Proceedings Paper:

Djemame, K orcid.org/0000-0001-5811-5263 (2021) Energy Efficiency in Edge Environments: a Serverless Computing Approach. In: Lecture Notes in Computer Science. 18th International Conference on the Economics of Grids, Clouds, Systems and Services, 21-23 Sep 2021, Online. Springer Verlag , pp. 181-184.

https://doi.org/10.1007/978-3-030-92916-9_15

Energy Efficiency in Edge Environments: a Serverless Computing Approach

Karim Djemame[1 [0000−0001−5811−5263]]

School of Computing, University of Leeds, UK K.Djemame@leeds.ac.uk
https://eps.leeds.ac.uk/computing/staff/187/professor-karim-djemame

Abstract. The paper revisits the Internet Architecture by leveraging Software Defined Networks (SDN) with Network Function Virtualisation (NFV) technologies to allow efficient and on-demand placement of Virtual Network Functions (VNF) on a serverless platform for energy-aware function provisioning in edge environments. Edge computing is seen as critical for supporting the next generation of services and applications that demand high speeds and low-latencies though energy consumption is a matter of concern. Serverless computing as a paradigm in virtualisation is considered as a low-latency and a rapidly deployable alternative to traditional virtualisation approaches. Event-triggered serverless functions incentivise energy efficient resource usage and provide granular reporting on a function level. The research will develop a new building block that satisfies the services performance while reducing the energy consumption in edge environments.

Keywords: Serverless computing · energy efficiency · Software Defined Networks · Network Function Virtualisation.

1 Context

There is proliferation of applications benefiting from edge computing solutions: scalability, reliability, cost-effectiveness, which are being adopted in various domains such as autonomous vehicles, traffic management, edge video orchestration, industrial Internet of Things (IIoT) to name a few. Edge computing pushes the intelligence, processing power and communication capabilities of an edge gateway or appliance directly into edge devices, ensuring it is closer to where the data originates from, e.g. sensors.

Virtualisation servers running containers (or unikernels) are usually deployed at multiple locations at the edge of the network. This virtualisation infrastructure hosts not only mobile application services to execute on edge and cloud nodes, but also other related services, namely Network Function Virtualization (NFV) and Software Defined Networking (SDN) to reserve and set up a portion of the underlying networking infrastructure appropriately for guaranteeing the desired runtime behaviours for each application operating on the edge. Such deployment would reduce the deployment costs, and provide a common management and orchestration infrastructure for all virtualised services.

SDNs facilitate the containerised applications and network traffic consolidation to optimise performance and energy consumption. Leveraging SDN together with NFV technologies allows for efficient and on-demand placement and chaining of VNFs, making orchestration and consolidation of services easy and dynamic deployment of network services possible. Moreover, the consideration of Virtualized Network Functions (VNFs) is key to enable 5G application use-cases with specific processing and networking capability requirements. Moreover, a serverless computing system [4] is an ideal solution to build and optimise any IoT operation with zero infrastructure and maintenance costs and little-to-no operating expense [3] as it allows IoT businesses to offload all of a server's typical operational backend responsibilities.

Energy consumption in the Internet architecture is one of the highest operating costs. Energy is becoming even more important due to climate change and sustainability considerations. The advent of 5G mobile-network technology is bringing a significant increase in data traffic and the infrastructure to support it, which consequently will consume more energy. However, applications' performance lies with not only efficient node-level execution but energy consumption as well as these applications may operate in a low energy computing environment. The energy increase coming from applications and infrastructure calls for action. Network load optimisation and efficient resource management are essential to ensure a reduction in total energy consumption.

2 Ambition

This research aims to reduce energy consumption of applications deployment and operation in edge computing by addressing the challenges in resource management to support disruptive applications through large scale connected devices operating in low energy environments. The proposed Internet architecture renovation will be able to automate the deployment, monitoring, scaling of containers running serverless functions ensuring interoperability in an energy-aware edge environment. To do so, it considers the SDN architecture, leveraged with NFV to enable the network to be intelligently and centrally controlled using software applications. Therefore, it addresses the control layer to configure the infrastructure and the application layers to support autonomous energy efficiency in edge computing. The serverless platforms does not take into account energy savings in resource management decisions, and to the best of our knowledge, there is currently no work that addresses performance concerns combined with availability and energy efficiency concerns in serverless computing.

The innovation lies in the incorporation of serverless architectures with 1) SDN controllers which are highly event-driven, modular, and concurrent (with minimal sharing of state between the modules) 2) NFV for orchestrating VNF as well as applications that require running short, on-demand tasks operating on data collected from the data plane. VNFs launched and orchestrated in a serverless manner incentivises efficient resource usage and provides granular reporting on a function level: functions take up the most execution time can be identified, which equals *cost*. This is essentially a *proxy* for energy usage as a unit of (serverless) compute, making VNFs instantiation and orchestration significantly energy and resource efficient.

The SDN controllers are a great fit for the serverless computing paradigm as they are highly event-driven, modular and parallel [1]. Moreover, Serverless computing provides a resource-efficient, low overhead alternative to Virtual Machines (VMs) and containers, and can effectively support the NFV architecture.

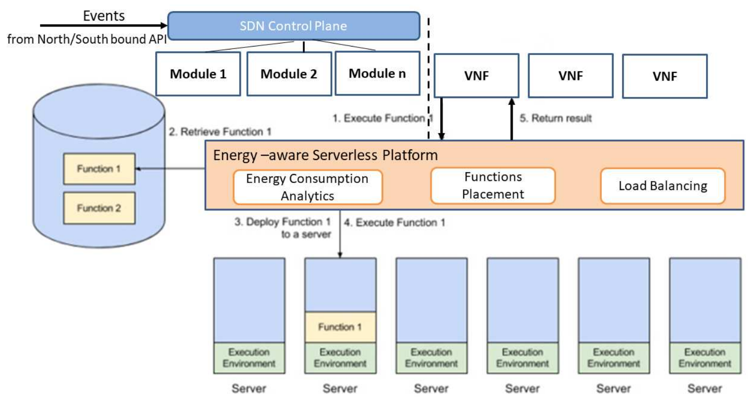

Fig. 1. Proposed Solution

3 Solution: Technical Approach

Architecture. The execution of SDN modules are triggered by events coming from the southbound API, e.g. OpenFlow as well as events received via the northbound API, which interfaces with network applications running on top of the controller, e.g. security services, applications orchestration across edge resources. The SDN controller modules in the SDN architecture are implemented as separate, stand-alone serverless functions (Figure 1), including the flow management on network switches, exercising direct control over the state in the switches via OpenFlow APIs. Serverless functions are used to orchestrate multiple VNFs for short-lived sessions. The open-source serverless platform (e.g. Apache Openwhisk) [2] will be extended to support resource mapping and load balancing to increase resource utilisation by distributing the function executions to available resources with the aim to minimise power consumption. A load balancing strategy considers functions interactions by assigning the function executions belonging to the same session to the same server. Latency sensitive communication services require careful placement of VNFs by allowing locality requirements for grouping functions as a single application. Containers image sizes are reduced to speed up the start of a function execution thus avoiding cold start.

Expected Results. The research new building block is made of 1) a methodology combining SDN, NFV and serverless architectures; 2) placement algorithms for serverless functions to minimise energy consumption; 3) the underlying software implementation. SDN and NFV make communication networks adaptive and scalable. Their combination with the serverless platform provides the required agility, robustness, and scalability for the services executed and will 1) match the demand of a service by scaling up fast to provision additional compute resources for the service (even if that traffic is increasing rapidly); 2) make efficient use of the available resources (services are never over-provisioned and idle service capacity is released immediately); 3) require minimal configuration and management from the developers; 4) isolate services and their provisioned resources from each other e.g. faults and load spikes. A target reduction in energy consumption as a Key Performance Indicator is envisaged in SDN/NFV-enabled networks following the incorporation of the serverless architecture. This comparison is drawn against a use case application as a baseline deployed and operated on traditional network resources (non-serverless platform).

Acknowledgements

The author would like to thank the Next Generation Internet Program for Open INTErnet Renovation (NGI-Pointer 2) for supporting this work under contract 871528 (EDGENESS Project).

References

1. Aditya, P., Akkus, I.E., Beck, A., Chen, R., Hilt, V., Rimac, I., Satzke, K., Stein, M.: Will Serverless Computing Revolutionize NFV? Proceedings of the IEEE **107**(4), 667–678 (2019)
2. Djemame, K., Parker, M., Datsev, D.: Open-source Serverless Architectures: an Evaluation of Apache OpenWhisk. In: 2020 IEEE/ACM 13th International Conference on Utility and Cloud Computing (UCC). pp. 329–335 (2020)
3. Großmann, M., Ioannidis, C., Le, D.T.: Applicability of Serverless Computing in Fog Computing Environments for IoT Scenarios. In: Proceedings of the 12th IEEE/ACM International Conference on Utility and Cloud Computing Companion. p. 29–34. UCC '19 Companion, Association for Computing Machinery, New York, NY, USA (2019)
4. Kritikos, K., Skrzypek, P.: A Review of Serverless Frameworks. In: 2018 IEEE/ACM International Conference on Utility and Cloud Computing Companion (UCC Companion). pp. 161–168 (2018)