FogPi: A Portable Fog Infrastructure through Raspberry Pis

Cristian Martín, Daniel R. Torres, Manuel Díaz and Bartolomé Rubio
Department of Languages and Computer Science
University of Málaga, Málaga, Spain
{cmf,drtorres,mdr,tolo}@lcc.uma.es

***Abstract*—Nowadays, Fog computing is facing the requirements of time-sensitive applications in the IoT-cloud continuum. These requirements are decisive for mission-critical applications like structural health monitoring. In this paper, a portable Fog computing infrastructure, known as FogPi, is presented. This infrastructure has been designed around Raspberry Pi, which offers a low-cost and scalable solution for running containerized applications. FogPi allows the deployment, management, and orchestration of Docker containers and is especially suitable for environments where the limited Internet connection and reduced budgets limit the adoption of Fog and Edge deployments.**

I. Introduction

The Internet of Things (IoT) [1] is a paradigm that is revolutionizing society thanks to its continuous expansion to many areas, ranging from Industry 4.0 to connected smart cities. As a result, this paradigm is paving the way to sense and actuate the physical world.

These resulting applications are prone to generate large amounts of data such as the structural health monitoring of infrastructures where high frequencies are required to detect precision damages. The limitations of the IoT devices make addressing current paradigms that require large capacities in terms of processing and storage, such as big data and deep learning, very complicated. For this reason, during the last years, cloud computing was integrated with the IoT to provide the capabilities needed to address these paradigms. However, this has resulted in an increase of bandwidth and latency between the IoT and the cloud, which is not admissible by situations where low latency and timely actions are required (e.g., detecting an earthquake).

To overcome the previous challenge, both Edge [2] and Fog computing [3] have been thoroughly studied in recent years. The main idea of both paradigms is to reduce the latency and bandwidth in IoT-cloud communications by moving the computation as close as possible to the IoT. Nonetheless, there are some situations where a limited connection to the Internet (e.g., isolated bridges) makes it difficult to adapt these architectures. Moreover, the deployment of enterprise servers for a Fog/Edge layer can incur high costs and extra maintenance like cold server rooms.

In this paper, we propose FogPi, a portable Fog infrastructure composed of a cluster of Raspberry Pis. Raspberry Pi has gained a lot of attention during the last years as a small single-board computer thanks to its low price and its large community. Other benefits are not only their small size, and high portability, but also their high computing capacity that sometimes it is comparable with personal computers. As a result, this solution, apart from being portable to environments with a limited Internet connection, provides a Fog infrastructure with a low-power consumption compared to enterprise Fog servers and at a low-cost.

The rest of the paper is organized as follows. Section II presents the FogPi architecture and its components for management and monitoring. In Section III, an evaluation of a modal analysis task is performed in FogPI and also in a public cloud platform. Related work is discussed in IV. Then, Section V describes the steps to be carried out in this demo. And finally, conclusions are presented in Section VI.

II. Portable Fog Architecture

There are conditions under which the Internet is limited, and there is a need for monitoring and processing large amounts of in-coming signals from isolated environments. These encourage investing in elevated-cost business fog structures, but they can not compensate economically.

FogPi is, however, an embedded, portable, low-cost, and low-power consumption Fog computing infrastructure, which has the potential to become the ideal solution for this problem in isolated environments that need a Fog infrastructure. This infrastructure runs Docker containers, which bring more benefits such as scalability, high availability, and vertical and horizontal migrations, than traditional infrastructures. Furthermore, FogPi does not require Internet access since containers can be deployed from images previously loaded and configured.

The FogPi software stack, shown in Fig. 1, comprises a cluster of 48 Raspberry Pis model 3 B+ (1GB RAM, 1.4GHz) with 64GB micro-SDs. Each Raspberry Pi board runs the Hypriot[1] Operative System (OS), which is a minimal Debian-based OS optimized and designed to run Docker containers. It includes the Linux Kernel 4.19.75 and Docker 19.03.5 ready for most of Raspberry Pi models. 3D models have been designed to allocate each layer of the FogPi cluster and the power distribution units where the Raspberries are connected and can be controlled.

Docker Swarm[2] enables the orchestration of Docker containers, providing load balancing and high availability through

[1] https://blog.hypriot.com/
[2] https://docs.docker.com/engine/swarm

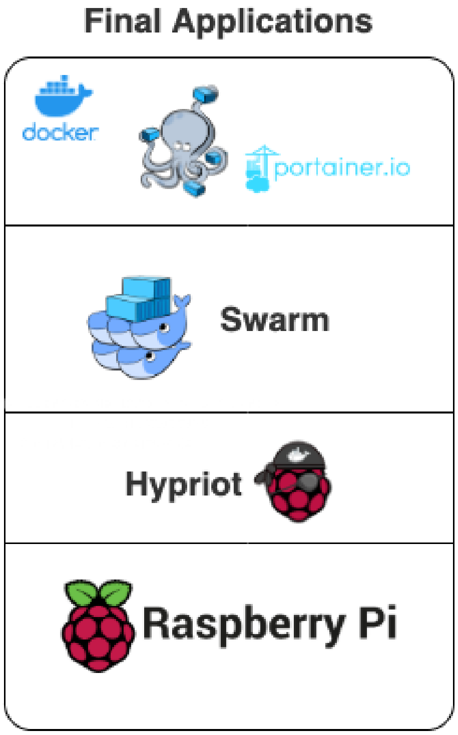

Fig. 1. FogPI software stack for running containerized applications

container replicas and monitoring. Docker Swarm monitors each Raspberry Pi in the infrastructure and allocates and distributes the containers based on their availability and their resource utilization. Therefore, in case a Raspberry fails or has overload, Docker Swarm automatically manages to allocate its components in an available node, enabling to easily distribute high availability mission-critical applications to the FogPi infrastructure. To have a quick view of the state of the nodes and the services that they are running, we use Docker Swarm Visualizer[3], which is an open-source project that provides a user-friendly Web UI for visualizing the nodes belonging to a Swarm cluster. It provides us with a simple interface to control them. However, we do not only need to watch nodes and their status but also need to be able to control services and containers to expand and configure our cluster to different environments. For this purpose, Portainer Community Edition[4] has shown to be a complete and user-friendly solution. Both Docker Swarm Visualizer and Portainer are single lightweight Docker containers. Therefore they can be easily deployed as services in the FogPi infrastructure. All of these open-source projects enable the distribution, management, visualization, and orchestration of Docker containers and nodes, and the allocation of high availability and load balancing micro-services for final applications like structural health monitoring.

Finally, we have used a command-line script to flash[5] the OS image to each micro-SD. This tool saved us a lot of time in burning the images since the configuration of each node such as hostname, IP and SSH keys can be configured in a filename, thereby after the burning process, each micro-SD is ready to go. All the components which comprise the FogPI infrastructure along with the 3D designs and some bash management scripts can be openly accessible in the FogPi GitHub repository[6]. The FogPi cluster is shown in Fig. 2.

[3]https://github.com/dockersamples/docker-swarm-visualizer
[4]https://www.portainer.io/
[5]https://github.com/hypriot/flash

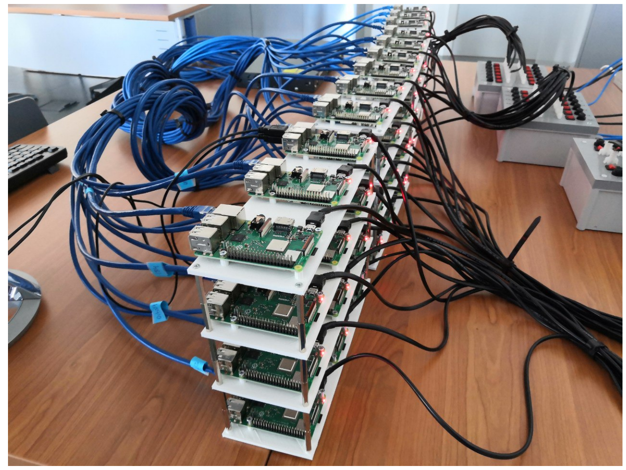

Fig. 2. FogPi: A Portable Fog Infrastructure through a Raspberry Pi cluster

One of the main problems in creating the cluster was the power requirements of Raspberry Pi. Raspberry Pi requires an AWG18 cable, and most of the USB cables do not support that. Therefore, we decided to produce Raspberry Pi power cables by connecting a power AWG18 cable with a micro USB connector. On the other hand, we had two options for powering the cluster: 1) a power supply of at least 120A (48x2,5A per Raspberry) working at 5V required; 2) utilizing USB hubs. In the end, we went for option 1) because the cost of the USB hubs highly increased the cluster costs, and many of them cannot supply the amps required by the cluster.

III. Evaluation

Figure 3 shows the average latency response in FogPi and Google Cloud performing modal analysis of structural health monitoring measures from a different number of edge devices in a tunnel deployment (around 2,9MB). In both the Fog and the Cloud, one container replica with 1GB of RAM has been deployed. These results demonstrate better performance in the portable FogPi, especially when the number of connected edge devices increases.

On the other hand, high availability and load balancing can be provided to applications through Docker container replicas managed by Docker Swarm. Figure 4 shows the average latency of the modal analysis application with a different number of replicas. As can be seen, with a higher number of replicas, FogPi can offer a higher quality of service (QoS). However, when increasing the number of replicas the overhead of the system can also reduce the performance. In this case,

[6]https://github.com/ertis-research/rpi-cluster

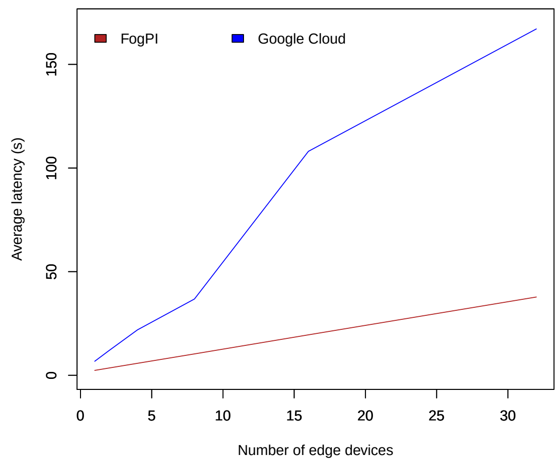

Fig. 3. Average latency response of modal analysis by multiple edge devices in FogPi and Google Cloud

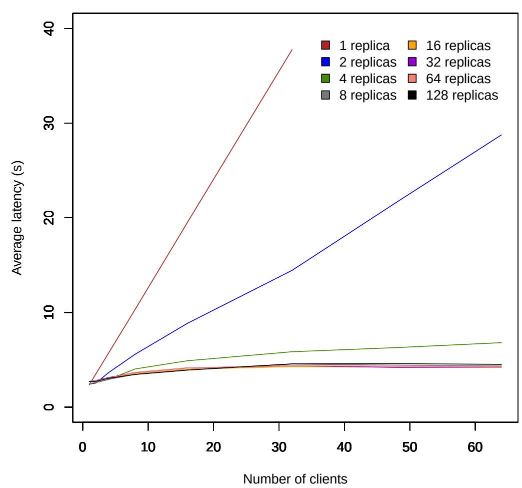

Fig. 4. Average latency response of modal analysis in FogPi with different numbers of replicas

the lowest latency is obtained with 16 replicas, and from 32 to 128 replicas the latency is stable or higher.

IV. Related work

This project is definitely inspired by other Raspberry PI clusters, like Iridis-Pi [4], which was one of the first Raspberry Pi computing clusters. Their low-price and computing power have made possible their application to smart grids, teaching, data mining algorithms and structural health monitoring like this case. Other Raspberry PI clusters have also been proposed for the Fog and Edge, like in [5], where a container-based edge Platform-as-a-Service is proposed. However, due to its age, it was not able to use current orchestration container platforms.

V. Demonstration

In the first part of the demonstration, users will see how the infrastructure is designed and can be configured. Moreover, the tools we use to monitor and manage the FogPi infrastructure will be presented. Secondly, how a container application can be deployed and easily scaled in the infrastructure for load balancing and high availability will be shown. Finally, users will see how the latency can be reduced concerning a cloud solution in the FogPi infrastructure with load balancing.

VI. Conclusions

Fog computing is considered a disruptive paradigm to provide the IoT and cloud with the optimization of latency and bandwidth required in mission-critical applications. This paradigm has been materialized through a portable, low-cost and low-power consumption solution with Raspberry Pi devices, which enable a container infrastructure with load balancing and high availability for isolated environments. This solution, known as FogPi, and the 3D modular cases and components are openly available on the Internet. The evaluation performed for modal analysis in a structural health monitoring application denotes a better performance of FogPi compared to a public cloud solution.

Cluster federation allows the coordination and management of multiple clusters from a single hosting cluster. This feature is required in the IoT-cloud continuum, where multiple layers like edge, fog and cloud can have its own container cluster for cluster isolation and avoiding a high latency in communications. Nevertheless, this is not provided in the community version of Docker Swarm. Therefore, other container orchestration mechanisms like Kubernetes will be explored and evaluated for FogPi as future work.

Acknowledgment

This work is funded by the Spanish project RT2018-099777-B-100 ("rFOG: Improving latency and reliability of offloaded computation to the FOG for critical services"). Thanks to Luis Alonso and Fran Martín that helped us in the design of the 3D models and finding electronic components respectively.

References

[1] M. Díaz, C. Martín, and B. Rubio, "State-of-the-art, challenges, and open issues in the integration of internet of things and cloud computing," *Journal of Network and Computer Applications*, vol. 67, pp. 99–117, 2016.

[2] C. Martín Fernández, M. Díaz Rodríguez, and B. Rubio Muñoz, "An edge computing architecture in the internet of things," in *2018 IEEE 21st International Symposium on Real-Time Distributed Computing (ISORC), 29-31 May, Singapore*, 2018, pp. 99–102.

[3] C. Martín, D. Garrido, M. Díaz, and B. Rubio, "From the edge to the cloud: Enabling reliable iot applications," in *2019 7th International Conference on Future Internet of Things and Cloud (FiCloud), August 26-28, Istanbul, Turkey*, 2019, pp. 17–22.

[4] S. J. Cox, J. T. Cox, R. P. Boardman, S. J. Johnston, M. Scott, and N. S. O'brien, "Iridis-pi: a low-cost, compact demonstration cluster," *Cluster Computing*, vol. 17, no. 2, pp. 349–358, 2014.

[5] C. Pahl, S. Helmer, L. Miori, J. Sanin, and B. Lee, "A container-based edge cloud paas architecture based on raspberry pi clusters," in *2016 IEEE 4th International Conference on Future Internet of Things and Cloud Workshops (FiCloudW)*. IEEE, 2016, pp. 117–124.