat{\nu_n}$ is the Gittins index function of $\widehat{\varphi_n}$, the empirical distribution.
- Finally, let $\widehat{\varphi_{n,\Delta}}$ be the shifted version of $\hat{\varphi}_n$. The shifted empirical index $\widehat{\nu_{n,\Delta}}$ is the Gittins index function of the distribution $\widehat{\varphi_{n,\Delta}}$.

Shifted distribution

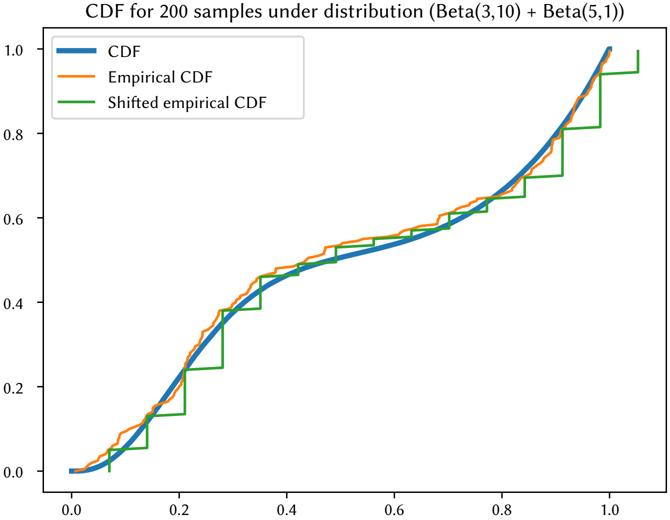

Shifted empirical index

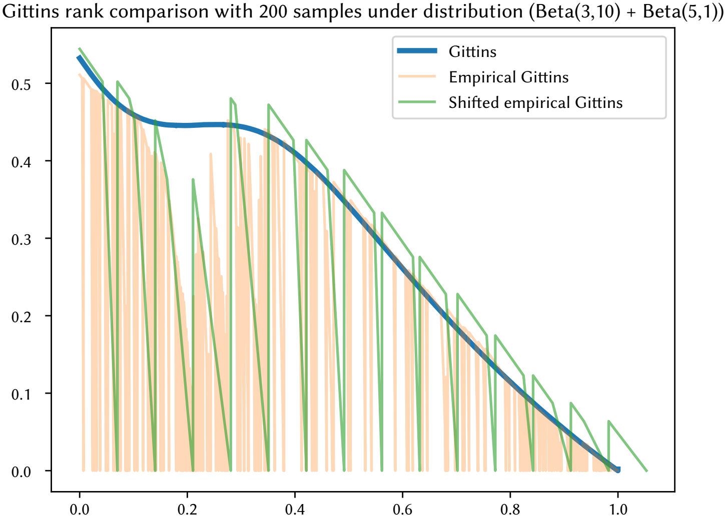

Why is the shift to the right a good thing?

What is the negative impact of using replacement of the shifted empirical Gittins instead of the real Gittins function?

The shift to the right can be seen as overestimating the size of jobs. Now if a job is selected to be executed at time t then any unexpected premature completion can only be a good thing for the response time.

On the other hand, shifting the samples to the left would have a much bigger negative impact: think of a distribution with two atoms, one at A and one at 1, where A is not a multiple of $1/\Delta$. Shifting the samples from A to the left would result in indices that become large just before time A and become small again at time A. The corresponding index policy will stop the execution of the current job before time A, missing an opportunity to complete the job (if the job is of size A).

Convergence of the shifted empirical index (part 1)

The proof has two main steps:

1. Gittins indices for the shifted distribution are close to Gittins indices.

Lemma

Let $\epsilon > 0$ *and* $\Delta \in \mathbb{N}$*. For all* $i \in \mathbb{N}$ *such that* $\frac{i}{\Delta} < 1 - \epsilon$*, then:*

$$\nu\left(\frac{i}{\Delta}\right)\frac{1 - \frac{L}{\Delta \ell \epsilon}}{1 + \frac{L}{\Delta \ell \epsilon}} \leqslant \nu_\Delta\left(\frac{i}{\Delta}\right) \leqslant \nu\left(\frac{i}{\Delta}\right).$$

Convergence of the shifted empirical indices (part 2)

2. Shifted empirical Gittins indices are close to Gittins indices for the shifted distribution.

Lemma

Let $n \in \mathbb{N}$ *be the number of samples. Assume* $n \geqslant \Delta^{2+4\alpha}$*, where* $1/2 > \alpha > 0$*. Under the foregoing assumptions, with probability greater than* $1 - 2\Delta \exp(-2n^{\alpha})$ *we get for all* $0 < i \leqslant \Delta$*,*

$$\frac{1 - \frac{n^{-\alpha(1-2\alpha)}}{L}}{1 + \frac{n^{-\alpha(1-2\alpha)}}{\ell}} \nu_{\Delta}(i/\Delta) \leqslant \widehat{\nu_{n,\Delta}}(i/\Delta) \leqslant \frac{1 + \frac{n^{-\alpha(1-2\alpha)}}{\ell}}{1 - \frac{n^{-\alpha(1-2\alpha)}}{L}} \nu_{\Delta}(i/\Delta).$$

Proof: Hoeffding inequality for the 1-Wasserstein distance between distributions:

$$\mathbb{P}(W_1(\widehat{f_{n,\Delta}}, f_{\Delta}) \geqslant A) \leqslant \sum_{i=0}^{\Delta-1} \mathbb{P}(|(S_n^i/n - \int_{i/\Delta}^{(i+1)/\Delta} f(s)ds)| \geqslant A) \leqslant 2\Delta \exp\left(-2nA^2\right).$$

Convergence of the shifted empirical index policy

A result from Scully's thesis shows that the performance of the 2 policies are close.

Theorem

By denoting K a bound on the expected response time of any work-conserving policy, there exists $\gamma_n \to 1$ and $p_n \to 0$

$$\mathbb{E}(T_{\textsc{shifted-empirical}}) \leqslant \gamma_n^2 \mathbb{E}(T_{\textsc{Gittins}}) + p_n K.$$

Proof: Bound on the response time in a M/G/1 queue:

Workload at t, $W(t)$ in the M/G/1 queue is the same under of any work conserving policy π, and so are the busy periods ($W(t) > 0$). The response time is bounded by the busy period $B(W_a + S)$ (W_a is the workload present when the job arrives and S is its size). This means $\mathbb{E}(T_\pi) \leqslant \mathbb{E}(B(W_a + S)) = \mathbb{E}(W_a + S)/(1 - \rho)$. Using PASTA, $W_a =_{db} W$. The Pollaczek - Khinchine formula gives $\mathbb{E}(W) = \dfrac{\lambda \mathbb{E}(S^2)}{2(1-\rho)}$. and $\mathbb{E}(T_\pi) \leqslant \dfrac{\mathbb{E}(S)}{1-\rho} + \dfrac{\lambda \mathbb{E}(S^2)}{2(1-\rho)^2}$.

Numerical comparisons with other policies

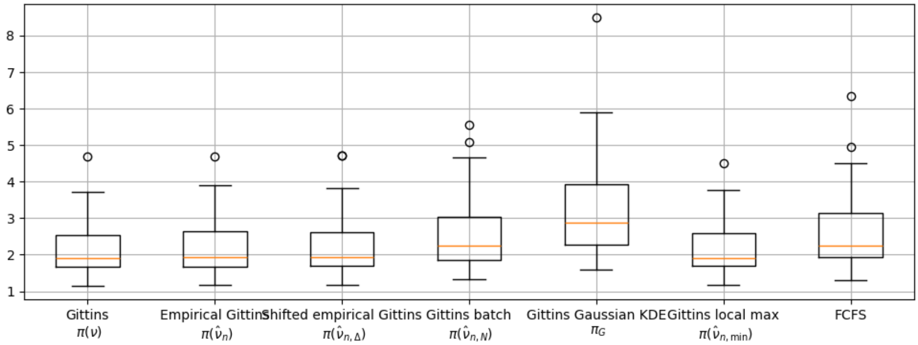

Figure: Boxplot of average service time under Pareto 5 distribution for several scheduling policies.