



HAL
open science

Towards an Energy-Aware Multi-objective Scheduling Policy for Server-based Edge-Cloud Continuum

Anderson Andrei Da Silva, Yiannis Georgiou, Michael Mercier, Grégory Mounié,
Denis Trystram

► **To cite this version:**

Anderson Andrei Da Silva, Yiannis Georgiou, Michael Mercier, Grégory Mounié, Denis Trystram. Towards an Energy-Aware Multi-objective Scheduling Policy for Server-based Edge-Cloud Continuum. COMPAS 2023 - Conférence francophone d'informatique en Parallélisme, Architecture et Système, Jul 2023, Annecy, France. pp.1-21, 2023. <hal-04244625>

HAL Id: hal-04244625

<https://inria.hal.science/hal-04244625v1>

Submitted on 16 Oct 2023

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



Distributed under a Creative Commons CC BY 4.0 - Attribution - International License

Towards an Energy-Aware Multi-objective Scheduling Policy for Server-based Edge-Cloud Continuum

Compas 2023

Anderson Andrei Da Silva^{1 2}, Y. Georgiou², M. Mercier², G. Mounié¹, D. Trystram¹.

¹Université Grenoble Alpes, ²Ryax Technologies

Contact: anderson-andrei.da-silva@inria.fr, silva.andersonandrei@gmail.com

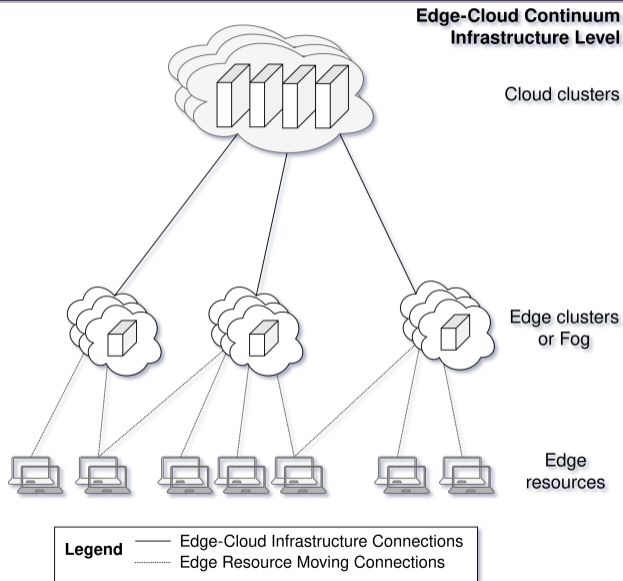
<https://andersonandrei.github.io/>

4-7 of July of 2023

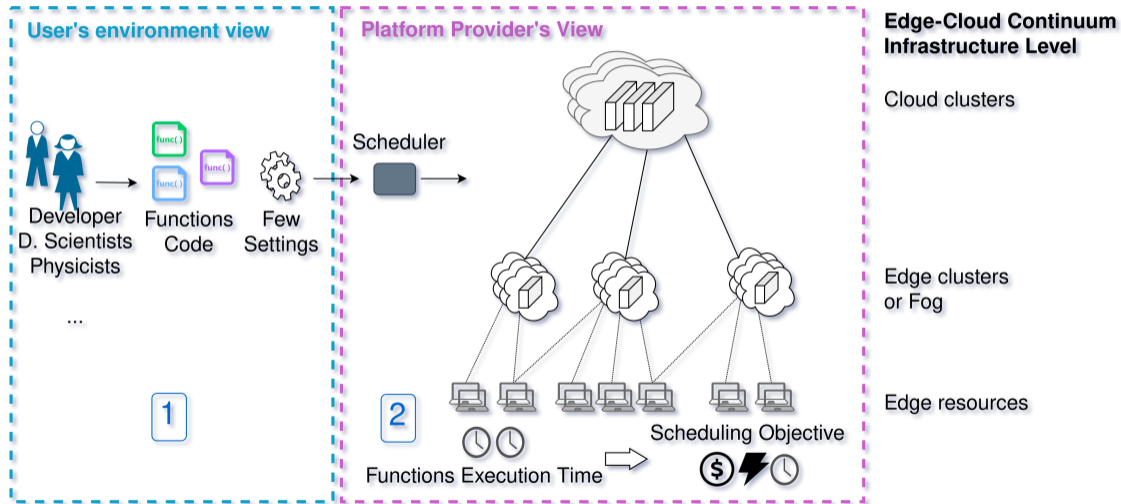


PHYSICS

The edge-cloud continuum



Serverless Computing and the Edge-Cloud Continuum



Towards a Multi-objective Scheduling Policy for Serverless-based Edge-Cloud Continuum

Towards a Multi-objective Scheduling Policy for Serverless-based Edge-Cloud Continuum

The 23rd IEEE/ACM International Symposium on Cluster, Cloud and Internet Computing
Bangalore, India — May 1-4, 2023

L. Angelelli¹, **Anderson Andrei Da Silva**^{1 2}, Y. Georgiou²,
M. Mercier², G. Mounié¹, D. Trystram¹.

¹Université Grenoble Alpes, ²Ryax Technologies

<https://andersonandrei.github.io/>


<https://gitlab.com/andersonandrei/foa-a-multi-objective-scheduling-policy-for-serverless>

Thanks to the Intelligence of the Multi-disciplinary Institute on Artificial Intelligence MIAI at Grenoble Alpes (ANR-19-P3IA-0003), to the European Union's Project H2020 PHYSICS (GA 101017047), and to the Grid'5000 testbed (<https://www.grid5000.fr>).




Reproducible
Artifacts


Motivation 1 and a multi-objective approach



FOA: Function Orchestration Algorithm



FOA: Function Orchestration Algorithm



Assumptions:

- 1) Functions and environments are independent and known in advance;
- 2) Their cost and processing time on each machine are known;
- 3) The dependency between functions and environments is known.

*The *Minimum Cost Integral Matching* phase uses the Dual-Approx. Algorithm of Shmoys and Tardos [D. B. Shmoys and E. Tardos, "An approximation algorithm for the generalized assignment problem," Math. Program., vol. 62, no. 1-3, p.461-474, Feb. 1993].

Experimental Protocol

- Adaptation of the benchmark **FunctionBench** ¹,
- Deployment of **OpenWhisk** ² on **GRID5000** ³,
 - Measurement and calibration of results.
- Evaluation of scheduling policies FOA, K8S ImageLocality (baseline):
 - In a **simulated environment**,
 - On top of **Batsim/ Simgrid** [H.Casanova et.al, 2014].

¹ <https://github.com/kmu-bigdata/serverless-faas-workbench> ,

² <https://openwhisk.apache.org/>,

³ <https://www.grid5000.fr> .

Parameters	Values
Workload Size	200, 600, 1000
Platform Size	100, 300, 500
Heterogeneity Level	3, 5, 7
Scheduling Policies	FOA, <i>K8S ImageLocality</i>
Random Seeds	30

Table 1 Design of Experiments. Total of 1620 experiments.

Design of Experiments


Parameters	Values
Workload Size	200, 600, 1000
Platform Size	100, 300, 500
Heterogeneity Level	3, 5, 7
Scheduling Policies	FOA, <i>K8S ImageLocality</i>
Random Seeds	30

Table 1 Design of Experiments. Total of 1620 experiments.

Function	Input Values	Unit
Chameleon	$2 \cdot 10^3$, $3 \cdot 10^3$	Matrix size
Float operation	$10 \cdot 10^7$, $20 \cdot 10^7$, $30 \cdot 10^7$	Operations
Image processing	60, 40	Image size in MB
Linpack	$5 \cdot 10^3$, $6 \cdot 10^3$	Matrix size
Matrix Multiplication	$3 \cdot 10^3$, $4 \cdot 10^3$	Matrix size
Model Training	50, 100	Dataset size
Pyaes	$3 \cdot 10^3$, $4 \cdot 10^3$, $5 \cdot 10^3$	Length of message
Video Processing	30, 50, 100	Video size in MB

Table 2: Functions adapted from FunctionBench.

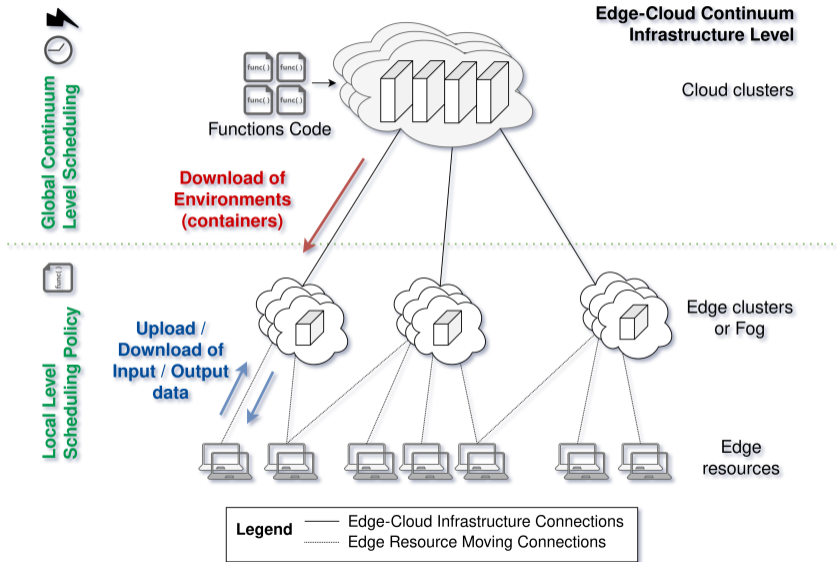
Experimental results: Linear program trade-off



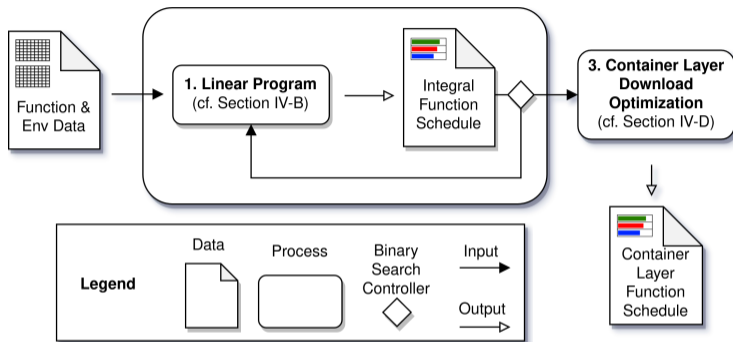
Conclusions - Part 1

- Cloud **greedy algorithms** may **not** profit from heterogeneity,
- FOA outperforms our baseline for both **data transfers** and **makespan** criteria, in addition to the **system utilization** by up to **two orders of magnitudes**,
- FOA minimizes **cold start delays** and **speed up** functions' **execution time**,
- However, FOA is very time-consuming.

Motivation 2 = Motivation 1 + Energy Consumption





FOA-e: Function Orchestration Algorithm




Assumptions:

- 1) Functions and environments are independent and known in advance;
- 2) Their cost, energy consumption and processing time on each machine are known;
- 3) The dependency between functions and environments is known.


Energy Consumption on Serverless Platforms



Energy Consumption on Serverless Platforms




Function — linpack_5000 — linpack_6000



Function — chameleon_2000 — chameleon_4000

Experimental results: FOA-e Energy Consumption



Conclusions - Part 2

- Cloud **greedy algorithms** may **not** profit from heterogeneity,
- FOA-e outperforms our baseline for both **data transfers** and **energy consumption** criteria, in addition to the **system utilization**,
- FOA-e minimizes **cold start delays** and **speed up** functions' **execution time**,
- FOA-e has a scheduling time in order of seconds.

Conclusions - Part 2

- Cloud **greedy algorithms** may **not** profit from heterogeneity,
- FOA-e outperforms our baseline for both **data transfers** and **energy consumption** criteria, in addition to the **system utilization**,
- FOA-e minimizes **cold start delays** and **speed up** functions' **execution time**,
- FOA-e has a scheduling time in order of seconds.

Work in progress:

- We have implemented FOA-e to be executed in K8S,
- We have contributed to an Open-Source project called Kepler that measures the energy consumption of k8s components (containers, pods, nodes, etc).

- To improve FOA's **model** and try other linear programs' **solvers**,
- To study applications that can be modeled as **workflows** of serverless functions,
- To continue the investigations towards the reduction of energy consumption on serverless,
- To add both schedulers in the PHYSICS platforms,
- To include and use Kepler for a continuous measurement of the energy consumption,
- To perform experiments of k8s platforms to compare with the simulations and with default k8s scheduler.

Thank you for your attention! Any question?

Towards an Energy-Aware Multi-objective Scheduling Policy for Server-based Edge-Cloud Continuum

Compas 2023, Annecy - France

Towards a Multi-objective Scheduling Policy for Serverless-based Edge-Cloud Continuum

The 23rd IEEE/ACM International Symposium on Cluster, Cloud and Internet Computing

Bangalore, India — May 1-4, 2023

L. Angelelli¹, **Anderson Andrei Da Silva**^{1 2}, Y. Georgiou², M. Mercier², G. Mounié¹, D. Trystram¹.

¹Université Grenoble Alpes, ²Ryax Technologies <https://andersonandrei.github.io/>
<https://gitlab.com/andersonandrei/foa-a-multi-objective-scheduling-policy-for-serverless>

Thanks to the Intelligence of the Multi-disciplinary Institute on Artificial Intelligence MIAI at Grenoble Alpes (ANR-19-P3IA-0003), to the European Union's Project H2020 PHYSICS (GA 101017047), and to the Grid'5000 testbed (<https://www.grid5000.fr>).



Reproducible
Artifacts



ryax
technologies



MIAI
Grenoble Alpes



PHYSICS