Experimenting with Cloud and Network Orchestration for Multi-Access Edge Computing

P. Valsamas, S. Skaperas, G. Violettas, T. Theodorou, S. Petridou, D. Vardalis, A. Tsioukas, L. Mamatas
Department of Applied Informatics, University of Macedonia, Thessaloniki, Greece
{xvalsama, sotskap, georgevio, theodorou, spetrido, emamatas}@uom.edu.gr, dvardali@ee.duth.gr, tsioukas@gmail.com

***Abstract*—We propose the Multi-Access Edge Computing for Content and Measurements (MECOM) experimentation platform which orchestrates cloud and network aspects for efficient content and IoT measurements' distribution to mobile users. MECOM brings together: (i) heterogeneous lightweight cloud technologies hosting the content and IoT data; (ii) multi-homing capabilities in the mobile nodes that select the best connectivity option for the service used, e.g., for low latency or high throughput; (iii) early content popularity change detection that scales the content distribution; and (iv) IoT routing protocol adjustments that reduce the delays of measurements' collection. Our proposal addresses diverse service requirements in two relevant scenarios, i.e., for content distribution and IoT services.**

I. Introduction

The emerging 5G Networks call for new network and cloud paradigms that bring elasticity in the network environment and enable high throughput or ultra-low latency services, such as on content or IoT data distribution. Along these lines, multiple research initiatives bring together networks and clouds, including the Multi-Access Edge and Fog Computing.

MECOM is a novel experimentation facility for integrated cloud and network orchestration targeting the specific scenarios of elastic content distribution and IoT measurements' collection. In this context, we integrate and optimize together a real mobile environment with lightweight cloud resources (e.g., Unikernels). In contrast to related experimentation platforms for elastic content distribution [1], [2] and Fog Computing [3], MECOM exhibits the following novel features:

- realizes intelligent and modular orchestration for both cloud (e.g., efficient virtual machine placement) and network aspects (e.g., dynamic load balancing and multi-homing connectivity control);
- supports heterogeneous lightweight virtualization in the network edge, e.g., Unikernels which are suitable for mobile environments due to their small footprint and rapid manipulation [1];
- implements novel content-popularity change point (CP) detection scaling the cloud resources;
- exploits mobile clients that reside at real moving buses, based on the novel MONROE platform [4]; and
- performs multi-homing capabilities and protocol adjustments for the mobile nodes and IoT devices, respectively.

MECOM performs real experimentation in two corresponding scenarios, i.e., for content-distribution and IoT services. In the following sections II and III, we present the MECOM platform and describe the demo, respectively.

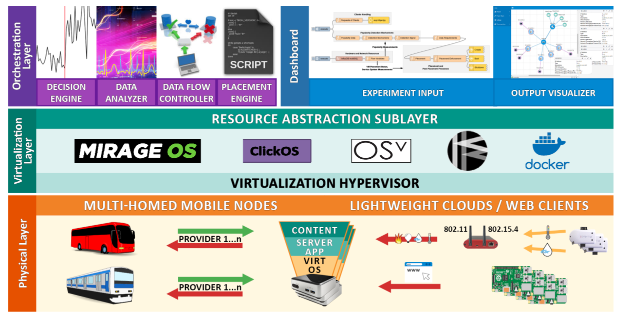

Figure 1. The MECOM experimentation platform

II. The MECOM Platform

Fig. 1 illustrates a high-level overview of the MECOM architecture. In a bottom-up approach, we highlight the following three layers: (i) the *Physical Layer* consisting of the lightweight edge clouds, IoT devices that collect measurements, and mobile clients with multi-homing capabilities; (ii) the *Virtualization Layer* enabling lightweight cloud capabilities through Virtual Machines (VMs) with "tiny" operating systems, such as the Mirage OS and Rump Kernel Unikernel technologies. A Resource Abstraction Sublayer (RAS) hides virtualization heterogeneity and exports a uniform interface for VM control; and (iii) the *Orchestration Layer* with the following features: a Data analyzer which performs CP detection to early "track" changes in the content evolution, a Decision engine which specifies either to deploy or remove lightweight VMs accommodating content and IoT data near the end-users, a Data flow controller that balances the traffic load among active VMs, and a Placement engine which determines the optimal location of the VMs. The MECOM dashboard takes the experimentation input through the Node-RED tool and provides the corresponding results (i.e., see Fig. 2).

The MECOM implementation details are summarized as follows. Unikernel-based Web servers offer the data to the mobile clients and are dynamically manipulated from the orchestration processes through the RAS. The latter exports a unified north interface to the former, while its south interface communicates via ansible scripts with different Unikernel flavors. For extensibility, the orchestration entities are implemented as Node-RED nodes. The MONROE testbed [4] provides real mobile nodes with multi-homing capabilities. A *Multi-homing mechanism* is deployed on real vehicles to

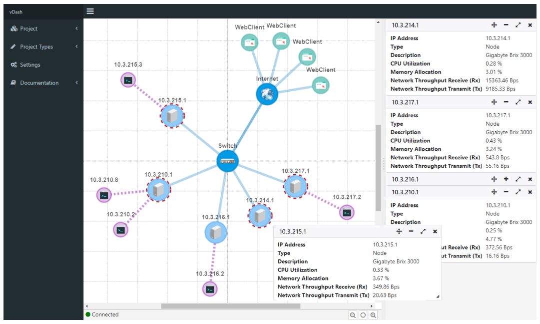

Figure 2. Visualization of a live experiment

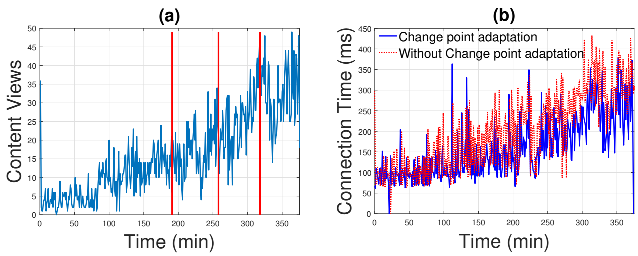

Figure 3. a) On-line CP detection (vertical line), b) equivalent web-client performance

improve the mobile broadband (MBB) connectivity of mobile users, in terms of throughput and delay, i.e., for the content distribution and IoT services scenario, respectively. Zolertia RE-Mote devices produce the periodic measurements, besides the complicated IoT setups implemented in the Cooja emulator. Finally, our *DNS-based load-balance* tool keeps track of the content popularity and realizes load-balancing through an SDN controller. An early version of a relevant content distribution scenario is presented in [5].

III. Demo Description

Our live demo highlights the aforementioned MECOM platform's novel aspects through implementing the content distribution and IoT services scenarios. We show our bespoke visualization tool in Fig. 2, while the demo video can be found at https://bit.ly/2KSO6HF.

The first scenario is on a novel elastic content distribution network. A number of low-end PCs host video content that may become viral and cause resource exhaustion at the edge cloud. Outside, mobile clients on MONROE buses suffer from MBB connection inefficiencies and delays due to the overloaded cloud. MECOM handles both issues with its orchestration mechanisms. The *Multi-homing mechanism* constantly monitors the MBB connectivity in terms of delay and bandwidth (e.g., by periodic "ping" and HTTP downloading) and allows the mobile users to dynamically and transparently switch among the available providers, i.e., selecting the best connectivity option for the service used. The *Data Analyzer* employs our novel *CP detection* which early detects significant content-popularity changes. Notifications, such as the ones illustrated in Fig. 3a, are provided to the *Decision Engine*

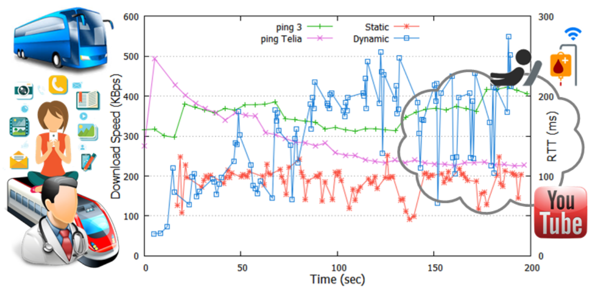

Figure 4. Multi-homing for mobile users

which initiates the deployment of additional content-caches, in the form of Unikernels. In response, the *Placement Engine* defines the location of VMs while the DNS-based load balancer assigns the users' requests to different Unikernels. The improvements in users' connection time as a result of content distribution orchestration are depicted in Fig. 3b.

In the IoT scenario, a medical doctor accesses up-to-date IoT vital sings of his patients using his smartphone; the service requires ultra-low latency. To achieve this: (i) the mobile users are accessing IoT measurements cached in Unikernels; (ii) the *Multi-homing mechanism* improves mobile communication, e.g., achieves the highest downloading speed by dynamically switching between two Swedish providers, namely "3" and "Telia", as shown in Fig. 4; and (iii) IoT routing protocol adjustments reduce the delays in the IoT network.

Our results show that orchestrating network resources with lightweight clouds brings advantages in mobile users' experience. Next steps include implementing relevant network slicing techniques and evolving the MECOM experimentation facilities to a complete MANO-inspired architecture, as the one defined in the NECOS project [6].

Acknowledgements

This work is partially supported by the H2020 NECOS (grant agreement no. 777067) and the MONROE (grant agreement no. 644399) projects.

References

[1] S. Kuenzer, A. Ivanov *et al.*, "Unikernels everywhere: The case for elastic cdns," in *Proc. of the 13th ACM SIGPLAN/SIGOPS Int. Conf. on Virtual Execution Environments*. ACM, 2017, pp. 15–29.

[2] G. Siracusano, R. Bifulco, M. Trevisan *et al.*, "Re-designing dynamic content delivery in the light of a virtualized infrastructure," *IEEE J. on Selected Areas in Communications*, vol. 35, no. 11, pp. 2574–2585, 2017.

[3] H. Gupta, A. Vahid Dastjerdi, S. K. Ghosh, and R. Buyya, "ifogsim: A toolkit for modeling and simulation of resource management techniques in the internet of things, edge and fog computing environments," *Software: Practice and Experience*, vol. 47, no. 9, pp. 1275–1296, 2017.

[4] Ö. Alay, A. Lutu, M. Peón-Quirós *et al.*, "Experience: An open platform for experimentation with commercial mobile broadband networks," in *Proc. of the 23rd Ann. Int. Conf. on Mobile Computing and Networking*. ACM, 2017, pp. 70–78.

[5] P. Valsamas, S. Skaperas, and L. Mamatas, "Elastic content distribution based on unikernels and change-point analysis," in *Proc. of the 24th European Wireless Conf.*, 2018.

[6] F. Silva, M. Lemos, A. Medeiros *et al.*, "Necos project: Towards lightweight slicing of cloud-federated infrastructures," in *Workshop on advances in slicing for softwarized infrastructures (S4SI) IEEE NETSOFT*, 2018.