A Two-level Multi-objective Energy-Aware Scheduling Policy for Serverless-based Edge-Cloud Continuum

Anderson Andrei Da Silva[1,2], Yiannis Georgiou[2], Michael Mercier[2], Gregory Mounié[1], Denis Trystram[1]

[1] Univ. Grenoble Alpes, CNRS, Inria, Grenoble INP, LIG, Grenoble, France
[2] Ryax Technologies, Lyon, France

anderson-andrei.da-silva@inria.fr, yiannis.georgiou@ryax.tech, michael.mercie@ryax.tec, gregory.mounie@imag.fr, denis.trystram@inria.fr

Keywords : *Scheduling of functions, Multi-objective, Energy Consumption, Serverless Computing, Edge-cloud Continuum*

In the last years, we have been witnessing the emergence of edge computing to complement the well-established technology of cloud computing. It led to the paradigm of an edge-cloud computing continuum [1] which can better address the challenges brought by the new generation of applications. These applications involve massive data, a much shorter time for action, along with security and privacy vulnerabilities, that the cloud alone could not handle. The edge-cloud continuum, comprises a multi-layer architecture where cloud clusters, edge clusters (also known as fog) and edge resources are interconnected. This continuum comprises the heterogeneity of resources while proposing common abstractions and mechanisms per cluster, enabling a unified control. It provides a platform with heterogeneous resources to better address the needs of modern applications, but it also introduces further complexity, particularly in resource management and scheduling where it requires more efforts to deal with such heterogeneous multi-layered infrastructures. More specifically, standard cloud scheduling policies are based on greedy algorithms that do not efficiently handle platforms' heterogeneity and do not optimize data transfers. In parallel, serverless technology has gained popularity as the new way to program and deploy applications on clouds [2]. Serverless computing enables lightweight deployments of small functions with a short execution time. It is a perfect fit for the edge-cloud continuum because it allows quick adaptations to any move toward the edge level while keeping the applications' footprints low. However, serverless also brings new challenges such as managing heterogeneous platforms and applications that deal with massive data, as well as deploying complex software environments such as the ones needed for machine learning and artificial intelligence applications. Solving such challenges brought by the current generation of applications and data producers relies on managing heterogeneous platforms [3].

We are focused on reducing resource usage, cold-start delays and energy consumption through a multi-objective scheduling policy designed for the continuum. Whenever a function is triggered, its container image is required. If the function already has its container deployed, it is called a warm start-up, and the function will be shortly initialized. If not, it is called a cold start-up (or cold start delay). Shahrad et al. showed that the deployment of containers can cause an overhead of up to 20x of the platform slowdown, with cold start delays up to 10x longer [4]. Even so, containers are composed of layers, and due to this composition, containers can profit from a sharing mechanism of layers (or caching of layers), which can be used to speed up their deployment [5]. Although energy is a hot societal topic, there are only few tools that accurately measure its consumption on cloud platforms. This task becomes even more challenging on serverless platforms due to its nature of providing resources to the user with none or little information and control. All of that makes it impossible to access physical sensors on the machines, and consequently, makes it a hard task to measure the energy consumption of such platforms. Solutions are present in the state of the art with a few different tools and projects such as Kepler [6]. However, they are mostly used relying on estimations, and it is not of our knowledge any solution that accurately performs such measurements of energy consumption of serverless functions through physical sensors. Therefore, we investigate solutions for accessing accurate measurements relying on physical sensors and bare-metal platforms. Finally, due to our findings that the study of serverless platforms in the edge-cloud continuum is not an easy task, we are focused on providing reproducible techniques and accessible data, through controlled and reproducible environments. Combining GRID5000, OpenWhisk, FunctionBench, and Batsim/ Simgrid, we build a controlled and reproducible environment by having respectively, a bare-metal infrastructure, a serverless platform, a set of serverless functions, and a batch scheduler simulator.

We propose a multi-objective scheduling policy, called *FOA*-energy, that enables the allocation of batches of serverless functions. The scheduling policy aims to minimize energy consumption, makespan, and data transfers. We show that considering the heterogeneity of platforms at the scheduling phase impacts a lot in the efficiency of serverless platforms in the edge-cloud continuum. *FOA*-energy outperforms the baseline, Kubernetes *ImageLocality*, for energy consumption, makespan, data transfers, and number of machines used by up to three orders of magnitude. In addition, we propose a) an evaluation methodology, with reproducible setup and artifacts, that can be extended to other serverless-based heterogeneous edge-cloud platforms [7]; b) the study of energy consumption on serverless platforms, though bare-metal infrastructures with physical sensors, or cloud-designed tools, going one step further in this direction; and c) A serverless-function energy-based benchmark that evaluates several combinations of heterogeneous machines.

References

[1] Daniel Rosendo, Alexandru Costan, Patrick Valduriez, and Gabriel Antoniu. "Distributed intelligence on the Edge-to-Cloud Continuum: A systematic literature review". In: *J. Parallel Distributed Comput.* 166 (2022), pp. 71–94 (cit. on p. 1).

[2] Eric Jonas, Johann Schleier-Smith, Vikram Sreekanti, Chia-Che Tsai, Anurag Khandelwal, Qifan Pu, Vaishaal Shankar, Joao Carreira, Karl Krauth, Neeraja Yadwadkar, Joseph E. Gonzalez, Raluca Ada Popa, Ion Stoica, and David A. Patterson. "Cloud Programming Simplified: A Berkeley View on Serverless Computing". en. In: *arXiv:1902.03383 [cs]* (Feb. 2019). arXiv: 1902.03383 (cit. on p. 1).

[3] Luc Angelelli, Anderson Andrei da Silva, Yiannis Georgiou, Michael Mercier, Gregory Mounié, and Denis Trystram. "Towards a Multi-objective Scheduling Policy for Serverless-based Edge-Cloud Continuum". In: *2023 IEEE/ACM 23rd International Symposium on Cluster, Cloud and Internet Computing (CCGrid)*. 2023, pp. 485–497 (cit. on p. 1).

[4] Mohammad Shahrad, Jonathan Balkind, and David Wentzlaff. "Architectural Implications of Function-as-a-Service Computing". In: *Proceedings of the 52nd Annual IEEE/ACM International Symposium on Microarchitecture*. MICRO '52. Columbus, OH, USA: Association for Computing Machinery, 2019, pp. 1063–1075 (cit. on p. 1).

[5] Thomas Rausch, Alexander Rashed, and Schahram Dustdar. "Optimized container scheduling for data-intensive serverless edge computing". In: *Future Generation Computer Systems* 114 (Aug. 2020) (cit. on p. 1).

[6] Marcelo Amaral, Huamin Chen, Tatsuhiro Chiba, Rina Nakazawa, Sunyanan Choochotkaew, Eun Kyung Lee, and Tamar Eilam. "Kepler: A Framework to Calculate the Energy Consumption of Containerized Applications". In: *2023 IEEE 16th International Conference on Cloud Computing (CLOUD)*. 2023, pp. 69–71 (cit. on p. 2).

[7] *FOA Energy - GitLab Repository*. `https://gitlab.com/andersonandrei/foa-energy-journal`. Accessed: November-2023. 2023 (cit. on p. 2).