Machine Learning for HPC Scheduling

Raphael Y. de Camargo (UFABC - Federal University of ABC)

UFABC

Obtaining Dynamic Scheduling Policies with Simulation and Machine Learning

Danilo Carastan-Santos

Raphael Y. de Camargo

- Simulations to generate examples (features, label)
- Training of non-linear regression
- Trained and tested on synthetic and real datasets

Nominated for Best Paper (SC'17)

$$AVEbsld(T) = \frac{1}{|T|} \sum_{t \in T} \max\left(\frac{w_t + r_t}{\max(r_t, \tau)}, 1\right)$$

Name	Function
FCFS	$score(t) = s_t$
SPT	$score(t) = r_t$
WFP3	$score(t) = -(w_t/r_t)^3 \cdot n_t$
UNICEF	$score(t) = -w_t/(\log_2(n_t) \cdot r_t)$

ID	Nonlinear Function
F1	$\log_{10}(r) \cdot n + 8.70 \cdot 10^2 \cdot \log_{10}(s)$
F2	$\sqrt{r} \cdot n + 2.56 \cdot 10^4 \cdot \log_{10}(s)$
F3	$r \cdot n + 6.86 \cdot 10^6 \cdot \log_{10}(s)$
F4	$r \cdot \sqrt{n} + 5.30 \cdot 10^5 \cdot \log_{10}(s)$

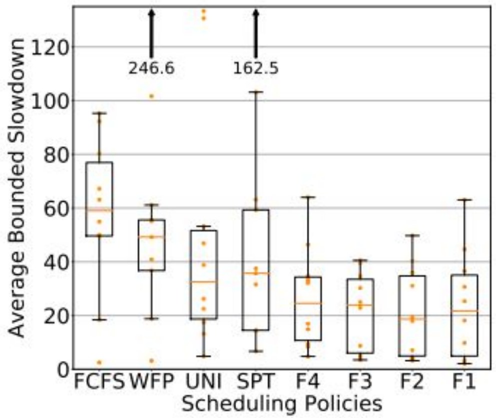

(a) Curie

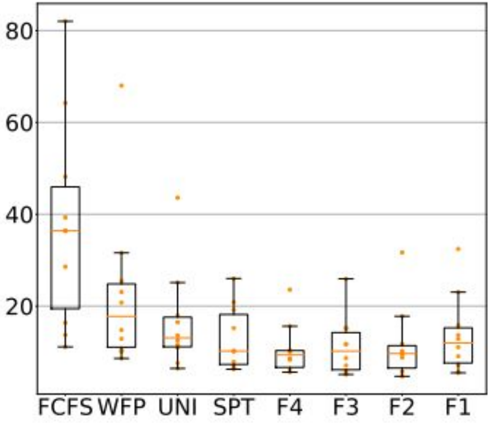

(c) SDSC Blue

One can only gain by replacing EASY Backfilling: A simple scheduling policies case study

Danilo Carastan-Santos, Raphael y de Camargo, Denis Trystram, Salah Zrigui

- SAF is difficult to beat

Best-Paper Award (CCGrid)

Name	Description	Function
FCFS	First-Come-First-Served [27]	$f(t) = r_t$
SPF	Smallest Estimated Processing Time First [28]	$f(t) = \tilde{p}_t$
SQF	Smallest Resource Requirement First	$f(t) = q_t$
SAF	Smallest Estimated "Area" First	$f(t) = \tilde{p}_t \cdot q_t$

$$AVEbsld(T) = \frac{1}{|T|} \sum_{t \in T} \max\left(\frac{w_t + r_t}{\max(r_t, \tau)}, 1\right)$$

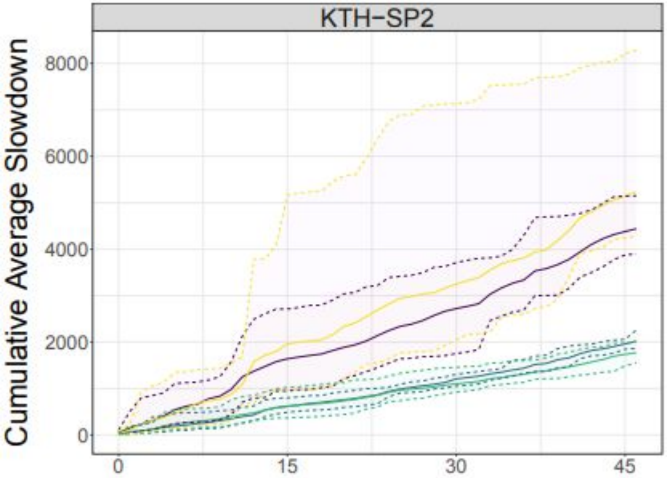

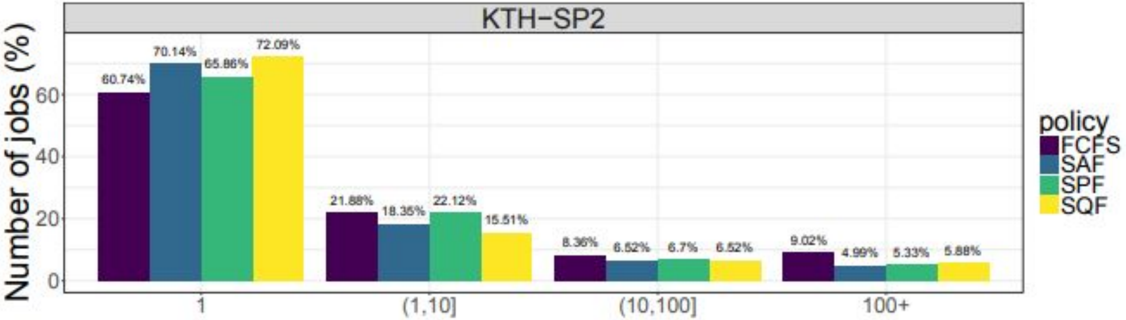

Can We Train an Agent to Perform Scheduling?

Reinforcement Learning

- Must make decisions quickly (worth only for larger tasks)
- State:
    - queue: waiting time, resource requirements, workflow graph
    - machine: active processes, idle resources
- Actions:
    - next job to execute, which policy to use
- Reward:
    - Minimize execution time, waiting time, energy consumption

Fan, Y., et al., DRAS: Deep Reinforcement Learning for Cluster Scheduling in High Performance Computing, *IEEE Transactions on Parallel and Distributed Systems*, vol 3, no 12, 2022.

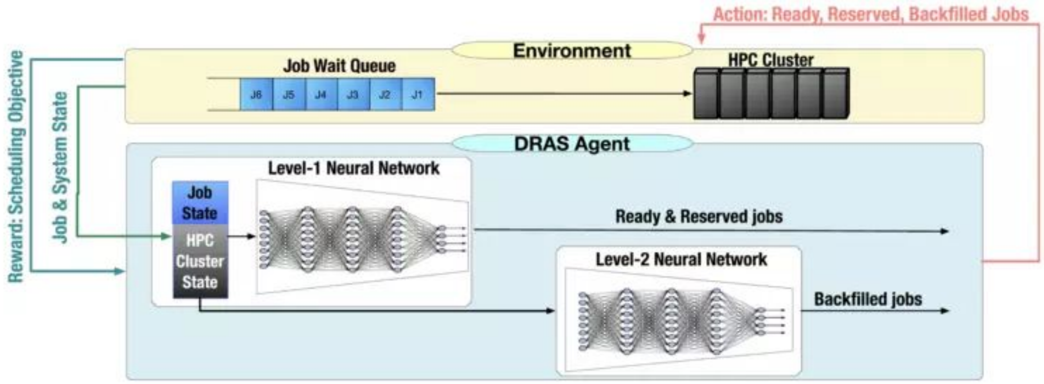

Can We Train an Agent to Perform Scheduling?

Train the algorithm using simulation

- Q-function: map state and action to reward
- Policy function: map state into action

Pre-train Q-function and/or policy function

- Using good scheduling functions (e.g., SAF)
- Fine-tuning using RL

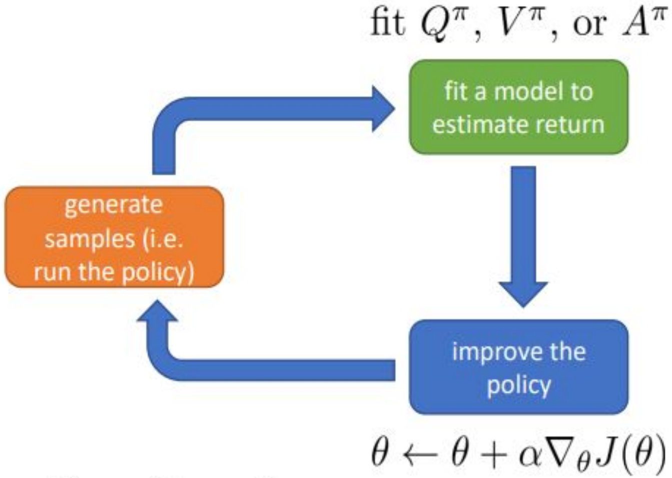

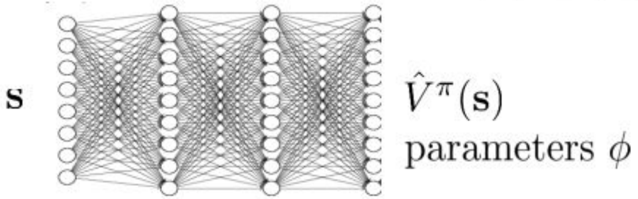

Reinforcement Learning for Last Mile Problem

Use of Demand Responsive Transit (DRT)

- Take passengers to and from Hubs
- Microtransit with 8 to 16 seats

Reinforcement Learning

- Decide on which direction to go from Hub
- Outbound trip: take passenger going toward target site
- Inbound trip: take passenger going towards the hub

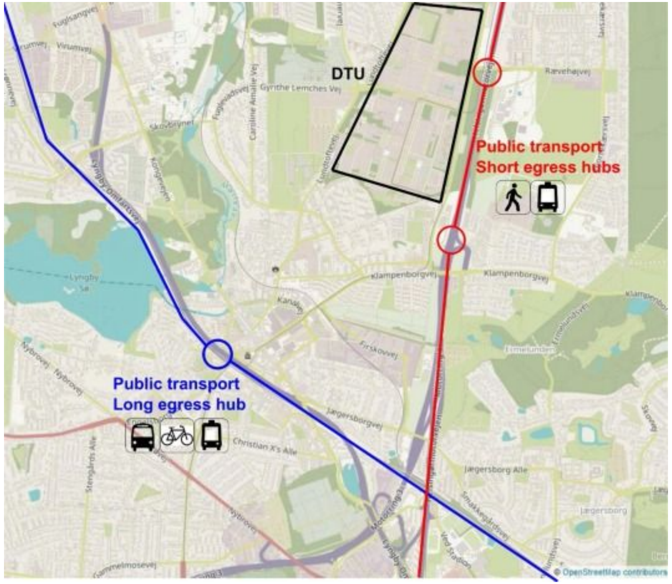

Rick Grahn, Sean Qian, Chris Hendrickson
Optimizing first- and last-mile public transit services leveraging transportation network companies (TNC) Transportation. 2023