A 5G Platform for Future Interactive Media Systems

D. Trossen and S. Robitzsch
InterDigital Europe
London, England
Email: dirk.trossen@interdigital.com

M. Boniface and S. Phillips
University of Southampton
Southampton, England
Email: m.j.boniface@soton.ac.uk

A. Betzler, C. Fernandez, D. Guija
M. Catalan and P. S. Khodashenas
Fundacio i2CAT, Barcelona, Spain
Email: august.betzler@i2cat.net

***Abstract*—5G promises an increasingly networked society with features that go significantly beyond current 4G offerings. The characteristics of low latency and high throughput, greater adaptability of highly distributed compute, storage and communications infrastructure, and configuration and management of a wide range of network services will transform how future applications are conceptualised, provisioned and utilised. This paradigm shift in network operations and usage will create a tighter integration of infrastructure with services, content and demand in the creative industries and transport, energy and public sectors. 5G will deliver innovative services that improve the Quality of Experience (QoE) for citizens whilst controlling cost for service providers as infrastructure demands increase. Future Interactive Systems are ideal 5G services to validate the acceptance, viability and performance of 5G solutions; as such systems have considerable capacity and performance requirements, the content that is produced and distributed (e.g. Augmented Reality and Virtual Reality) exhibits challenging variability characteristics that impact resource consumption and drive service and infrastructure adaptation. Personalisation, interaction, mobility and localisation all require responses from services and the infrastructure if QoE is to be delivered and costs controlled. In this paper, we present a 5G service delivery platform FLAME that address the needs of future interactive systems. We outline the requirements for future interactive systems, describe the conceptual design of the FLAME platform and provide insights from deployment in the city of Barcelona.**

I. Introduction

The networked society is increasingly dependent on interactive media systems. Today many systems are based on Over-The-Top (OTT) content distribution approaches where the delivery of audio, video, and other media over the Internet is achieved without the involvement of a network operator in the control or distribution of the content. Such approaches tend to result in non-optimal resource allocations driving over-provisioning costs for network operators or poor Quality of Experience (QoE) for consumers. Dedicated content delivery networks (CDNs) offered by global companies remain the predominate way to cope with the constantly growing demand for multimedia delivery through the in-network placement of dedicated storage and delivery resources.

However, over the last decade visualized infrastructures have increasing supported the requirements of interactive systems through techniques for slice-based isolation, programmability, and monitoring. Real-time requirements for media service have been specified using visualized service networks and translated dynamically to resource specifications in operator clouds, including Quality of Service (QoS) characteristics bandwidth, latency, packet loss and jitter [1]. Whilst the emergence of SDN technologies such as Open Flow [2] has fostered transformational networking and distributed computing models such as fog and mobile edge computing [3] [4] . SDN has enabled application specific controllers to dynamically control topology and QoS [5] increase observability of network traffic [6], and supported end-to-end communication channels over multiple mediums to connected distributed data centres [7].

In addition to the increasing softwarization of infrastructure, content formats and production and consumption patterns are continuously changing as users demand improved Quality of Experience (QoE). This trend [9] requires adaptation processes that respond to aspects such as personalisation, localisation, interactivity and mobility. Personalisation adapts content or content distribution processes to meet individual preferences, devices, activities and situations (e.g. presenting music videos according to user profiles). Localisation adapts to requirements of groups of users, often related to a specific geographic location (e.g.tailoring content for a specific market segment).Interactivity is the process by which a system in response to user actions (e.g. presenting a specific video segment in response to a click navigation or presenting views of a virtual world in response to gaze). Mobility is the process of dynamically distributing content in response to proximity of highly mobile users (e.g. geographical migration of data according to crowd movements).

In this paper, we describe a 5G platform "FLAME" that optimises media content delivery through direct interaction with underlying SDN management functions. FLAME promotes a content delivery solution that addresses emerging demand trends through cross-layer integration between Virtualised Service Networks (VSNs) and Information-Centric Networking (ICN) [8]. We firstly describe the two-sided platform architecture, then experiences of southbound integration with an urban-scale infrastructure, and finally northbound usage by exemplar future interactive systems.

II. FLAME Platform Architecture

The ambition of the FLAME Platform is to offer new dynamic content production and delivery platform supporting optimal delivery of future interactive systems using distributed computing models that combine distant media cloud with mobile edge (i.e., support for many points of presence) and deployed in software-defined infrastructure within minutes, supporting flexible service endpoint management, routing and switching. [8] provides a short discussion on Networked Multimedia Systems outlining the changes to user experience and the relationship with advanced infrastructures responsible for storing, processing and transmitting multimedia content. In this case, the user shifts from the role of observer to participant, where their input (e.g., on-demand navigation and streaming) determines the output.

An important scoping factor is that we are exploring the relationship between media and the infrastructure resources used to store, process and transmit the content. Traditionally, interactive media systems are developed without the direct involvement of the network operator in terms of utilising in-network resources, usually referred to as Over-The-Top (OTT) service provision. FLAME instead explores how to bring media service providers and operators together in a way that allows the infrastructure to be dynamically optimised based on the demand for interactive multimedia content with the constraint for optimising the quality of interactive media content against the available infrastructure resources. A definition of Interactive Media is given in [11] as "Interactive media is the integration of digital media including combinations of electronic text, graphics, moving images, and sound, into a structured digital computerised environment that allows people to interact with the content for appropriate purposes. The digital environment can include the Internet, telecoms and interactive digital television.". Based on that, we must model and analyse factors influencing the relationship between content and the communication system used to deliver it. Our initial analysis highlighted personalisation, interactivity, mobility and localisation as stated above are the main factors that can influence adaptation of a solution over the infrastructure [10].

A. FLAME Platform Components

Given this scope, we base our resulting platform on an emerging view of the underlying infrastructure in alignment with future 5G systems. Hence, we assume support for isolation of resources through virtualisation techniques at the compute, storage and communication level. More specifically, we assume an adoption of the emerging platform technologies being used within FLAME, specifically those surrounding NFV [12] as well as SDN [13]. Virtualisation technologies will expose virtual resources to software instances of functions and services of the platform as well as at the application level, so called Virtual Network Functions (VNF). In this context, HW resources are provided within the so-called slice of the infrastructure. This brings us to an important aspect of the emerging resource management view of future infrastructures, namely to provide the management of resources and services provided. Slicing techniques are being used to establish the perception of perfect resource isolation to other deployments happening within the infrastructure.

Once operating within said resource isolation of an infrastructure slice, FLAME will establish its own resource management scheme within the constraints of the underlying (finite) HW resources. Such constraint of finite resources is particularly relevant when considering infrastructure in physically constrained environments, such as those found in street furniture deployments or generally what we consider "edge deployments of resources". Such resources very often comprise of a limited number of processors, memory and network capability, many of which are already taken by platform functions that need to take priority in resource allocations over non-essential functions (such as media services). Platform-wide resource management must consider such constraints and expose resources to media service providers in an appropriate manner that ensures execution of the platform functions underneath. Suitable orchestration templates ensure such management to happen.

An important consideration for the resource management at the infrastructure level is that of the point-of-presence (POP) of resource pools. Many infrastructure deployments nowadays assume a single POP, e.g., in a local data centre, from which resources will be provided to the service provider. In FLAME, we assume a multi-POP approach, i.e. many pools of resources that are distributed over the municipal infrastructure. Even more so, we assume a multi-tier POP deployment that ranges from resources based in street furniture, allowing for the service provider to claim resources in a specific geo-spatial location (e.g., near a building or a landmark) and possibly near end users, complemented by resources in regional offices, such as serving a geographically larger area like a city district, ultimately complemented yet again through the distant cloud, representing the today's well-known scaled data centres of cloud providers such as Amazon and alike. Through such multi-tiered approach, we assume geo-spatial constraints to govern the reservation of resources within FLAME and its underlying infrastructure.

In addition to the representation of resources at a single or multiple presences, the relation to the infrastructure(s) is crucial. In FLAME, we assume that one FLAME platform instance is provided by a single infrastructure provider. However, infrastructure brokers or exchanges can be used to expose a single view over multiple physical infrastructures, e.g., in different municipalities. In that case, the (distributed) infrastructure would appear as a single one towards FLAME albeit provided by two different infrastructure providers.

Figure 1 shows the FLAME platform architecture [14], operating on top of an infrastructure. In this figure, we focus on the main layers of the overall architecture, including the FLAME platform, while showing crucial inter-layer interfaces for explanation of the overall workings.

At the very bottom, we assume the existence of the infrastructure (provider), exposing an OpenStack compliant interface to the FLAME platform for resource management at the wholesale level, i.e. the FLAME platform reserves platform resources in the compute, storage and networking domain. Resources of the infrastructure provided to the FLAME platform are in turn provided as retail resources to the media services at the top of the platform through management interfaces exposed to media services. The management interfaces allow to initiate the orchestration of (retail) resources from the media service provider side. Furthermore, FLAME relies on the OpenFlow

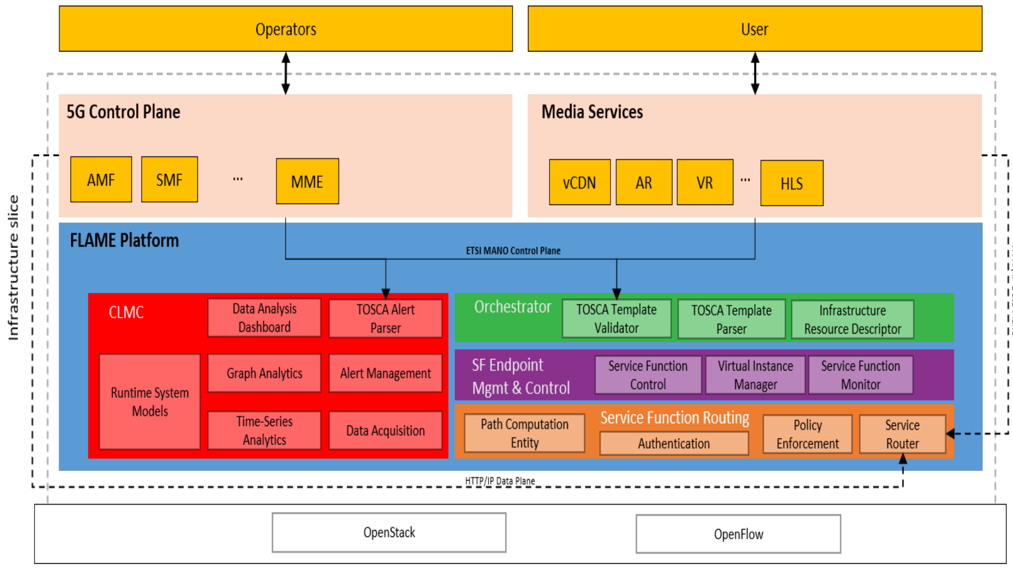

Fig. 1: FLAME Platform Architecture

protocol (via the OpenDaylight controller [15]) to configure the suitable communication resources via forwarding rules in the underlying infrastructure network.

We consider the realisation of the media services outside the scope of the FLAME platform itself. We do assume, however, that a media service is realised through a set of media components, each communicating with each other through an HTTP/IP-compliant data plane interface, while utilising the management and monitoring interfaces to the FLAME platform to facilitate and enable the deployment of those media components through the FLAME platform, i.e. media components are (computing and storage) resources from a FLAME platform perspective. Towards the end user, we see media components utilising service-specific interfaces and interaction methods. Media components are implemented utilising FLAME platform provided resources, such as servers, connectivity components (e.g., switches) and others, while also utilising resources outside of the scope of FLAME, such as end user devices and Internet-of-things components. With that in mind, we position FLAME as a distributed programmable resource platform, which can be used by media services for the fulfilment of a desired experience towards users.

The Orchestration component of the FLAME platform interfaces with the infrastructure resource management to administer the compute/storage/network resources at the retail level, i.e. towards the media service provider, while it utilises a platform-internal interface towards the Service Function (SF) Endpoint Management and Control (SFEMC) component to realise the orchestration-level management [16] policies as well as to set suitable shorter-term control policies for service function endpoints. For the realisation of the configured service function endpoint policies, the SFEMC will utilise an FQDN registration interface towards the service function routing component to control the registration and deregistration of the service endpoints. With this, the visibility of specific service function endpoints, addressed via specific a FQDN, can be controlled, allowing for service requests to be routed to said service function endpoint. The Service Routing component, in turn, will use the OpenFlow interface (e.g., via suitable platforms such as OpenDaylight) to suitably configure the switching fabric of the underlying infrastructure, while utilising the data plane Layer 2 Ethernet interface to forward the media service data traffic, together with any platform internal traffic.

Particular consideration is given in our platform to the gathering of information across various layers and enabling demand/supply analysis for controlling the flow of content in the network. This functionality is realised by the Cross-Layer Management and Control (CLMC) component. The gathered data and the performed analytics over said data is useful and needed for control-level decisions, such as the activation of service endpoints by realising the specific policy decisions expressed in the orchestration template provided by media service providers, it also provides a rich pool of data for the experiments, either as insights towards the creation of further experimentation or insights that reflect directly on the ongoing experiments, e.g., adjusting crucial longer-term strategies such as those for content placement or media adaptation. FLAME taps into existing monitoring frameworks provided by the individual sub-systems (such as those provided directly by switching platforms such as OpenDaylight or ONOS [17], obtained through the aforementioned OpenFlow [18] interface towards the underlying infrastructure), while also developing solutions for analytics and knowledge management beyond those existing ones.

B. *FLAME Platform Benefits*

With the design of the FLAME platform, we are able to improve on performance of interactive media systems whilst managing costs associated with infrastructure resources. In particular, we consider the following platform benefits:

- **Reduce latency**: Latency has long been recognised as a major impact on user experience, leading not only to the deployment of content delivery networks but many past and ongoing protocol improvements (e.g., introduction of QUIC [19] aiming at browsing latency improvements). Reducing the service path length is an important target for FLAME through utilising an intelligent service endpoint management and flexible routing solutions, allowing for services provided at minimum one hop length when utilising cabinet-level deployments of infrastructure.
- **Stem unicast proliferation**: The emergence of HTTP as the de-facto streaming protocol in the Internet results in infrastructure providers being currently incapable of utilising in-network multicast capabilities in order to stem the linear cost explosion that the unicast delivery model of HTTP creates. Through its innovative service function routing component, the FLAME platform realises the capability to deliver HTTP response through efficient Layer2 multicast, therefore providing a unique capability to significantly reduce costs for multi-viewer scenarios.
- **Differentiate services**: Virtualisation opens up the capability to differentiate services by placing service endpoints throughout the network with localized and personalized behaviour. This, however, requires the network to provide a dynamic service routing capability that directs traffic to the most appropriate local service instance. Also, a failover mechanism is required to indirect service requests if a local instance is unable to provide sufficient service response. Furthermore, adaptive media delivery is crucial for differentiation of services, allowing for adapting services, for instance, to different user device requirements by adding transcoding capabilities to the service path for specific users. FLAME provides exactly these capabilities through its innovative service function routing component.
- **Localise traffic**: Reduction of network traffic is often realised through localising traffic wherever possible, also addressing the aforementioned latency reduction. Capitalising on FLAME capabilities to dynamically route requests to the most appropriate service instance achieves a likely significant reduction of traffic being sent over longer paths. It also allows for keeping data local in terms of information security, therefore possibly exposing traffic to fewer parties involved. This ability to localise traffic needs to be balanced in a real-life deployment with the possibly higher operational costs for the distributed servers in comparison to centralized data centres. FLAME provides this ability to trade off these aspects towards a commercially viable offering.
- **Cross layer analytics**: Matching the supply of resources against time-varying demand for content is a difficult task under the pressing cost constraints that every media service provider experiences. FLAME provides unique insights into data points that help expressing demands and, more importantly, the fulfilment of critical performance indicators for such demand through the CLMC component that provides insights into data points across all layers of the system, including low level infrastructure but also high-level service data points.
- **Remove insecure content access**: FLAME provides the capability to elevate content delivery from intermediary Content Delivery Networks (CDNs) to fully secured surrogate service endpoints. This capability provides stronger security through the ability to not only host the content but also the security context in which this content is delivered. Through this, insecure content references could be removed by allowing content to be hosted at surrogate service endpoints with minimal computational authorisation functionality. This ensures that content is not exposed to unauthorised parties. In addition, any delegated secured content is delivered along a direct path, removing the need for triangular routing to origin servers as required in standard IP routing systems.

III. Urban Deployment Details and Insights

After presenting the FLAME solution and its benefits, the following section focuses on the considerations for an infrastructure deployment and then a real-life city deployment in Barcelona, which offers FLAME capabilities at urban scale.

A. Considerations for Urban Infrastructure Deployments

To enjoy the benefits of FLAME, the city deployments should involve a variety of infrastructure mechanisms depending on physical constraints of the city setting. Installations should be planned based on coverage and performance required; which will dictate where equipment should be placed. In any case, at the equipment level the infrastructure should provide [20]:

- **Wireless and Radio equipment** which can be deployed on lampposts. In these locations, equipment shall be ruggedised for outdoors and there has to be an easy way to access these devices, yet remote access or reset is recommended. There are power, aesthetic and weight considerations to mounting equipment on a lamppost which are usually defined by the lamppost manufacturers specifications. Radio equipment can be either mounted in a street cabinet at the base of the lamppost for accessibility and security reasons, or they can be mounted on top of the lampposts within safely attached boxes.
  Wireless equipment can also be installed directly on buildings. These installations are easy to access and maintain but must be installed in a secure location.
- **Active nodes / switching capabilities / firewalls / monitoring equipment** should be installed in a secure room within a building. Staff should be able to access equipment for regular planned maintenance but also be able to gain rapid access in case of an emergency.
- **Fibre / radio infrastructure**, which should be deployed as per local regulations; but will involve civil works and

these can be costly. The preferred local government contractor is usually recommended as they will be familiar with the assets across the city.
- **Smart City Technologies** can include many different high and low-level sensors and actuators, depending on the use cases for the specific implementation. The minimum requirements for use cases would require Wi-Fi connectivity to the test-bed network and FLAME platform. However, IoT sensors might be required depending on the use case.

On the technology side, the following enablers/technologies need to be in place:

- **SDN Fabric**: FLAME is built on an SDN-enabled networking fabric and implements a stateless switching solution which requires the switches and controller(s) to be at least OpenFlow 1.3-compatible.
  As there is no capability verification alliance for OpenFlow (e.g. Wi-Fi alliance) and the OpenFlow 1.3 features are being considered as experimental, it is highly recommended to double-check with the vendor of the fabric if the following two features are supported: 1) switches support arbitrary bitmask matching via semantically overloaded IPv6 fields; 2) controller supports handling (read and insert) of arbitrary bitmask matching rules.
  On a software-based switching fabric, it is recommended to use OpenvSwitch, which seems not to have compatibility issues. However, hardware switches implement the actual switch in their ternary content-addressable memory (TCAM) tables, which have an OpenFlow-compatible API and only one switch is known to support arbitrary bitmask matching, i.e. PICA8 [21].
  As mentioned before, the chosen SDN controller must a) accept the rules communicated via the REST API and b) insert them into the switches. The following controllers have been successfully tested: Floodlight [22] and OpenDaylight. ONOS does not support arbitrary bitmasks yet.
- **Connectivity**: different forms of connectivity should be considered depending on experiment needs: 1) fibre connectivity, which provides high data bandwidth. Installation is typically static and costly as fibre has to be installed in ducts in the street and might have to run for some distance. Different topologies can be used, such as mesh, point to point, and star topologies (or a combination). Fibre is typically terminated in patch panels and a suitable media converter / small form-factor pluggable (SFP) must be used; 2) microwave links, which give point-to-point connections and RAN connectivity. Links should be planned to allow for buildings and topography and will typically provide up to 1 Gbps aggregated (depending on the technology used). Installation cost is usually low, as physical install is only needed at either end; 3) millimetre-wave links, which give point-to-point connections. Links should be planned to allow for buildings and topography and will typically provide up to 4 Gbbps aggregated (depending on the technology used). These devices are typically less robust than microwave links as they are a newer technology. Installation cost is usually low, as physical install is only needed at either end; 4) fibre optic switches so as to interface between the backhaul technologies and access technologies. A switching fabric is required. Devices should support SDN to allow all devices to mesh together; support VLANs to allow traffic separation and be sized according to fibre / electrical split per location. Normally fibre is used in the ring and Ethernet towards edge devices. Required reliability / resilience can be provided by multiple switches or by more expensive switches.
- **Access Technology**: a method for connecting devices into the FLAME platform is required. Different forms of access should be considered depending on experiment needs. For FLAME, the following should be considered: 1) Wi-Fi access, suitable access points that will allow different network slices to connect with different Service Set Identifier (SSID). Ideally, an SSID has its own unique password that connects to a specific VLAN that is isolated from other slices. Domain Name System (DNS) and Dynamic Host Configuration Protocol (DHCP) will be handled within that network slice rather than by a generic server; 2) LTE / wireless access, where different phones should have access to their own slice so only an Access Point Name (APN) allows specific Public Data Network (PDN) contexts. Ideally, access will be achieved through a private test network with a dedicated virtual Evolved Packet Core (EPC) / EPC cluster that can be interfaced through SDN to other technologies. An alternative proposal would be for devices to access through a commercial cellular network with secure applications that can be connected through the internet to the test network. Typically, devices will either authenticate to a test network or a commercial network, but not both of them; 3) millimetre-wave, included in this section as a potential access technology. Selection of appropriate test devices should be considered to use this access technology; 4) LORA [23], included in this section as a potential access technology. This allows coverage of large areas with multiple devices, however each can only transmit for short periods each day. Typical applications would be to read a meter where the value does not change very often, and the amount of data sent per device is low; and 5) Li-Fi, indoor solution based on visible light communications that can be used to provide coverage in specific areas.
- **Core Network**: the FLAME platform requires compute devices to host VNFs. The compute solution deployed should support integration with orchestration software and SDN controllers. FLAME recommends OSM release THREE [24] as a potential orchestration tool. OpenStack can be integrated with OSM and SDN controllers.
  OpenStack Ocata [25] environment configured for self-service networks and Distributed Virtual Routing (DVR)

with the Neutron OpenvSwitch agent is also recommended. High availability should be available for the project virtual routers that reside on both the controller/network node and the compute nodes. VRRP is used to provide automatic failover of SNAT services to a backup virtual router running on a different node. Only VMs with floating IP addresses will maintain full network connectivity in the case of a network node failure.

- **Edge computing capabilities**: required by the FLAME platform. There are two approaches that can be adopted: individuals can provide edge devices or, alternatively for FLAME, the edge computer could be a shared resource provided by the infrastructure provider. It is important to allow access to the slice within OpenStack, where each customer VNF should live within the respective customers slice. Additional switching will be required in the edge, next to the edge computing resource. The edge switch will need to support VLAN and SDN switching in order to secure exclusive access per slice. The switch will need to be housed and sized appropriately, considering environmental and practical constraints and computing resource capabilities.
- **Infrastructure slicing**: each infrastructure slice is an OpenStack project. VLANs are uniquely assigned to projects aside from where multiple projects require connectivity to the same VLAN. In such cases the VLANs can be setup as shared networks within OpenStack, allowing multiple projects to access the same VLAN. Each OpenStack project is usually created with at least 2 VLAN networks by the administrator. The FLAME OpenStack project should have a VLAN for each of the networks: management, SDN control and data plane. These networks are real city VLANs on the networking equipment and are trunked around the city. This allows non-OpenStack equipment to be able to be connected into a project. Examples are clients' own kit plugged into city racks or wireless devices connected to Wi-Fi AP SSIDs around the city, which are configured to tag traffic with a project's VLAN id. Devices can get DHCP configured according to the project VLAN network's DHCP settings, and will have connectivity to the project's Virtual Machines, subject to the project's internal firewall rules and such.
  Project VLANs can route between each other, and to/from VxLAN networks created in the project, by creating a virtual router and connecting the networks to it. Giving the router a gateway to an external network allows project networks to connect to shared resources in other projects via the "public" network. If needed, project VLANs can also route to other subnets via the main firewall or router, by creating an interface on the firewall in the VLAN (all VLANs are trunked into the firewall). An SDN controller is deployed within the network that can configure VLANs and routing through the system and that will link to the overall FLAME orchestration functionality.

B. Insights from the Barcelona Deployment

The implementation of the FLAME architecture in Barcelona consists of (1) the on-street deployment that provides Radio Access Network (RAN) capabilities, (2) the Multi-Access Edge Computing (MEC) installations to provide light services close to the edge, and (3) the main datacenter (DC) deployment in i2CAT facilities. Main DC IT resources are used to provide heavy computational / storage services, e.g. high definition video content, video transcoding, quality of service and consumption analytics, as well as resource orchestration and management logic, e.g. OpenStack, OpenDaylight, DHCP servers, etc. The Barcelona FLAME infrastructure consists of three main sites that are interconnected by a private network:

- the Omega building, close to the i2CAT premises hosting the main DC,
- the on-street deployment in Pere IV, which hosts the edge equipment and the wireless nodes, and
- the Institut Municipal d'Informàtica (IMI) premises that host networking equipment and are a concentration point of the fibre connections from the other two sites.

Figure 2 shows a high-level topology overview depicting the three main sites.

The upper-left corner shows the Omega building that hosts the main DC infrastructure. The main DC is composed of three servers connected to each other following a star topology via a stack of two switches, which gives us a degree of failure tolerance. The main DC is accessible over a fibre optic link from the i2CAT office and externally from the Internet.

In the lower left part of Figure 2, we see the IMI network that connects the Omega building with the on-street deployment. While the connection from Omega building to IMI happens over an optical network with maximum capacity of 8 x 10 Gbps, the connection between IMI and the on-street deployment goes over a 10 Gbps wire to Pere IV, where the FLAME cabinet setup is located. The on-street deployment is depicted on the right-hand side of Figure 2, also showing the fibre connections between the lampposts and the cabinet.

The on-street deployment consists of the wireless nodes mounted on lampposts that provide connectivity for user equipment over Wi-Fi. The Single Board Computers (SBCs) holding the wireless interfaces are integrated in adapted metal boxes containing also other crucial electronic and networking components, such as AC/DC converters (220V to low voltage for the SBC), fibre media converters (optics to electric), etc. The lampposts are connected via optical fibre with the FLAME edge infrastructure. In Barcelona, the edge infrastructure is deployed within a street cabinet, consisting of an edge server to enable ICN routing and providing VNF capabilities, as well as networking devices that aggregate traffic coming from the lampposts and provide connectivity towards the main DC. The connection between the edge cabinet and the main DC has an intermediate hop in the IMI facilities at Glòries area (Barcelona). The connectivity is guaranteed between such node and the cabinet setup at the Pere IV street via a fibre connection operated by IMI. Moreover, a private fibre optical link

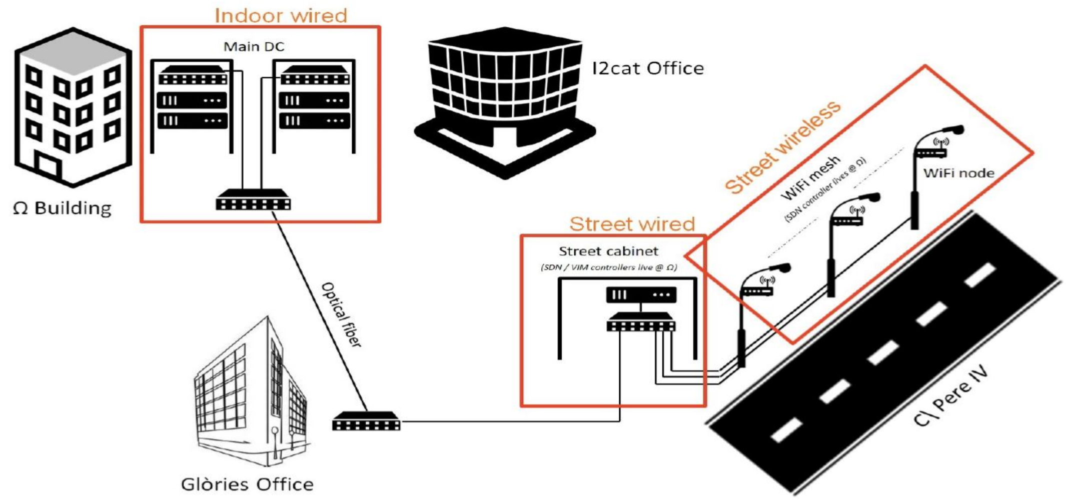

Fig. 2: High-level view of the FLAME infrastructure in Barcelona

(owned by i2CAT / IMI) is placed between the Glòries node and the main DC in the Omega building (Zona Universitaria, Barcelona). In the following, each of the edge deployments (on-street wireless access and edge DC) are detailed.

On-street deployment: wireless nodes on lampposts

For the on-street deployment, the Pere IV street (district of Sant Martí) was chosen. Within this street, a segment of around 400-500m hosts the deployment of the wireless nodes that provide RAN capabilities. Along this segment of the street, a total of 5 lampposts are manually picked in such a way that the nodes are more or less equidistant to each other and they follow a zig-zag pattern, switching from one side of the street to the other.

The lampposts provide both mains power and fibre connections to the wireless nodes. For the deployment of the wireless nodes, a third party designed a casing that fulfils the following requirements:

- Weather-resistant
- Capable of switching from electrical to optical networks
- Capable of converting from 220V mains power (power line connectivity for standard "household" devices) to 48V to power the wireless nodes
- Providing a module that allows remote (hard) reset
- Providing a battery that activates in case of loss of mains power, so the equipment can be turned off safely
- Providing fans for ventilation, to keep the temperature in the casing below any critical threshold

At each case installed on lampposts there are a remote reboot module, the media converter, the AC/DC converter and the battery module. The empty space in the casing left intentionally is reserved for the Gateworks Ventana (GW) 5410 SBC [26], the platform that hosts the wireless interfaces. That is equipped with either 2 or 3 wireless network interfaces of type WLE650V5 (Compex), implementing IEEE 802.11ac standard with backwards compatibility for the IEEE 802.11 a/g/n standards. One interface is always used for the RAN, i.e. it is used to instantiate wireless Access Points (APs), whereas the remaining 1 or 2 interfaces provide wireless backhaul connectivity from each lamppost to its neighbours (thus edge nodes only require one backhaul interface and intermediate nodes require two interfaces to point in two directions). For the RAN, omnidirectional dipole antennas are used; while for the backhaul directive panel antennas are used (both supporting 2x2 MIMO). The SBC also has two Ethernet ports that are both connected to the fibre media converter which enables a wired connection over fibre to the edge cabinet.

Edge deployment: cabinet server and networking devices

The edge server offers application developers and content providers cloud-computing capabilities close to the end users; since having services closer to the end user should improve the user experience. As an example, resources on the edge computing server might be used to support video analytic applications, location services, IoT, augmented reality applications, optimised local content distribution and data caching. In FLAME, a portion of the edge server resources should be allocated for the instantiation of Service Routers (SRs) in the form of one VM per lamppost. SRs are mandatory elements to realise FLAME routing solution. In the Barcelona deployment, the edge cabinet server is a 12-core multi-threading CPU mini-tower server with 128 GB RAM and 2 TB of storage capacity. This machine is registered as a compute node into the OpenStack controller hosted in the main DC. Secure data and control / management communication lines between the main DC and the cabinet server are established.

Besides the cabinet server, the FLAME cabinet router (Cisco ASR920) is mounted in the street cabinet. It provides enough ports to connect each lamppost to the edge server and also the FLAME streep setup with the main DC. Each fibre coming from a lamppost terminates into a Gigabit SFP at the Cisco router side. The FLAME setup, i.e., the edge server and the router with fibre connectivity to the wireless nodes on the lampposts, is connected to IMI over a 10 Gbps fibre. According to Cisco, the ASR 920 series is a full-featured converged access platform which provides a comprehensive and scalable set of layer-2 and layer-3 VPN services. It offers high throughput and low power consumption which makes it ideal for mobile backhaul, business services, video and data applications.

The Barcelona deployment represents a cost-effective city installation where the FLAME solution could provide a significant leap forward for media delivery supporting personalized, interactive, mobile and localized workflows. Leveraging on 5G-enabled programmable infrastructure, FLAME advantages create room for a significant reduction in the overall costs while ensuring fast availability of services towards end users.

IV. Platform Validation

FLAME is validated through a series of experiments and user-centric trials. Over 25 trials of which four initial urban-scale trials are described below. Each trial is designed to capitalise on the benefits of the FLAME Platforms use of software-defined infrastructure, mobile edge computing, flexible service provisioning and routing, and cross layer management and control. The four scenarios demonstrate the potential of the FLAME solution and are representative of the opportunities:

1) **Participatory Media for Live Events**: Media content is a fundamental to experiences live events, from logistics through to recording performances and broadcasting content to friends and other online audiences. The organiser must make sure that all practical information is clear and visitors can move safely and meet with friends, whilst media companies want to should engage audiences on location and on line through user generated and professional content. Future interactive systems must to deliver interactivity that engages audiences whilst considering targeting services to communities based on user profiles, preferences, location, etc. The challenge is to engage and interact with audiences, stimulating them to contribute media related to a large-scale live event (e.g. friends they attend the event with, fans of a certain band, people living in the same district, etc).
2) **Personalized Media Mobility in Urban Environments**: Many broadcasters have developed multichannel television offers based on closed interactive digital platforms with high definition at fixed places (e.g. home) and complementary streaming applications to access contents on the move. Personalised streaming from the closed home networks to broadcasting is now an increasing business sector that is dominated by the OTT (i.e. service providers providing their services across the Internet). These cloud-based media service providers are progressively merging with broadcasters (or taking over their traditional market role) and capturing the emerging interest of users. Examples include Netflix, Amazon Prime Instant Video, HBO Go. The challenge meet changing consumers demand to personal media on the move (personal VoD and CCTV recording/streaming) by offering Personalised Media Mobility where personal content follow users on the move.
3) **Augmented Reality Tourist Guides**: Augmented 3D characters embedded into real world locations can offer a range on interactive support to citizens. In this case the challenge is to offer tourists a story-based city-wide quest that helps citizens experience and learn about urban environments through interaction with 3D historical characters. Each user or user group is provided with a personalised narrative according to interests allowing the participant to influence narrative progression, media types delivered and response to the interactions throughout the storytelling.
4) **Augmented Reality Location-based Gaming**: The proliferation of mobile devices that combine computation and graphics processing with video and GPS sensors holds great potential to develop enhance augmented reality games. The challenge is to weave the magic of gameplay into cultural locations such as the cities through location-based games in ways that encourages real-world interactions and gamifies daily commuting activities in a city.

Each application is designed by media service provider who develops a Service Function Chain (SFC) for the required content workflow in each case. Each SFC consists of a set of connected Service Functions (SFs) implemented as virtual machines or containers. The entire SFC is described using a TOSCA [27] template which includes the resource specification for each SF and policies for placing, booting and connecting SFs within the network in response to service demand. The triggers for policies are also described using a separate TOSCA specification and used to configure CLMC processes for monitoring and alerting trigger events. The infrastructure in the cities now have the potential to meeting the ranging demands of different stakeholders by offering localised infrastructure that allows cloud-based applications to be provisioned at the heart of the city itself. FLAME, with the proximity of network hubs providing a full mesh platform, removes the need for media companies to bring local broadcast infrastructures and other services no longer need to rely on global Internet companies to reach audiences.

V. Conclusion

Deploying an urban-scale 5G-enabled media platform is a challenge. In this paper, we presented the approach of the FLAME project to address this challenge by outlining the founding cornerstones of our FLAME platform as well as the use cases that can be covered by it and the benefits returned. We also provided insights into the real deployment at urban scale, not only by outlining a number of considerations every infrastructure provider will needs to address but also by providing insights from our actual deployment in Barcelona. Current work focus on the formalization of those insights in the form of specifications and deliverables, standard contributions (for crucial platform technologies) and trial insights. The latter are planned to commence in early 2019 in the Barcelona and Bristol (UK) deployments.

Acknowledgment

The research leading to these results has been supported by the EU funded H2020 project FLAME (no. 731677) and Spanish national projects ONOFRE-2 (no. TEC2017-84423-C3-1-P) and 5G-CITY (no. TEC2016-76795-C6-2-R). Authors also would like to thank Gregg Young and Mike Matton (Vlaamse Radio- en Televisieomroep - VRT), Gino Carrozzo (Nextworks), Fabio Zünd (ETH Zurich), Steven Poulakos and Robert W. Sumner (Disney Research Zurich - DRZ) for their contributions on the description of the selected FMI scenarios; as well as the Institut Municipal d'Informàtica de Barcelona - IMI for their support on the urban deployment of the equipment that supports the platform.

References

[1] Boniface, M., Nasser, B., Papay, J., Phillips, S.C., Servin, A., Yang, X., Zlatev, Z., Gogouvitis, S.V., Katsaros, G., Konstanteli, K. and Kousiouris, G., 2010, May. *"Platform-as-a-service architecture for real-time quality of service management in clouds"*. In Internet and Web Applications and Services (ICIW), 2010 Fifth International Conference on (pp. 155-160). IEEE.

[2] McKeown, N., Anderson, T., Balakrishnan, H., Parulkar, G., Peterson, L., Rexford, J., Shenker, S. and Turner, J., 2008. *"OpenFlow: enabling innovation in campus networks"*. ACM SIGCOMM Computer Communication Review, 38(2), pp.69-74.

[3] Hu, Y.C., Patel, M., Sabella, D., Sprecher, N. and Young, V., 2015. *"Mobile edge computing A key technology towards 5G"*. ETSI white paper, 11(11), pp.1-16.

[4] Bonomi, F., Milito, R., Zhu, J. and Addepalli, S., 2012, August. *"Fog computing and its role in the internet of things"*. In Proceedings of the first edition of the MCC workshop on Mobile cloud computing (pp. 13-16). ACM.

[5] Gorlatch, S., Humernbrum, T. and Glinka, F., 2014, February. *"Improving QoS in real-time internet applications: from best-effort to Software-Defined Networks"*. In Computing, Networking and Communications (ICNC), 2014 International Conference on (pp. 189-193). IEEE.

[6] Middleton, S.E. and Modafferi, S., 2016. *"Scalable classification of QoS for real-time interactive applications from IP traffic measurements"*. Computer Networks, 107, pp.121-132.

[7] Channegowda, M., Nejabati, R., Peng, S., Amaya, N., Zervas, G., Shu, Y., Rashidifard, M. and Simeonidou, D., 2013, September. *"Design and demonstration of multi-domain, multi-technology software defined networks for high-performance cloud computing infrastructure"*. In Optical Communication (ECOC 2013), 39th European Conference and Exhibition on (pp. 1-3). IET.

[8] M. Boniface, D. Trossen, and M. Calisti, *"User-Centric Media Demands through Software Defined Infrastructures"*, NEM Summit, Portugal. 23-25 Nov 2016.

[9] Poulakos, S., Sumner, R., Crowle, S., Boniface, M., Phillips, S., Young, G., Matton, M., Carrozzo, G., Zuend, F., Trossen, D. and Garcia, J., 2017. FLAME D3. 1: FMI vision, use cases and scenarios. Available at https://www.ict-flame.eu/deliverables/

[10] Boniface, Michael, Modafferi, Stefano, Voulodimos, Athanasios, Osborne, David Salama and Murg, Sandra (2017) EXPERIMEDIA a multi-venue experimentation service supporting technology innovation through new forms of social interaction and user experience. In, Chen, K.C. and Shukla, Sandeep (eds.) Building the Future Internet through FIRE 2016 FIRE Book: a Research and Experimentation based Approach. pp. 267-286.

[11] Definition for interactive media. Available at http://www.atsf.co.uk/atsf/interactive_media.pdf

[12] ETSI, Network Functions Virtualisation. Available at http://www.etsi.org/technologies-clusters/technologies/nfv

[13] Open Networking Foundation, Software-defined Networking (SDN) Definition. Available at https://www.opennetworking.org/sdn-resources/sdn-definition

[14] Trossen, D., Robitzsch, S., Boniface, M., Edo, M., Alberto, C., Carrozo, G., Chatain, J., Tomas, A., Matton, M., Faca, F.M. and Poulakos, S., 2017. D3. 3 FLAME platform architecture and infrastructure specification V1. Available at https://www.ict-flame.eu/deliverables/

[15] OpenDayLight controller. Available at https://www.opendaylight.org/

[16] Specification of Management and Orchestration. Available at https://www.etsi.org/deliver/etsi_gs/NFV-MAN/001_099/001/01.01.01_60/gs_NFV-MAN001v010101p.pdf

[17] ONOS controller. Available at https://onosproject.org/

[18] OpenFlow. Available at https://www.opennetworking.org/sdn-resources/openflow

[19] QUIC, a multiplexed stream transport over UDP. Available at https://www.chromium.org/quic

[20] FLAME D5.1: FLAME Replication Process v1. Available at https://www.ict-flame.eu/deliverables/

[21] Pica8, PICA8: Programmable Internetworking and Communication Architecture, Infinite(8). Available at http://pica8.com

[22] Floodlight controller. Available at http://www.projectfloodlight.org/floodlight/

[23] LoRaWAN specification 1.1. Available at https://www.lora-alliance.org/what-is-lora

[24] OSM release THREE. Available at https://osm.etsi.org/wikipub/index.php/OSM_Release_THREE

[25] OpenStack Ocata. Available at https://www.openstack.org/software/ocata/

[26] Gateworks Ventana 5410 Datasheet. Available at http://www.gateworks.com/imx6-single-board-computers-gateworks-ventana-family/item/ventana-gw5410-network-processor

[27] TOSCA Simple Profile in YAML Version 1.1. Available at http://docs.oasis-open.org/tosca/TOSCA-Simple-Profile-YAML/v1.1/os/TOSCA-Simple-Profile-YAML-v1.1-os.html