Achievable Performance of Blind Scheduling Policies

(How to work through your to-do list?)

Nikhil Bansal[1,2] (N&O), Bart Kamphorst[1], Bert Zwart[1,2] (ST)

[1] Centrum Wiskunde & Informatica, the Netherlands
[2] Technische Universiteit Eindhoven, the Netherlands

CWI Scientific Meeting, Amsterdam

May 13, 2016

www.alamy.com - AAANG8

To-do list

- Colleagues arrive according to (random) process; rate λ
- Every colleague gives a task; mean size $\mathbb{E}[T]$
- $\rho := \lambda \mathbb{E}[T] < 1$
- Pre-empting is allowed
- No deadlines
- Minimise average waiting time $\mathbb{E}[W]$ for colleague

To-do list

- Colleagues arrive according to (random) process; rate λ
- Every colleague gives a task; mean size $\mathbb{E}[T]$
- $\rho := \lambda\mathbb{E}[T] < 1$
- Pre-empting is allowed
- No deadlines
- Minimise average waiting time $\mathbb{E}[W]$ for colleague

To-do list

First Come First Serve (FCFS)

To-do list

First Come First Serve (FCFS)

To-do list

First Come First Serve (FCFS)

To-do list

First Come First Serve (FCFS)

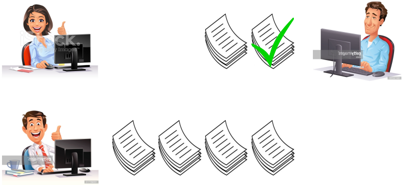

To-do list

First Come First Serve (FCFS)

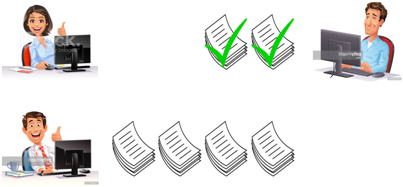

To-do list

First Come First Serve (FCFS)

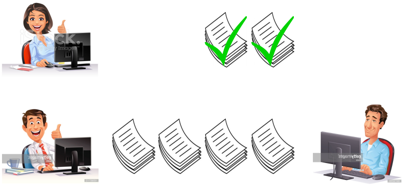

To-do list

First Come First Serve (FCFS)

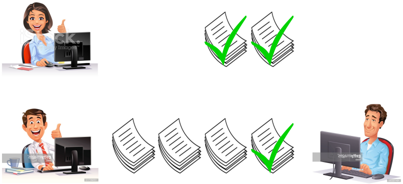

To-do list

First Come First Serve (FCFS)

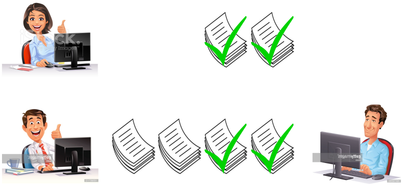

To-do list

First Come First Serve (FCFS)

To-do list

First Come First Serve (FCFS)

To-do list

First Come First Serve (FCFS)

To-do list

First Come First Serve (FCFS)

To-do list

First Come First Serve (FCFS)

To-do list

- Need to make a schedule
- Depends on size (duration) of tasks
    - Omniscient scheduler (SRPT)
    - Blind scheduler $(\mathrm{FCFS}, \mathrm{LCFS}, \mathrm{RMLF})$
- How is performance affected by knowledge?

© Can Stock Photo - csp13962659

Two CWI departments: N&O and ST

- N&O: competitive analysis
- ST: stochastic analysis

Competitive analysis	**Stochastic analysis**
Worst case performance	Average case performance
Relative benchmark	Absolute benchmark

Two CWI departments: N&O and ST

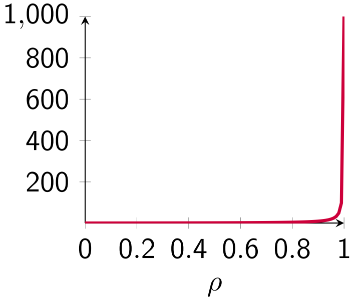

- $\mathbb{E}[W_{\mathrm{FCFS}}] = \frac{\rho c}{1-\rho}$
- Arbitrarily worse than optimal

Two CWI departments: N&O and ST

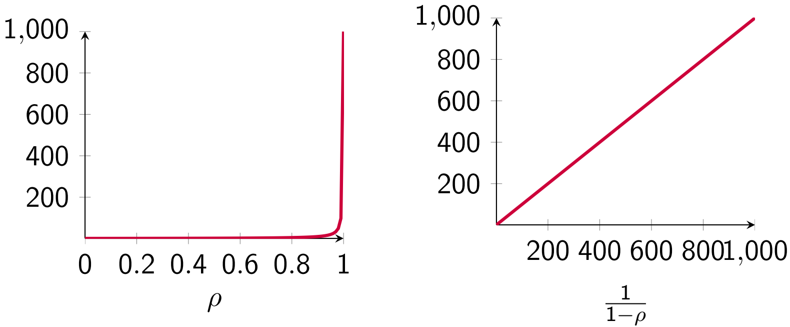

- $\mathbb{E}[W_{\mathrm{FCFS}}] = \frac{\rho c}{1-\rho}$
- Arbitrarily worse than optimal

The oblivious scheduler

- Input: sizes known
- Scheduling policy: Shortest Remaining Processing Time (SRPT)

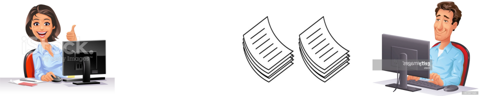

The oblivious scheduler

- Input: sizes known
- Scheduling policy: Shortest Remaining Processing Time (SRPT)

The oblivious scheduler

- Input: sizes known
- Scheduling policy: Shortest Remaining Processing Time (SRPT)

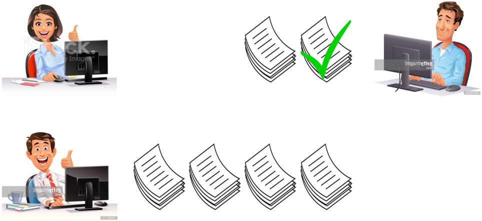

The oblivious scheduler

- Input: sizes known
- Scheduling policy: Shortest Remaining Processing Time (SRPT)

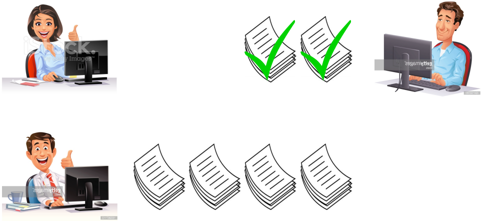

The oblivious scheduler

- Input: sizes known
- Scheduling policy: Shortest Remaining Processing Time (SRPT)

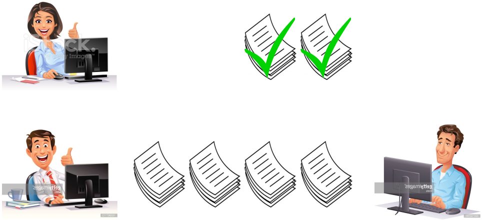

The oblivious scheduler

- Input: sizes known
- Scheduling policy: Shortest Remaining Processing Time (SRPT)

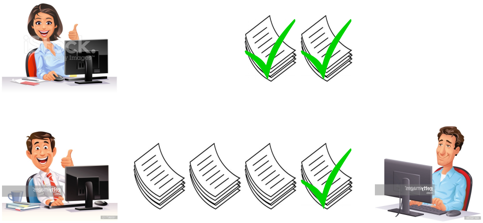

The oblivious scheduler

- Input: sizes known
- Scheduling policy: Shortest Remaining Processing Time (SRPT)

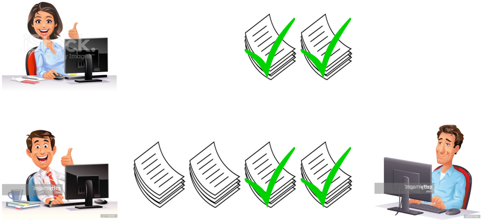

The oblivious scheduler

- Input: sizes known
- Scheduling policy: Shortest Remaining Processing Time (SRPT)

The oblivious scheduler

- Input: sizes known
- Scheduling policy: Shortest Remaining Processing Time (SRPT)

The oblivious scheduler

- Input: sizes known
- Scheduling policy: Shortest Remaining Processing Time (SRPT)

The oblivious scheduler

- Input: sizes known
- Scheduling policy: Shortest Remaining Processing Time (SRPT)

The oblivious scheduler

- Input: sizes known
- Scheduling policy: Shortest Remaining Processing Time (SRPT)

The oblivious scheduler

- Input: sizes known
- Scheduling policy: Shortest Remaining Processing Time (SRPT)

The oblivious scheduler

Performance: optimal

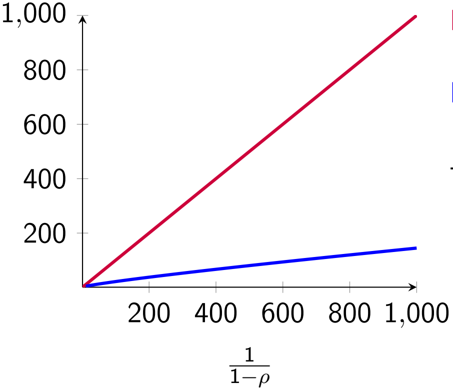

Red: $\mathbb{E}[W_{\mathrm{FCFS}}]$
$$\approx \frac{1}{1-\rho}$$

Blue: $\mathbb{E}[W_{\mathrm{SRPT}}]$
$$\approx \frac{1}{1-\rho}\frac{1}{\log\frac{1}{1-\rho}}$$

Task sizes: Exponential(1)
$$F(x) = 1 - e^{-x}$$

The oblivious scheduler

Performance: optimal

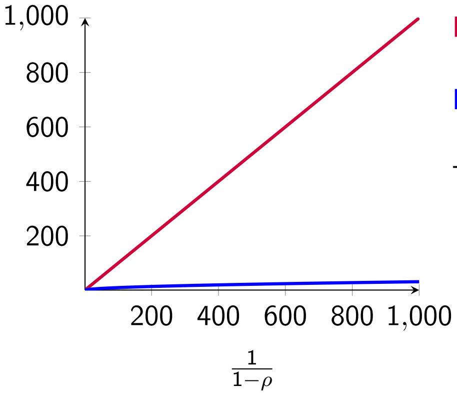

Red: $\mathbb{E}[W_{\mathrm{FCFS}}] \approx \frac{1}{1-\rho}$

Blue: $\mathbb{E}[W_{\mathrm{SRPT}}] \approx \frac{1}{\sqrt{1-\rho}}$

Task sizes: Pareto(3)

$$F(x) = 1 - x^{-3}$$

The blind scheduler

- SRPT: short tasks first
- Task sizes unknown
- Randomised Multilevel Feedback scheduling policy

The blind scheduler

- Input: sizes unknown
- Scheduling policy: Randomised Multilevel Feedback (RMLF)

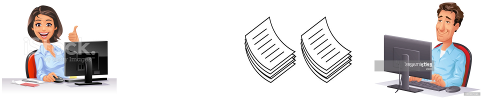

The blind scheduler

- Input: sizes unknown
- Scheduling policy: Randomised Multilevel Feedback (RMLF)

The blind scheduler

- Input: sizes unknown
- Scheduling policy: Randomised Multilevel Feedback (RMLF)

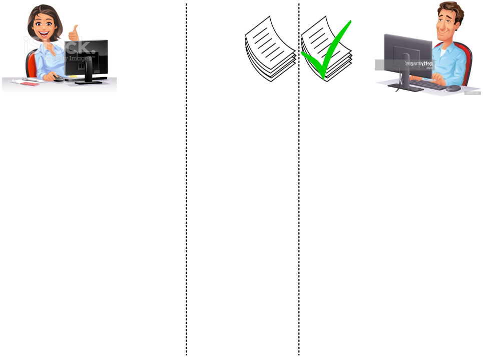

The blind scheduler

- Input: sizes unknown
- Scheduling policy: Randomised Multilevel Feedback (RMLF)

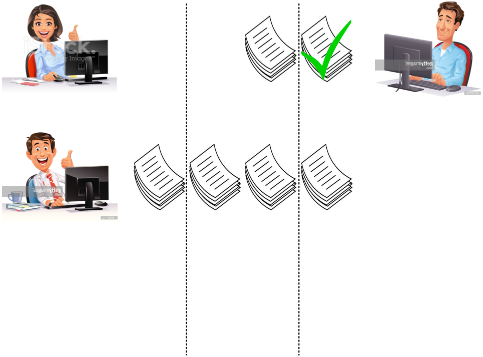

The blind scheduler

- Input: sizes unknown
- Scheduling policy: Randomised Multilevel Feedback (RMLF)

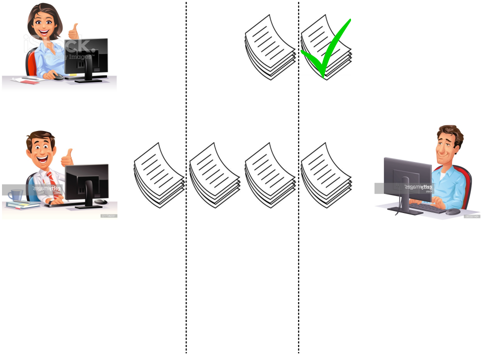

The blind scheduler

- Input: sizes unknown
- Scheduling policy: Randomised Multilevel Feedback (RMLF)

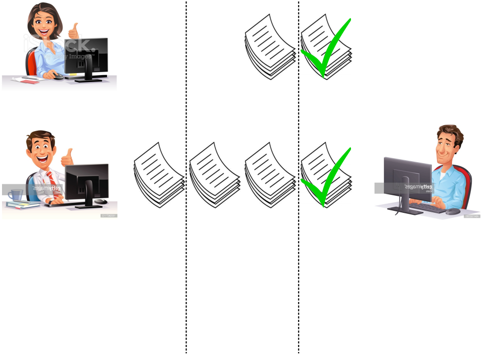

The blind scheduler

- Input: sizes unknown
- Scheduling policy: Randomised Multilevel Feedback (RMLF)

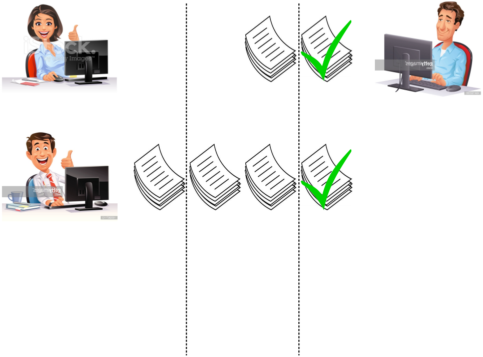

The blind scheduler

- Input: sizes unknown
- Scheduling policy: Randomised Multilevel Feedback (RMLF)

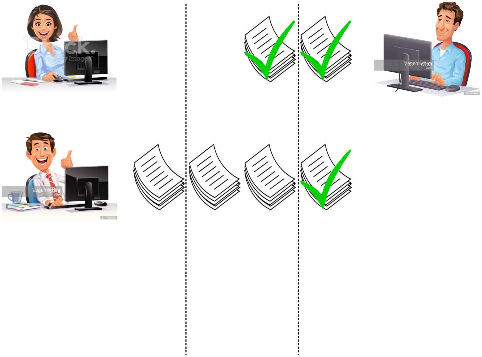

The blind scheduler

- Input: sizes unknown
- Scheduling policy: Randomised Multilevel Feedback (RMLF)

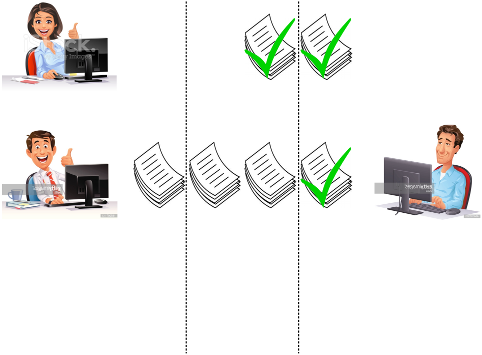

The blind scheduler

- Input: sizes unknown
- Scheduling policy: Randomised Multilevel Feedback (RMLF)

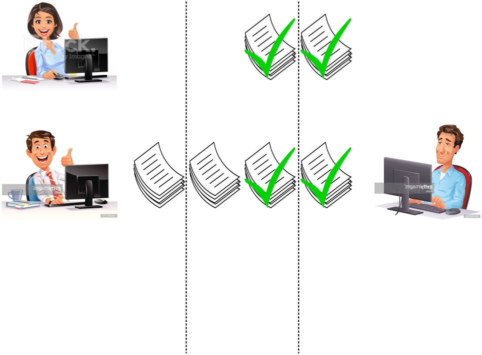

The blind scheduler

- Input: sizes unknown
- Scheduling policy: Randomised Multilevel Feedback (RMLF)

The blind scheduler

- Input: sizes unknown
- Scheduling policy: Randomised Multilevel Feedback (RMLF)

The blind scheduler

- Input: sizes unknown
- Scheduling policy: Randomised Multilevel Feedback (RMLF)

The blind scheduler

- Input: sizes unknown
- Scheduling policy: Randomised Multilevel Feedback (RMLF)

The blind scheduler

- Input: sizes unknown
- Scheduling policy: Randomised Multilevel Feedback (RMLF)

The blind scheduler

- Input: sizes unknown
- Scheduling policy: Randomised Multilevel Feedback (RMLF)

The blind scheduler

- Input: sizes unknown
- Scheduling policy: Randomised Multilevel Feedback (RMLF)

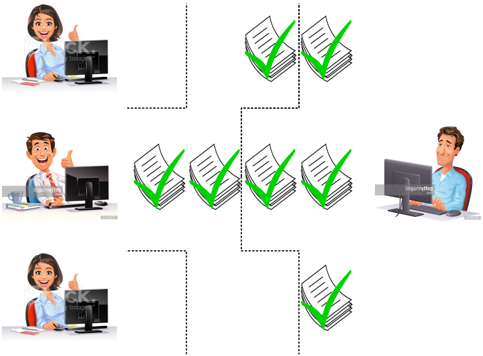

The blind scheduler

- Input: sizes unknown
- Scheduling policy: Randomised Multilevel Feedback (RMLF)
- Our theorem:

$$\mathbb{E}[W_{\mathrm{RMLF}}] \leq c \log\left(\frac{1}{1-\rho}\right) \mathbb{E}[W_{\mathrm{SRPT}}]$$

The blind scheduler

Performance: at most factor $c \log\left(\frac{1}{1-\rho}\right)$ from optimal

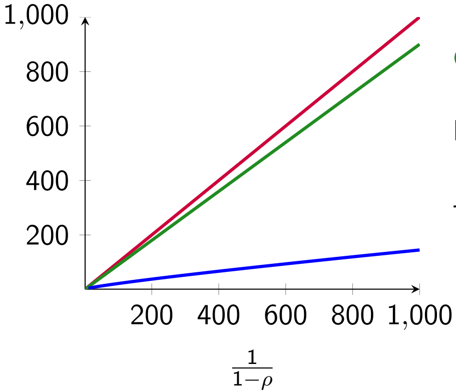

Red: $\mathbb{E}[W_{\text{FCFS}}]$

$$\approx \frac{1}{1-\rho}$$

Green: $\mathbb{E}[W_{\text{RMLF}}]$

$$\approx \frac{1}{1-\rho}$$

Blue: $\mathbb{E}[W_{\text{SRPT}}]$

$$\approx \frac{1}{1-\rho} \frac{1}{\log \frac{1}{1-\rho}}$$

Task sizes: Exponential(1)

$$F(x) = 1 - e^{-x}$$

The blind scheduler

Performance: at most factor $c \log\left(\frac{1}{1-\rho}\right)$ from optimal

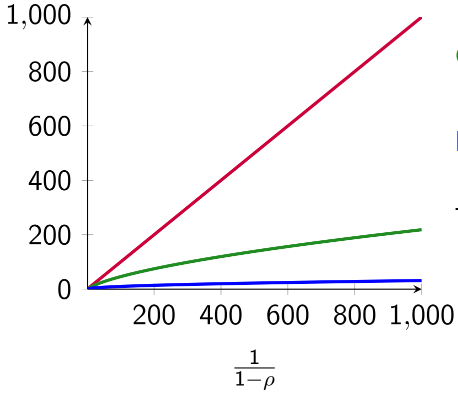

Red: $\mathbb{E}[W_{\mathrm{FCFS}}]$
$\approx \frac{1}{1-\rho}$

Green: $\mathbb{E}[W_{\mathrm{RMLF}}]$
$\approx \frac{1}{\sqrt{1-\rho}} \frac{1}{\log \frac{1}{1-\rho}}$

Blue: $\mathbb{E}[W_{\mathrm{SRPT}}]$
$\approx \frac{1}{\sqrt{1-\rho}}$

Task sizes: Pareto(3)
$F(x) = 1 - x^{-3}$

Takeaway:

- Do not apply FCFS when you're busy
- N&O: RMLF is close to optimal scheduling policy
- ST: bounds on waiting time under RMLF
- Proof needs techniques from both

To do:

- Extend concept to other models

Thank you!

N. Bansal, B. Kamphorst, B. Zwart
Achievable Performance of Blind Policies in Heavy Traffic
http://arxiv.org/abs/1512.07771

Bansal, N. (2005). On the average sojourn time under M/M/1/SRPT. *Operations Research Letters*, 33(2):195–200.

Becchetti, L. and Leonardi, S. (2001). Non-clairvoyant scheduling to minimize the average flow time on single and parallel machines. In *Proceedings of the Thirty-third Annual ACM Symposium on Theory of Computing*, STOC '01, pages 94–103, New York, NY, USA. ACM.

Kalyanasundaram, B. and Pruhs, K. R. (2003). Minimizing flow time nonclairvoyantly. *Journal of the ACM (JACM)*, 50(4):551–567.

Lin, M., Wierman, A., and Zwart, B. (2011). Heavy-traffic analysis of mean response time under shortest remaining processing time. *Performance Evaluation*, 68(10):955–966.