

Multi-Access Edge Computing for Efficient Content Distribution and IoT Services

Polychronis Valsamas, Sotiris Skaperas, George Violettas, Tryfon Theodorou, Sophia Petridou, Dimitris Vardalis, Antonios Tsioukas, Lefteris Mamatras
Department of Applied Informatics, University of Macedonia, Greece

ABSTRACT

We demonstrate the Multi-Access Edge Computing for Content and Measurements (MECOM) platform which orchestrates content distribution and IoT measurements' collection for improved mobile user experience. MECOM targets systemically the latter by bringing together: (i) heterogeneous lightweight clouds hosting content and IoT data near the end-users; (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.

1 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.

Our goal is to integrate and optimize together mobile environments with lightweight cloud resources (i.e., Unikernels or containers), and investigate the novel associated problems. For example, MECOM attempts to improve the communication of mobile users through orchestrating different network edge aspects, including content or measurements provisioning, connectivity options, and involved network protocols.

In contrast to the relevant platforms [2] and [4], MECOM: i) realizes intelligent and modular orchestration; ii) integrates heterogeneous lightweight visualization and IoT technologies; iii) supports multi-homing capabilities and mobility aspects; and iv) performs real experimentation.

More precisely, the proposed platform and demo integrates: i) modular orchestration features, i.e., efficient virtual machine (VM) placement, dynamic load balancing and multi-homing connectivity control; ii) Unikernel-based lightweight cloud technologies for the network edge, i.e., appropriateness for mobile environments due to their small footprint and rapid manipulation [2]; iii) novel content-popularity

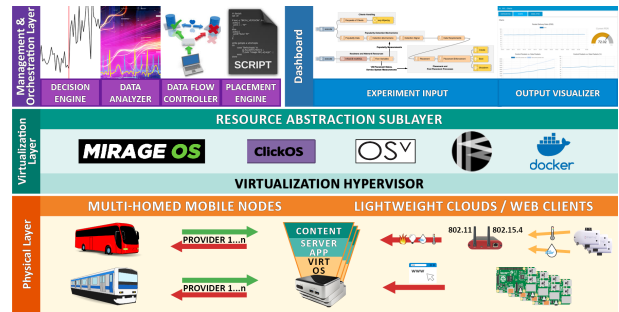


Figure 1: The architecture of the MECOM platform

change point (CP) detection scaling the cloud resources; iv) mobile clients based on the novel MONROE platform [1], i.e., residing at real moving buses; and v) content-distribution and IoT functionalities through two corresponding scenarios. In the following sections 2 and 3, we present the MECOM platform and describe the demo, respectively.

2 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 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 *Management and 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 passes the experimentation input through the Node-RED tool and provides the corresponding results.

The implementation details of the MECOM platform can be briefly summarized as follows: The *VM orchestration* tool places/monitors/removes Unikernel-based Web servers, i.e.,

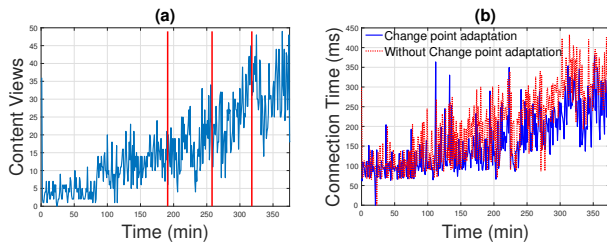


Figure 2: a) On-line CP detection (vertical line), b) equivalent web-client performance

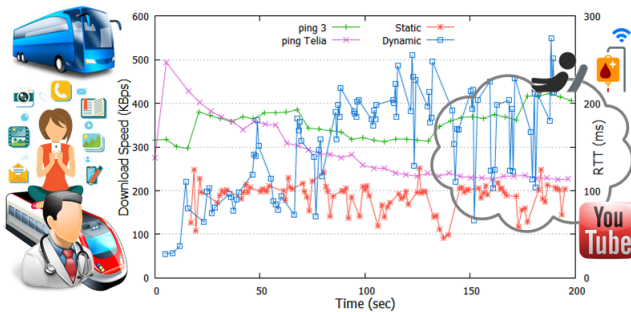


Figure 3: Multi-homing for mobile users

providing content and IoT data. Orchestration processes were simulated as Node-RED nodes whose unified north interface communicates with the hypervisor through the RAS, while their south interface communicates via ansible scripts with different Unikernel flavors. The *Multi-homing mechanism* is deployed on real vehicles to 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. Finally, the *DNS-based load-balance* tool keeps track and redirects the end-users' Web requests to particular content caches (i.e., hosted by Unikernels), in a round-robin fashion. Further implementation details are described in [5] and the demo video (<https://bit.ly/2wZj37f>).

3 DEMO DESCRIPTION

Our demo highlights the aforementioned MECOM platform's novel aspects in two relevant scenarios, namely for content distribution and IoT services. A major contribution is the real experimentation setup for both scenarios offered by the integration of the MONROE testbed [1] (i.e., providing real mobile nodes with multi-homing capabilities) with our SWN testbed (i.e., implementing the heterogeneous lightweight clouds and the real IoT network extension). The test-bed facility collaborates with the Cooja emulator for the implementation of more complicated IoT setups.

The first scenario is on a novel elastic content distribution network. A number of mini 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 *VM orchestrator* employs our novel *CP detection* which early detects significant content-popularity changes. Notifications such as the ones illustrated in Fig. 2a are provided to the *Decision Engine*, 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. As a result of this resource orchestration, the improvements in users' connection time are depicted in Fig. 2b.

In the second IoT services scenario, the mobile users are accessing IoT measurements cached in Unikernels. We assume that a medical doctor has to access up-to-date IoT vital signs of his patients using his smartphone, hence, the service requires ultra-low latency. The MECOM's *Multi-homing mechanism* achieves the highest downloading speed by dynamically switching between two Swedish mobile service providers, namely "3" and "Telia", as shown in Fig. 3.

Our results show that orchestrating network resources with lightweight clouds brings advantages in mobile users' experience. Next steps include implementing relevant network slicing techniques using the NECOS architecture [3].

ACKNOWLEDGEMENTS

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

REFERENCES

- [1] Ö. Alay, A. Lutu, M. Peón-Quirós, and others. 2017. 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, 70–78.
- [2] S. Kuenzer, A. Ivanov, and others. 2017. Unikernels Everywhere: The Case for Elastic CDNs. In *Proc. of the 13th ACM SIGPLAN/SIGOPS Int. Conf. on Virtual Execution Environments*. ACM, 15–29.
- [3] F.D. Silva, M. Lemos, A. Medeiros, and others. 2018. NECOS Project: Towards Lightweight Slicing of Cloud-Federated Infrastructures. In *Workshop on advances in slicing for softwarized infrastructures (S4SI) IEEE NETSOFT*.
- [4] G. Siracusano, R. Bifulco, M. Trevisan, and others. 2017. Re-designing Dynamic Content Delivery in the Light of a Virtualized Infrastructure. *IEEE J. on Selected Areas in Communications* 35, 11 (2017), 2574–2585.
- [5] P. Valsamas, S. Skaperas, and L. Mamatas. 2018. Elastic Content Distribution Based on Unikernels and Change-Point Analysis. In *Proc. of the 24th European Wireless Conf.*

4 DEMO REQUIREMENTS

Our demo requirements follow:

- **Space needed:** The table space provided by the organizers is enough to cover our needs.
- **Setup time:** We need approximately 30 *min* to setup up our demo.
- **Additional facilities:** Our demo requires 220 V power connections for four (4) devices.
- **Poster Size:** A0 paper size in portrait mode (841 × 1189 *mm*).