Performance and Stability Analysis of the Task Assignment based on Guessing Size Routing Policy

Josu Doncel

University of the Basque Country, UPV/EHU

joint work with E. Bachmat and H. Sarfati (Ben-Gurion University)

October 22, 2019

Parallel-Server Systems

K homogeneous FIFO queues

Poisson arrivals

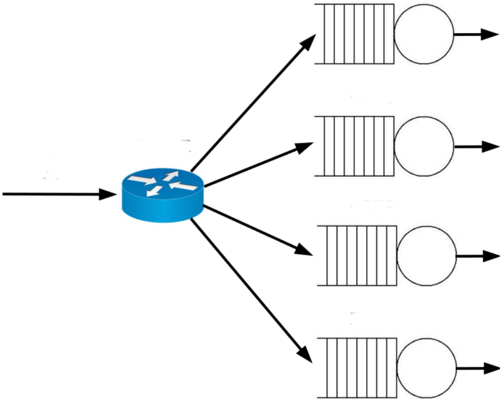

Question?

How to balance the load optimally?

Application

Heavy-tailed distribution

A small fraction of jobs make up the half of the load

Heavy-tailed distribution

A small fraction of jobs make up the half of the load

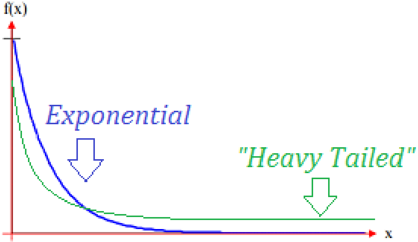

Heavy-tailed distribution

A small fraction of jobs make up the half of the load

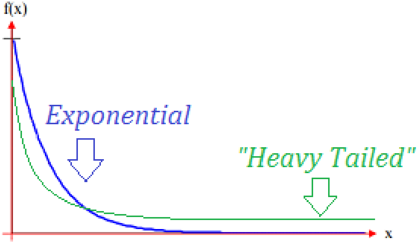

Example: Bounded Pareto (1,r,α)

$$f(s) = \frac{\alpha s^{-\alpha-1}}{1 - r^{-\alpha}}.$$

Known optimality results

JSQ: each incoming job is sent to the server with less number of jobs

Po2: pick d servers at random $\Rightarrow$ JSQ

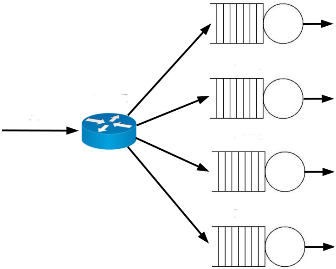

Disadvantage

Many observations $\Rightarrow$ Not practical

Open-loop Routing Policies

Heavy-tailed: bad performance

- Round-Robin
- Random Splitting

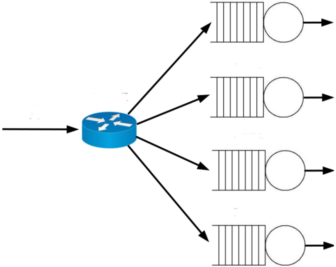

Open-loop Routing Policies

Heavy-tailed: bad performance

- Round-Robin
- Random Splitting

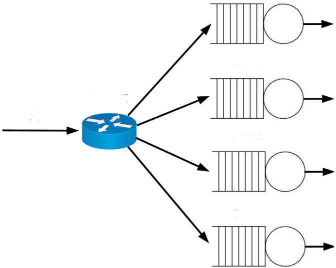

Heavy-tailed: good performance

- SITA: job duration knowledge

Open-loop Routing Policies

Heavy-tailed: bad performance

- Round-Robin
- Random Splitting

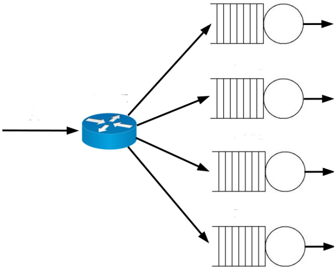

Heavy-tailed: good performance

- SITA: job duration knowledge
- Task Assignment with Guesing Size (TAGS)

TAGS Policy[1]

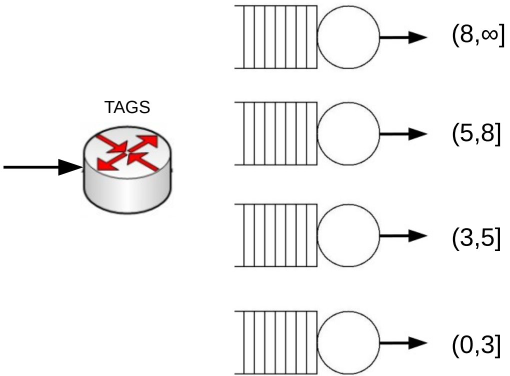

[1] M. Harchol-Balter. Task assignment with unknown duration. J. of the ACM, 2002.

TAGS Policy[1]

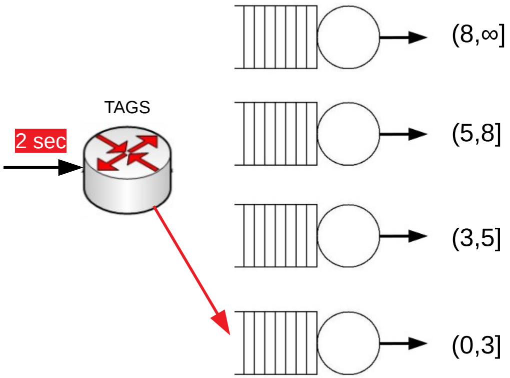

[1] M. Harchol-Balter. Task assignment with unknown duration. J. of the ACM, 2002.

TAGS Policy[1]

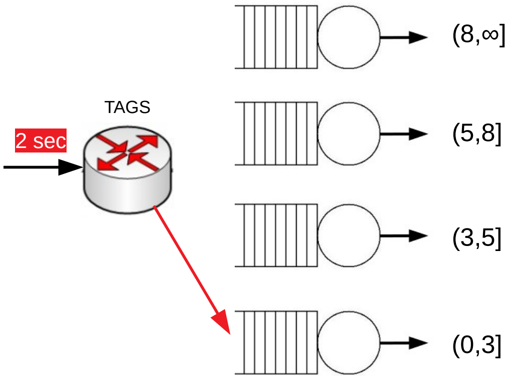

Waiting time: W_1

[1] M. Harchol-Balter. Task assignment with unknown duration. J. of the ACM, 2002.

TAGS Policy[1]

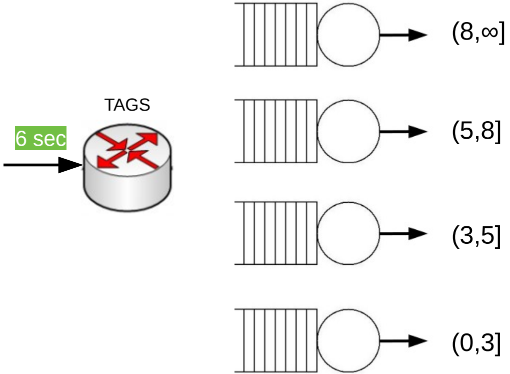

Waiting time: W_1

[1] M. Harchol-Balter. Task assignment with unknown duration. J. of the ACM, 2002.

TAGS Policy[1]

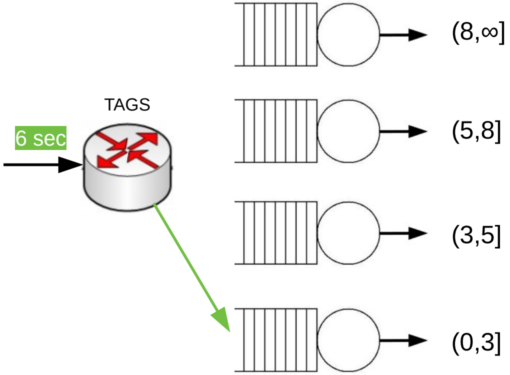

Waiting time: W_1

[1] M. Harchol-Balter. Task assignment with unknown duration. J. of the ACM, 2002.

TAGS Policy[1]

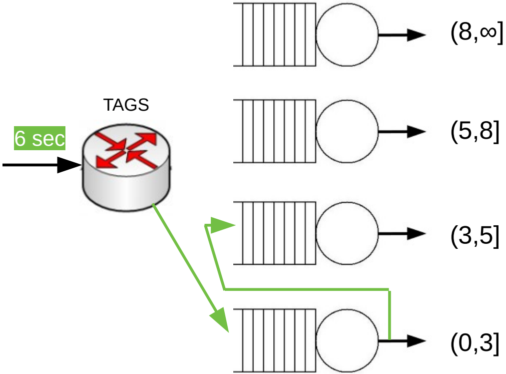

Waiting time: W_1

[1] M. Harchol-Balter. Task assignment with unknown duration. J. of the ACM, 2002.

TAGS Policy[1]

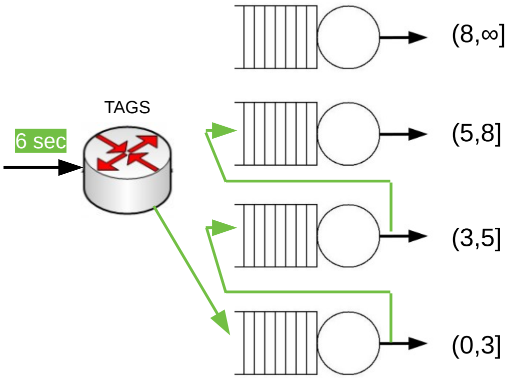

Waiting time: W_1

[1] M. Harchol-Balter. Task assignment with unknown duration. J. of the ACM, 2002.

TAGS Policy[1]

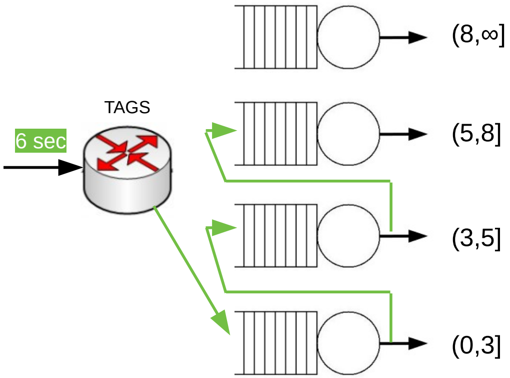

Waiting time: W_1

Waiting time: $\mathbf{W_1 + 3 + W_2 + 5 + W_3}$

[1] M. Harchol-Balter. Task assignment with unknown duration. J. of the ACM, 2002.

TAGS Policy[1]

Advantages

- No signaling
- Heavy-tail: good performance
- Durations unknown

[1] M. Harchol-Balter. Task assignment with unknown duration. J. of the ACM, 2002.

TAGS Policy[1]

Advantages

- No signaling
- Heavy-tail: good performance
- Durations unknown

Summary of contributions

1. Stability
2. Optimal performance
3. (Bounded) Pareto distribution
   - Comparison with SITA
   - Optimal performance analysis

[1] M. Harchol-Balter. Task assignment with unknown duration. J. of the ACM, 2002.

TAGS Policy

Advantages

- No signaling
- Heavy-tail: good performance
- Durations unknown

Summary of contributions

1. **Stability**
2. Optimal performance
3. (Bounded) Pareto distribution
   - Comparison with SITA
   - Optimal performance analysis

Stability

Proposition

Let $\rho = \lambda \mathbb{E}[X]$. The TAGS system is stable if and only if

$$\rho < \frac{\mathbb{E}[X]}{M(X)}$$

where $M(X) = sup_s s(1 - F(s))$.

$\Rightarrow$ Critical load $\rho_{crit}(X)$

Stability

Proposition

Let $\rho = \lambda \mathbb{E}[X]$. The TAGS system is stable if and only if

$$\rho < \frac{\mathbb{E}[X]}{M(X)}$$

where $M(X) = sup_s s(1 - F(s))$.

$\Rightarrow$ Critical load $\rho_{crit}(X)$

Proposition

Let X be a distribution in $[1, r]$

$$\rho_{crit}(X) \leq 1 + \log r.$$

Critical load

Bounded Pareto (1,r,α)

If $\alpha \neq 1$,

$$\rho_{crit} = (1 - r^{\alpha-1})(1-\alpha)^{-1/\alpha}.$$

If $\alpha = 1$

$$\rho_{crit} = \frac{r \log r}{r-1}.$$

Critical load

Bounded Pareto (1,r,α)

If $\alpha \neq 1$,

$$\rho_{crit} = (1 - r^{\alpha-1})(1-\alpha)^{-1/\alpha}.$$

If $\alpha = 1$

$$\rho_{crit} = \frac{r \log r}{r-1}.$$

Tight distribution

$f(x) = 1/x^2$, for $x \in [1, r]$

$\Rightarrow$ Dirac delta at r with mass r^{-1}

TAGS Policy

Advantages

- No signaling
- Heavy-tail: good performance
- Durations unknown

Summary of contributions

1. Stability
2. **Optimal performance**
3. (Bounded) Pareto distribution
    - Comparison with SITA
    - Optimal performance analysis

Bound of the Optimal Performance

Proposition

Let s^{que} be the vector of cutoffs that minimizes the maximum mean queue length of the servers. Then, in a system with h hosts,

$$\mathbb{E}[W(s^{que})] \leq h\mathbb{E}[W^*] + \mathbb{E}[X](h-1)$$

$\Rightarrow$ Upper-bound of $\mathbb{E}[W(s^{que})]$.

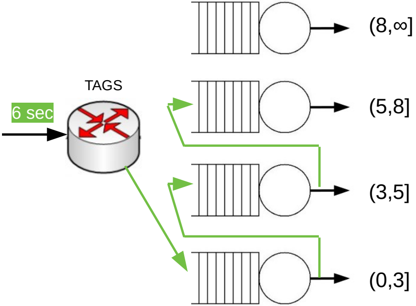

TAGS Policy

Advantages

- No signaling
- Heavy-tail: good performance
- Durations unknown

Summary of contributions

1. Stability
2. Optimal performance
3. **(Bounded) Pareto distribution**
    - Comparison with SITA
    - Optimal performance analysis

Bounded Pareto (1,r,α)

Assumption

Poisson arrivals to all the servers

$\Rightarrow$ Accuracy validation (over-estimation) numerically

Bounded Pareto (1,r,α)

Assumption

Poisson arrivals to all the servers

$\Rightarrow$ Accuracy validation (over-estimation) numerically

Let $\rho = \lambda \mathbb{E}[X]$. When $r \to \infty$ and $\rho < 1$

Proposition

The mean waiting time in a TAGS system with optimal cutoffs is at most two times larger than the mean waiting time of a SITA system with optimal cutoffs.

Bounded Pareto (1,r,α)

Bounded Pareto (1,r,α)

Let $\rho = \lambda \mathbb{E}[X]$. When $r \to \infty$ and $\rho < 1$

Proposition

The mean waiting time in a TAGS system with optimal cutoffs is at most two times larger than the mean waiting time of a SITA system with optimal cutoffs.

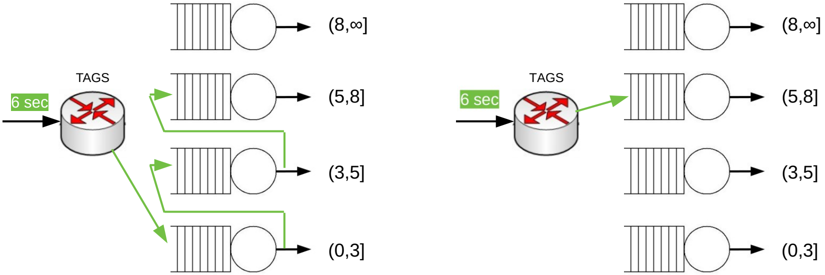

Bounded Pareto (1,r,α)

Let $\rho = \lambda\mathbb{E}[X]$. When $r \to \infty$ and $\rho < 1$

Proposition

The mean waiting time in a TAGS system with optimal cutoffs is at most two times larger than the mean waiting time of a SITA system with optimal cutoffs.

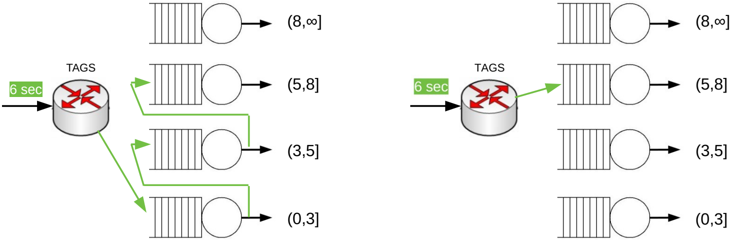

Penalty for not knowing the duration of the jobs is at most 2

Bounded Pareto (1,r,α) when $\rho > 1$

Big performance difference (stable?)

Bounded Pareto (1,r,α) when $\rho > 1$

Big performance difference (stable?)

$\tilde{h} = h - i + 1$: number of spare servers
$\Rightarrow$ i: minimum number of servers for stability

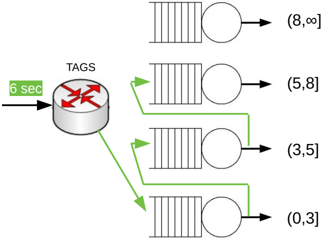

Bounded Pareto (1,r,α) when $\rho > 1$

Big performance difference (stable?)

$\tilde{h} = h - i + 1$: number of spare servers
$\Rightarrow$ i: minimum number of servers for stability

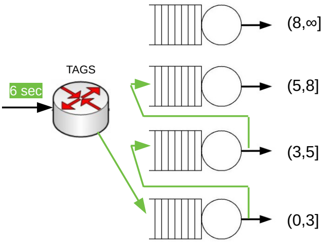

Proposition

When $r \to \infty$ and $\rho > 1$, the order of magnitude of the optimal mean waiting time of TAGS depends on $\tilde{h}$ and not on h.

T+W Policy

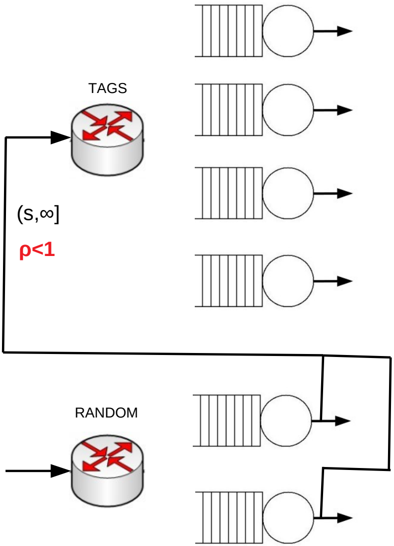

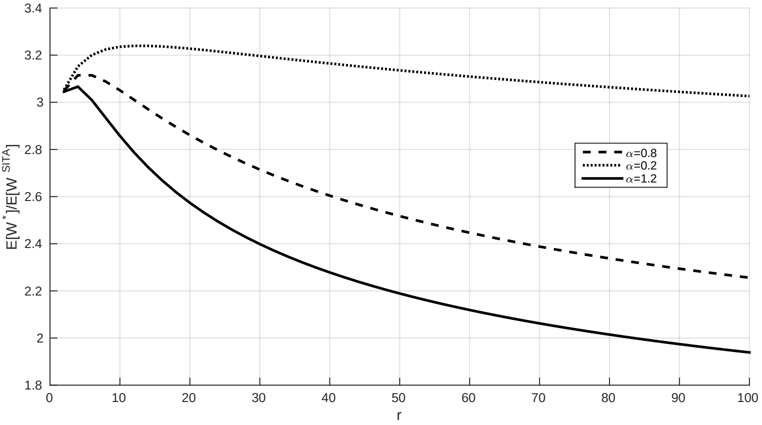
3.4
3.2
3
2.8
2.6
2.4
2.2
2
1.8
E[W*]/E[W SITA]
0
10
20
30
40
50
60
70
80
90
100
r
α=0.8
α=0.2
α=1.2

Conclusions

Advantages of TAGS

- No signaling
- Heavy-tail: good performance
- Durations knowledge not required

Our contributions

1. Stability
2. Optimal performance
3. (Bounded) Pareto distribution
    - Comparison with SITA
    - Optimal performance analysis

Conclusions

Advantages of TAGS

- No signaling
- Heavy-tail: good performance
- Durations knowledge not required

Summary of contributions

1. Stability ⇒ Critical load
2. Optimal performance
3. (Bounded) Pareto distribution
   - Comparison with SITA
   - Optimal performance analysis

Conclusions

Advantages of TAGS

- No signaling
- Heavy-tail: good performance
- Durations knowledge not required

Our contributions

1. Stability $\Rightarrow$ Critical load
2. Optimal performance $\Rightarrow$ Bounds
3. (Bounded) Pareto distribution
    - Comparison with SITA
    - Optimal performance analysis

Conclusions

Advantages of TAGS

- No signaling
- Heavy-tail: good performance
- Durations knowledge not required

Our contributions

1. Stability $\Rightarrow$ Critical load
2. Optimal performance $\Rightarrow$ Bounds
3. (Bounded) Pareto distribution
   - Comparison with SITA $\Rightarrow$ when $\rho < 1$
   - Optimal performance analysis $\Rightarrow$ when $\rho > 1$

Future Research

TODOs in our model

Non-asymptotic analysis and comparison with other policies

Future Research

TODOs in our model

Non-asymptotic analysis and comparison with other policies

Extensions to energy networks

EPN

On-off servers

Future Research

TODOs in our model

Non-asymptotic analysis and comparison with other policies

Extensions to energy networks

EPN

On-off servers

THANKS