

Flexible Distributed Computing Across End-Devices, the Edge and the Cloud

Alexandros PATRAS¹ and Spyros LALIS
*Electrical and Computer Engineering Department
University of Thessaly*

Abstract. The emergence of the edge/fog computing paradigm has increased the programming complexity of applications so that they can work seamlessly in the new distributed and heterogeneous system landscape. In this paper, we investigate a structured dataflow approach which simplifies application development and offers great flexibility regarding the deployment of the application across end-devices, edge computing infrastructure and remote cloud systems. Our prototype is built on top of the Node-RED framework, via a few extensions in order to support the distributed execution of application flows. We use a real-world application example to illustrate our approach, as well as to explore the performance of different deployment scenarios on real computing systems.

Keywords. Internet of Things, edge computing, distributed dataflows, Node-RED.

1. Introduction

A large number of IoT and pervasive computing application scenarios revolve around a rather simple architectural approach, whereby low-end and/or mobile devices send data to and receive actuation/control requests from powerful server machines in the cloud. This approach, though straightforward to implement, has several drawbacks. Firstly, a large amount of low-level data are routed over the Internet to remote machines, leading to scalability issues. Secondly, due to the large latency of the Internet, it may not be possible to support control/feedback loops with tight real-time constraints. Last but not least, privacy-sensitive data ends-up in the cloud from where it may leak to third parties, either intentionally for business purposes or unintentionally as a result of attacks.

An alternative approach is to adopt a more complex system architecture, which allows part of the data processing and decision making to be performed at the edge, on machines close to the end-devices, or in part even directly on the end-devices themselves. However, writing applications that span across end-devices, edge computing infrastructure and remote cloud systems is a non-trivial task: (i) the developer has to structure the application in different parts; (ii) each part must be written/prepped so that it can run on the target host machine(s); (iii) the interface/interaction between the different parts of the application has to be cleanly defined and implemented so that it can be performed over a network; (iv) finally, each part must be installed/instantiated on its host and properly linked together with the other parts of the application. And this process

¹ Corresponding author: Electrical and Computer Engineering Department, University of Thessaly, Gklavani 37, 38221 Volos, Greece; E-mail: patras@uth.gr.