FACILITY FOR LARGE-SCALE
ADAPTIVE MEDIA EXPERIMENTATION

ICT-FLAME.EU
@ICT_FLAME

FLAME

24 MONTHS OF OPEN CALL EXPERIMENTATION

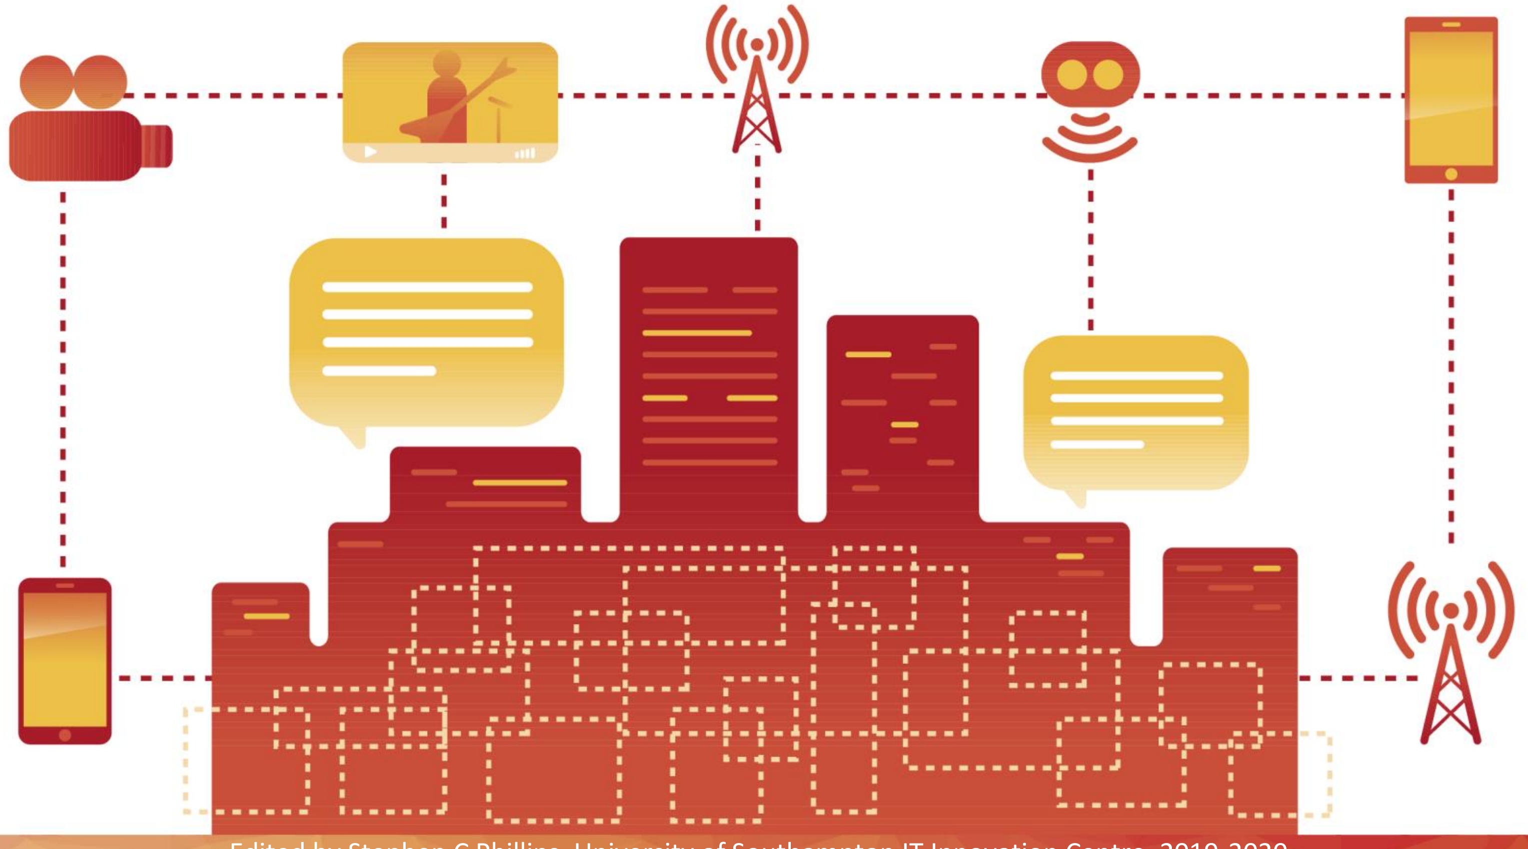

Edited by Stephen C Phillips, University of Southampton IT Innovation Centre, 2019-2020

24 Months of Open Call Experimentation

November 2018 – October 2020

FLAME received funding from the European Union's H2020 research and innovation programme under grant agreement No 731677

24 Months of Open Call Experimentation

November 2018 – October 2020

Introduction

Over 24 months, FLAME hosted over 20 experiments funded through four open calls with awards ranging from €50k to €130k. Four start-ups, twelve SMES and four industrial partners were funded in nine verticals.

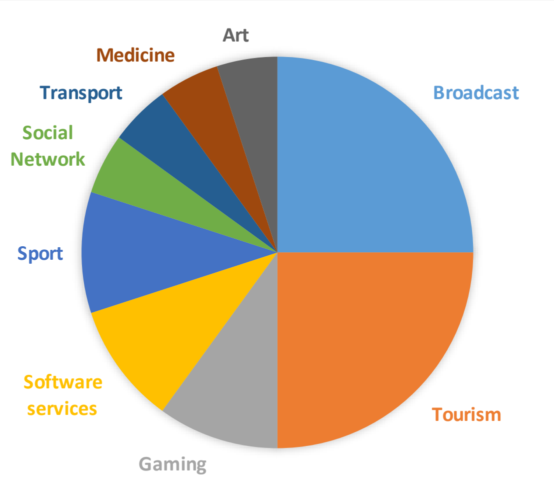

Benefits to Experimenters

Experimenters made full use of the opportunity to develop applications for FLAME with many getting a first exposure to micro-service architectures, mobile edge computing and 5G.

The focus on controlled experimentation and trials with end-users in physical locations made experimenters carefully evaluate their work with respect to QoS and QoE and test in realistic situations.

The results of FLAME influenced the development of a product or service in all cases with plans to adapt 7 existing products and 16 new ones.

Lessons Learnt

Supporting 20 companies in developing and deploying dynamic service function chains in a distributed compute environment was a challenge. FLAME's success in this endeavour was in large part down to clear processes.

The dev-ops workflow to go from the desktop to user trials via FLAME-in-a-box and the FLAME sandpit testbed was essential, and this approach feeds into the software engineering and the experimentation community NGI-EXP.

FLAME's powerful yet relatively simple approach to configuring an experiment's resource allocation was welcomed by experimenters. The average '*service creation time*' is one of the main KPIs defined by the European 5G PPP initiative targeting a reduction '*from 90 hours to 90 minutes*' and the FLAME project is contributing to it.

Experimenter Recommendations

As a result of early experiences and requests, FLAME added emulated UEs to each replica to allow experimenters to test their services in situ before travelling to the site. Over time documentation and processes were improved in response to suggestions.

Some recommendations, such as the addition of GPUs could not be accommodated but should be considered for future deployments.

Benefits to FLAME

The quality and variety of innovative experiments using the FLAME platform has validated the concept of strong integration between media services and programmable network and computing infrastructures.

The funded and unfunded experimentation activities produced a body of work demonstrating technology readiness (TRL7) of the FLAME platform components service routing, cloud native orchestration and lifecycle management for the telco cloud. The remarkable number of customer-facing trials is key in the standardisation work in IETF, NGMN and ITU demonstrating feasibility and readiness.

FLAME received funding from the European Union's H2020 research and innovation programme under grant agreement No 731677

Open Call

1

FLAME received funding from the European Union's H2020 research and innovation programme under grant agreement No 731677

FLAME's first open call provided awards of €70k for SMEs and €130k for more extensive work by industry partners (TNO and Ficosa) with trials to be performed in Bristol and Barcelona.

Eight Bells Ltd: RAISE

Augmented reality for POIs and navigation with content created by users and stored at the edge.

Ficosa: V2FLAME

Video of an incident relayed live from one car to others in the vicinity.

HOP Ubiquitous: BeMemories

Interactive tourist guide co-created by citizens with content placed at specific locations and only made available to mobile users at those locations.

MOG Technologies: MEMPHIS

Mobile users (citizen journalists) streaming live video to consumers on the internet. Edge services ingest and provide a video switcher to permit a producer to select the stream to be broadcast.

TNO: Prosumer Video Distribution

Mobile users at an event streaming video to other local users using video processing and AI to filter poor quality content, identify landmarks and create aggregated feeds. Consumers get personalized content suggestions.

UbiWhere: EdgeTube

Mobile users at an event streaming live video to viewers on the internet. Edge services ingest, analyse for quality and event coverage of camera shot compared to others with feedback to the user and streams started and stopped as required by the system.

FLAME received funding from the European Union's H2020 research and innovation programme under grant agreement No 731677

Summary & Objectives

RAISE: Augmented Reality Smart Guidance Application based on Edge Computing

The RAISE experiment proposed a novel mobile phone application offering enhanced guidance in urban environments and live augmented reality (AR) video streaming, thus opening the door to a whole new range of services.

Targeting an interactive, personalized AR experience, the RAISE application connects digital content with the physical world in order to empower 3D navigation, real-time geotagging and geofencing, as well as launching of adverts and multimedia when the user is at specific locations.

The RAISE application was designed and developed by integrating open source libraries for map data, geo-storage, camera posing and video display. The application leveraged the FLAME wireless infrastructure in Barcelona for off-loading complex operations such as camera pose tracking and AR frame compositing, and caching hyper-local data, including local video data placed by users as geotags, thus reducing network congestion and improving application performance.

Key Results & Insights

Development of the Android application was finalized, and all components were implemented using the open source libraries ARCore, Sceneform. The application was developed with the ability to display POIs from Foursquare, add and display custom user augmentations (storing them in a MongoDB database) and stream MPEG-4 video on AR surfaces, using the Android media player and the latest build of Sceneform.

FLAME Insights

Benefits

RAISE adopted a bandwidth of reference with a given video resolution and frame per second in order to understand the cost/benefit analysis of using M-JPEG with various compression levels. The video captured by the smartphone was transmitted to the FLAME MEC server either as raw or encoded data, using intra-frame or inter-frame compression techniques,

- Mar – Sandpit testing
- 4 Apr – On-site test
- 10 Apr – Trial, 6 people

Funded as part of open call 1 by FLAME
FLAME received funding from the European Union's H2020 research and innovation programme under grant agreement No 731677

experiencing low end-to end delay (in the order of 10ms).

At the same time, the energy consumption with no other applications open was at normal levels, i.e., no major reduction was observed at the battery usage of all the smartphones that were used in the demonstration phase.

Support

Even from the contract preparation phase, the FLAME project was very willing to start the discussions and to provide descriptions, advices, and guidelines relevant to the FLAME infrastructure and the Barcelona testbed where the RAISE experiment was targeting.

Before testing on site, an initial packaging of the developed services and its local deployment in the FLAME-in-a-box environment as well as on FLAME sandpit in Southampton has taken place without any problems, to verify that packaged SFC, bundled as a collection of virtual machines and a TOSCA specification, deploys, initializes and serves correctly using FLAME's advanced platform services and SDN based network.

Overall, the interaction with the FLAME project was very good, with the mentoring team being always available to answer all kinds of questions. The FLAME documentation has been used conveniently and efficiently.

Recommendations

The functional components of FLAME and of the Barcelona platform as a whole have been used, assessed and validated during the experiment setup and execution. No errors or imperfections have been encountered during the experimentation phase. Though, the MEC capabilities of the FLAME radio cells can be increased to allow faster service provisioning.

The open call process has been smooth and the reviewers' comments were valid.

Next Steps

The RAISE innovative application used the FLAME MEC infrastructure to achieve faster access to media and services, lower latency, and higher embracement of the experience.

8BELLS has decided to explore the first mover opportunity by developing an AR smart city application, exploring 5G MEC installations and giving rise to new and disruptive applications that will offer enhanced guidance in urban environments and live AR video streaming.

Targeting to an interactive, personalized AR experience, the RAISE application can now be be used to connect digital content with the physical world in order to empower applications like 3D navigation, real-time geotagging and geofencing, as well as launching of advertisements and multimedia when the user is at specific locations.

Funded as part of open call 1 by FLAME
FLAME received funding from the European Union's H2020 research and innovation programme under grant agreement No 731677

V2FLAME | FICOSA

May 2019 – Jan 2020

Summary & Objectives

V2FLAME experiment merges Cellular V2X and FLAME networks to provide better, localized information of a potentially risky situation to a driver.

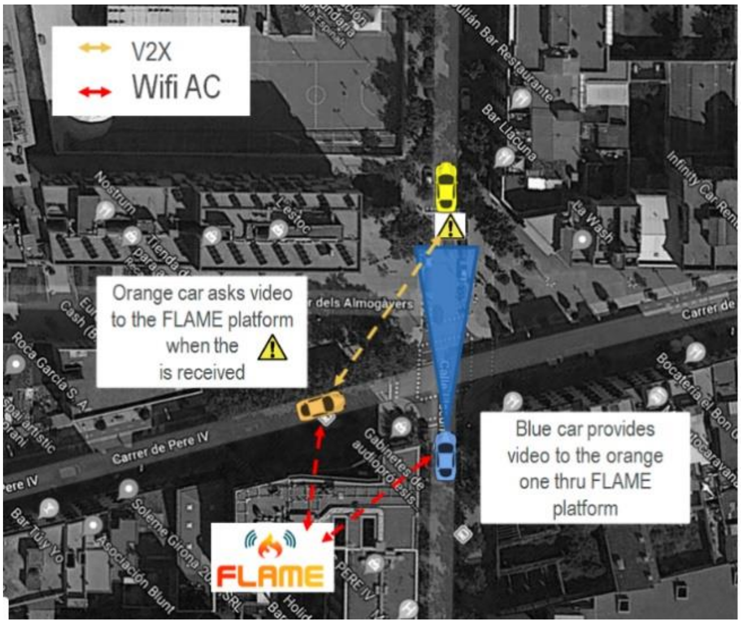

V2X technologies provide cars with a mechanism to share information between them and with infrastructures. However, V2X is designed to share small messages and is not able to broadcast video feeds due to the stringent use of resources required.

Merging V2X technologies and media capable networks such FLAME, we can have the best of both worlds: cars can receive warnings of any incident ahead, while having available a live feed video broadcasted from the infrastructure and linked to the incident. This video can be played to the driver whenever the car is stopped, and thus providing him/her relevant information (number of lanes blocked, the cause, etc.).

In V2FLAME experiment, Cellular V2X technology is used to communicate a disturbing event (stationary vehicle) to the car. The event information is showed in a map and a video feed showing what is going on.

With that information, drivers can decide to either go for an alternative route or be ready for moving to a different lane to avoid the obstacle.

The experiment performed by FICOSA uses Barcelona site on Pere IV Street where:

a) A stationary vehicle (yellow car) sends a C-V2X message when Vehicle Warning event takes place.

b) An approaching vehicle (orange car) receives the message in the C-V2X stack.

The FLAME architecture is then used to identify and broadcast the video feed retrieved from an in-car camera looking forward, covering the stationary car (blue car). This video is generated locally, and multicast to the cars which are heading towards it.

c) C-V2X stack asks FLAME for video feed covering the coordinates provided in the C-V2X message.

Funded as part of open call 1 by FLAME
FLAME received funding from the European Union's H2020 research and innovation programme under grant agreement No 731677

d) FLAME provides the video feed coming from blue car to the approaching vehicle.

e) In the orange car, all the information is shown to the driver in the developed Human Machine Interface (HMI).

Key Results & Insights

The experiment has demostrated so far in the first tests that FLAME can enhance V2X use cases, providing discovery of the best camera available and the right video feed in V2FLAME experiment

FLAME Insights

Benefits

Broadcasting real-time videos is not possible with C-V2X networks due to its high-demand for resources.

FLAME platform enables local, low-latency, high-bandwidth multicast to automotive sector, thanks to its integration with C-V2X network.

FLAME platform allows providing videos to many drivers at the same time without overloading the network due to data exploitation is closer to the source, making communications faster than using centralized Cloud infrastructures.

Moreover, this opens the way to new, more immersive applications where the driver benefits from AR-like scenarios such as "see-through" where a stopped vehicle occludes a relevant part of the field of view. The occlusion is minimized or removed by providing a video that can be overlapped to the occluding area.

FLAME takes advantage of the MEC concept, providing not only media capabilities but also hosting other services.

Support

The FLAME project gave support by assigning a Mentor from the project to the experiment. His guidance has been very useful to properly generate the experiment.

Recommendations

FLAME has proven to be a great way to deliver video content and providing intelligence in the network. However, the complexity of developing and deploying a service can be a bit high for people not used to deal with MECs and services.

Next Steps

FICOSA is a multinational company developing, among others communication units for cars (TCUs).

Traditionally oriented to low bandwidth services on 2G and 3G, nowadays TCUs are the main gateway for communications of the car including functionalities like Wi Fi or V2X technologies

This experiment helps improve V2X functionality and use cases, influencing the requirements of V2X units, which will use future networks such as 5G and services like FLAME.

Funded as part of open call 1 by FLAME
FLAME received funding from the European Union's H2020 research and innovation programme under grant agreement No 731677

BeMemories | HOP Ubiquitous

Nov 2018 – Jun 2019

Summary & Objectives

Be Memories builds a new communication channel between the city and citizens/visitors through the smartphone, allowing users to exchange the immaterial heritage harboured in the memory of the people, in a multimedia-centred way. Be Memories shares and visualizes one-minute videos, created by citizens for the visitors based on stories related to the Points of Interest (POIs) of the city.

The concept has successfully been deployed in Bilbao and Ceutí using the HOPU hardware. In FLAME the custom hardware is replaced with the edge and 5G technologies that FLAME provides. The experiment was to compare the approaches and understand the effect of technical differences on the user experience, development facilities and functionalities.

Key Results & Insights

Time to start playing video

A key metric for users is the time taken to start playing a video from the start of the application loading. In the FLAME trial this was 435ms (slightly lower than in Ceutí). The Bristol infrastructure was superior when multiple users were connected. In both approaches the delay is imperceptible to the user.

Deployment time

The FLAME orchestrator deploys images (from nothing to started) to all edge nodes in 8 minutes. In Ceutí this process takes more than 30 minutes for each node.

Delivery of high-quality video

The FLAME infrastructure supported the delivery of Full HD/50fps video to 9 simultaneous users with almost no interruptions. The customer infrastructure of Ceutí recorded many interruptions in the same scenario. The few interruptions observed in Bristol averaged 370ms compared to 1285ms in Ceutí.

Enhanced data privacy

In both Bristol and Ceutí the video data is held locally at the point of viewing. This enhances the privacy aspects of the data compared to conventional cloud solutions such as YouTube.

Good user acceptance

98% of respondents said they liked the concept of local cultural content channels 50% rated the application with 5 stars.

Mar — Sandpit testing

1 Apr — On-site test

3-5 Jun — Trial 50 people

Funded as part of open call 1 by FLAME
FLAME received funding from the European Union's H2020 research and innovation programme under grant agreement No 731677

FLAME Insights

Benefits

The main advantage of FLAME in this scenario is the remote deployment of the service function endpoints into the edge devices. The TOSCA file sent to the orchestrator makes it very easy to designate the clusters where the service functions were deployed, as well as the resources for each SFE.

The deployment of the different smart POIs can be generated with only one software image. The image is adapted to the deployment using FLAME's WhoAmI service. This process avoids the pain of replicating the image with minor changes. The videos used in each POI are held in central service function which supplies the data to the edge nodes.

The CLMC as a central metric storage system provided by the platform, helped the recovery of all data recorded during the execution of tests. The services deployed in FLAME report metrics to the CLMC from which they later can be extracted as CSV files to extract data from the platform.

Support

The dev-ops pipeline, from the media-service tutorial, FLAME-in-a-Box, through the sandpit to the Bristol replica itself worked to get the software ready for the deployment. The project documentation provided a solid understanding of the process and how to adapt to the FLAME platform but the mentor and IDE were on hand to resolve any specific issues.

Once the application was ready to test in the Bristol environment the University of Bristol staff provided physical and network access and helped test the application and liaised with the owners of Millennium Square to find the best opportunity for the trial.

Recommendations

FLAME could provide a simple way to ensure users were connected to the nearest access point (based on GPS position for instance).

Being able to access the interior of deployed service function endpoints (for debugging purposes) would have been an advantage.

A set of automated triggers to auto-scale resources would have been useful.

Next Steps

The work was disseminated at the Barcelona Smart City Expo, FI-WARE Summit and IoT Week as well as through a paper at the WinCOM Conference, Marrakech.

BeMemories is a product for smart tourist destinations in the Spanish market. The FLAME project has provided knowledge to help adapt the solution to international markets.

HOP Ubiquitous may take up the use of the TICK Stack data collection software used in the CLMC to keep ensure privacy of such metric data. The current deployment in Ceutí has been updated as a result of the FLAME project work.

Funded as part of open call 1 by FLAME
FLAME received funding from the European Union's H2020 research and innovation programme under grant agreement No 731677

MEMPHIS | MOG Technologies

Nov 2018 – May 2019

Summary & Objectives

In the aftermath of past and recent events (such as the incidents at the Bataclan, in November 2015, and the fire in Notre Dame, in April 2019), in which citizens revealed, first-hand, photographs and videos of those events, crowd journalism has been gaining increasing notoriety based on a common premise - ensuring a fruitful cooperation between citizens and media producers/ journalists to rapidly create news and/or to provide different personalized perspectives of breaking news.

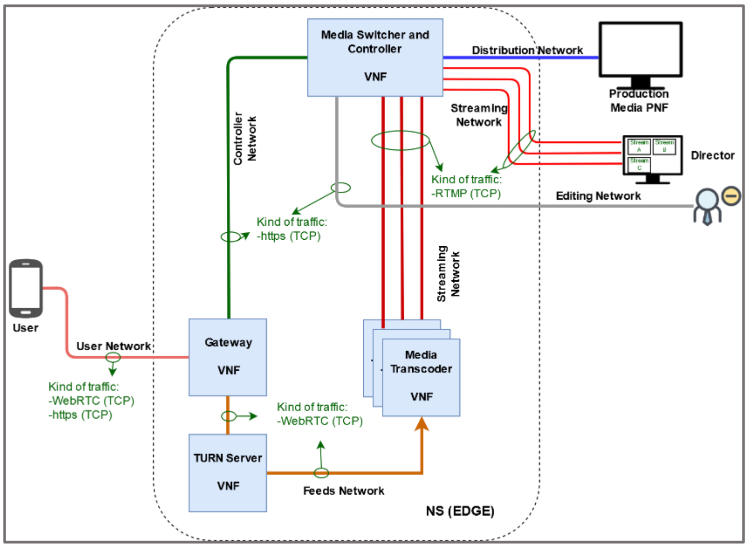

Considering the existing constraints to a wide implementation of crowd journalism as well as the limitations that current market solutions exhibit, **MEMPHIS - Massive Engagement of Media Producers for Highly Innovative Smart cities was designed to provide and test a crowd-oriented Cloud-based ecosystem that can be used by both citizens and media professionals to create added value live media.**

At a certain breaking news location, the media producer should ensure the existence of an appropriate wireless coverage infrastructure, with an adequate number of access points to which the citizens can connect their smartphones to start streaming live video feeds. Users, provided with an application, can connect to different access points, start capturing what they consider to be the most interesting images and streaming them out to the MEMPHIS application operating on top of the FLAME platform.

The multiple video feeds are then fed to a control room in which an editor/director can visualize them in a multiscreen view and choose the most appropriate to be distributed to the Internet or sent to large video screens spread out all around the event venue or to another suitable location, such as a TV station's headquarters.

MEMPHIS used the FLAME platform with the purpose of ensuring a way to quickly deploy an infrastructure (Cloud, edge), provide the necessary network slice, configure and run the ecosystem and deploy services that allow low latency and high availability with the necessary QoS for citizens, event organizers, journalists and final consumers.

15 Mar — Remote test

3 Apr — On-site test

7 May — Trial

Funded as part of open call 1 by FLAME
FLAME received funding from the European Union's H2020 research and innovation programme under grant agreement No 731677

MEMPHIS | MOG Technologies

Nov 2018 – May 2019

Key Results & Insights

Pre-trial results

(i) After performing network tests, poor results were observed on network stability and bandwidth when connected, except for the lamp post that was directly connected to the cluster over fibre;

(ii) The limitation to one lamp post could not be resolved, restricting the devices to be close to that lamp post;

(iii) Nevertheless, it was possible to guarantee that the complete ecosystem could work on top of FLAME and the connection could handle at least two devices and one operator.

Trial and post-trial results

(i) High packet drops were observed, even when connected to the well performing lamp post;

(ii) Despite many attempts and considerable efforts to solve the problem, new issues arose, such as a segmentation fault on FLIPS;

(iii) A new approach was applied involving offline recording of the interview on multiple smartphones, restreaming of the recordings to the MEMPHIS ecosystem and following the normal procedure for YouTube live;

(iv) In the end, it was not possible to conduct the trial as planned due to infrastructure and platform obstacles to perform a live demonstration of the MEMPHIS ecosystem. Nonetheless, these problems were identified and these tests provided valuable information to enhance and improve the FLAME platform.

FLAME Insights

Benefits

The FLAME platform was extremely important to understand how the crowd streaming use case could be coupled with an NGI-enabled smart city infrastructure.

FLAME also showed that it is possible to achieve decentralization of video services to the edge of the network and therefore to cause significant impact on the quality of experience by reducing latency, while lowering video production costs.

Support

All necessary support was offered by providing documentation with detailed information and guidelines as well as key technical assistance regarding the usage of both sandpit and on-site infrastructure.

Recommendations

It would be useful to replicate the FLAME sandpit in different infrastructures so that open call experimenters could continuously test their developments in the platform.

Solving on-site network instability issues would also be extremely beneficial.

Next Steps

MOG's participation in open call 1 contributed to the development of a new product - a crowd streaming ecosystem - that will be added to the company's product line.

In order to do that, MOG must extensively test and validate the product features and align it with the market needs and expectations.

Funded as part of open call 1 by FLAME
FLAME received funding from the European Union's H2020 research and innovation programme under grant agreement No 731677

Prosumer Video Distribution | TNO

Nov 2018 – Dec 2019

Summary & Objectives

The TNO Prosumer framework aims to leverage the FLAME Infrastructure to create a localized crowd-sourced multimedia production and distribution workflow to enable today's mobile producers of multimedia content to serve their content in real-time to other users, who may be both located nearby or remote to the producers, while at the same time drastically reducing the bandwidth utilization, both in uplink and downlink, as compared to traditional crowd-sourced multimedia production and distribution services. This will enable regular users in a city to share their experiences in real-time with others while not burdening the network backend.

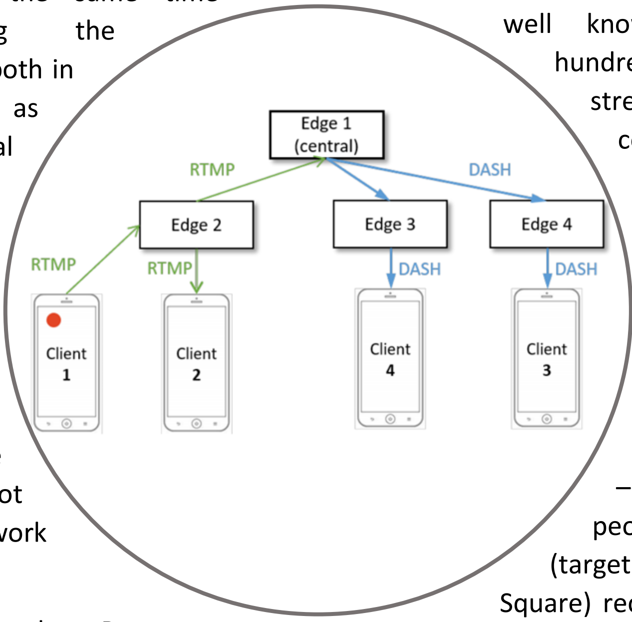

The architecture of the Prosumer framework exploits the edge processing capabilities of the FLAME platform and is based on the idea that the actuality of the content is more important for consumers who are near the action or a place of interest compared to consumers who are further away. With this in mind, the Prosumer framework uses RTMP streams for users watching live content close to the original producer and HTTP Adaptive Streaming (HAS) for users further away. Vicinity between producers and consumers is defined in terms of the edge node they access to: users accessing the same edge node are considered to be close by, while users accessing different edge nodes are considered to be farther away. The FLAME platform ensures that each user accesses the edge node closest to them. Transcoding finds place in one of the edge nodes, which is in our architecture acts as "central node".

The importance of content discoverability is well known; with potentially hundreds or thousands of streams uploaded constantly to such a prosumer frameworks, finding relevant and interesting content is essential for end users. To enable this, we have developed an Image processing module which performs – in real-time – object, people and landmark (targeting Bristol's Millennium Square) recognition on each video keyframe. The resulting information is then sent alongside the video metadata to the Prosumer app, and users can browse the labelled content.

Key Results & Insights

The tests conducted in August show that the solution is able to categorize videos correctly in real-time and that client latency is indeed dependant on consumers' location with respect to producers.

Funded as part of open call 1 by FLAME
FLAME received funding from the European Union's H2020 research and innovation programme under grant agreement No 731677

The measurements also show that there are significant bandwidth savings when using our framework on the FLAME platform, since not all videos are sent throughout the network clearing up bandwidth for the truly valuable video content shared by users. The automatic real-time tagging at the edge also allows for a simple video classification for end users, who get actual and useful metadata to choose the videos they are interested in.

FLAME Insights

Benefits

The platform, when deployed in Millennium Square, allows the testing of the edge computing capabilities in a real city environment, allowing experimentation with actual users during public events. Additionally, the ability to test solutions in advance, first in the sandpit, then in remote on the actual Bristol instance, have been very useful in troubleshooting our solution and the FLAME platform itself from remote, without requiring all the parties to be physically present at the deployment location. To summarize, the real-life deployment combined with the phased testing approach are unique elements of FLAME and, once streamlined and mature, will be greatly beneficial to service providers.

Support

Support was required both for what concerns the (software) platform itself as well as for what concerns the hardware where it was deployed. Also, local support for the organisation of the trials was needed, e.g. in order to find suitable public events, obtaining the consent from WTC (the museum owning the square where the platform was deployed) for trials, and recruiting participants.

The support on the (software) platform was excellent, with quick response times and problem solving occurring in a collaborative spirit. The support on the hardware issues was delivered at a slower pace; since most hardware issues are, by their nature, only detected once we were on site for tests and trials, we were not always able to perform all the planned tests and experiments for a specific onsite visit as a result. Support on obtaining consent from the site owner for testing and trials in Millennium Square was timely and support for recruiting participants for the final trial has been offered well in advance. Partial support was provided on finding and selecting the right public events for trials.

Recommendations

Expectations management, especially in the early phase of an open call project, can be improved, with respect to both the level of support that can be provided, as well as the state of maturity of the platform itself. More sandpit capacity and availability would also be welcome, in order to enable all the concurrent experimenters to thoroughly test their solutions before moving to the test phase. Handover among access points is not yet fully supported and it would be a valuable addition for experimenting in such a large space.

Next Steps

The open call allowed us to improve our technology to a state beyond that of just proof of concept. We were able to test and improve the performance of various of our components and to drastically improve scalability, as well as to demonstrate the benefits of edge computing. As a result of this, we are already in talks with some industrial parties that are interested in using our technology in their service offerings. Additionally, the experience gained within this open call has increased TNO's knowledge position in the areas that combine 5G / softwarized networks and media, as well as edge computing and AI.

Funded as part of open call 1 by FLAME
FLAME received funding from the European Union's H2020 research and innovation programme under grant agreement No 731677

EdgeTube | Ubiwhere

Jan 2019 – May 2019

Summary & Objectives

EdgeTube is a real time video streaming system designed to choose the best set of video streams from a large number of streams supplied by event participants. The set of streams are chosen based on quality and, innovatively, on the geometric coverage of the area in question. That is, if two people standing side by side both stream video of the same event to the system (as judged by their geo-location and device orientation) the best quality stream will be selected. If a poor-quality stream is made available but from a unique angle, then it will be included.

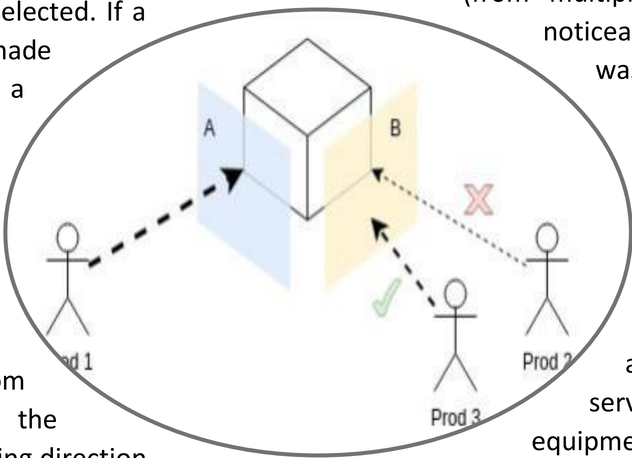

The purpose of EdgeTube is to guarantee that only the stream with best quality will be live from each direction and in the case of a stream changing direction, it will be replaced by another stream covering the same angle. Streams that are not required (lower quality duplications of an angle covered by another device) are stopped streams to ensure that they do not use the bandwidth unnecessarily.

Key Results & Insights

The experiment demonstrated that bandwidth consumption could be reduced for MNO's and distributors.

- Only one stream was published for each channel.
- The initial feedback latency to Producer was <= 6s
- The streaming engine was able to deliver multicast streams
- Service instantiation time was <= 15s after triggering of new scalability event
- Consumer QoE >= 3 (of 5) was confirmed
- The lag in real-time stream switching (from multiple producers) was not noticeable (but no transcoding was used).

FLAME Insights

Benefits

What was unique about the FLAME platform was the automatic placement of services closer to the user equipment and the automatic routing to make use of the close placement, ensuring an optimal latency and fewer hops.

Some of the same things could have been done in other platforms; however, as a whole, the platform provided a set of tools that facilitated the feasibility and viability analysis of the experiment with automated scaling based on a set of parameters, good monitoring solution and tools to facilitate local development. Due

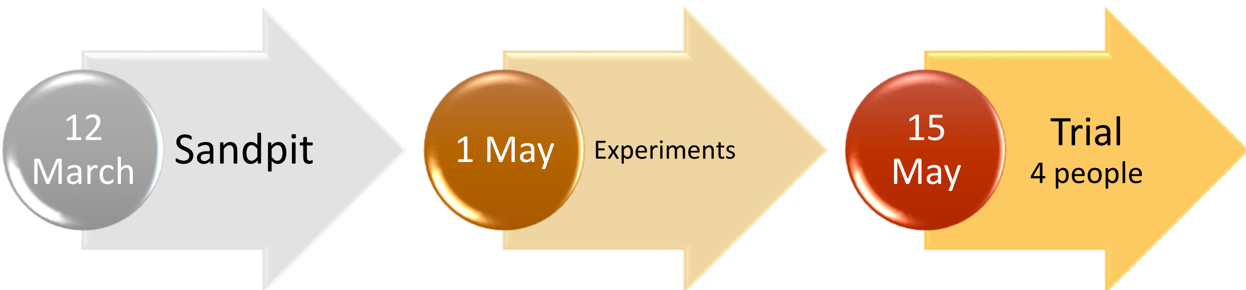

Funded as part of open call 1 by FLAME
FLAME received funding from the European Union's H2020 research and innovation programme under grant agreement No 731677

to the nature of the experiment, some key features of the platform such as intelligent placement and routing could not be exercised.

There were still things that could not be done on FLAME, e.g. real-time video ingestion. An open-source based solution (nginx-rtmp) had to be used to handle the automatic real-time ingestion of media. During the experiment, this feature was not available, only a single FMS for storage (file upload and download) and another service for streaming (assuming a file has already been ingested to disk).

With EdgeTube, FLIPS and the whole monitoring stack (CLMC) were fully utilised to scale the services accordingly with clearly defined metrics. Due to the application's real-time streaming requirements, however FLAME's media services such as storage, media quality analysis and streaming, as originally planned could not be used. Regardless, this had no impact on the project.

Support

Support was provided by IDE to deploy the application on the platform, and from the University of Bristol to execute trial and experiment in their infrastructure.

Recommendations

The missing real-time ingestion service could be of real value when added to FLAME. It would have allowed EdgeTube us to re-use most of FLAME's media service functions such as storage, quality analysis and streaming. Having access to such an ingestion mechanism would ease the development process and allow Ubiwhere focus on implementing a better and smarter algorithm for the scenario. Regardless this was not a major issue at all.

Next Steps

Ubiwhere, as a Portuguese Software development SME, which has co-developed a product called Smart Lamppost (https://smartlamppost.com), may incorporate the findings of the experiment in other similar services which may be of added value to Smart Lamppost in concrete deployment scenarios. As an example, whenever Smart Lamppost is deployed in crowded environments such as concert venues to add network capacity and coverage, the event organiser or sponsor may want to deploy an application such as EdgeTube to better engage with participants, leveraging Smart Lamppost-integrated Edge Computing environment. This integration of compute, storage and network elements inside street furniture such as lampposts in within Smart Lamppost's roadmap. This incorporation may happen, based on Smart Lamppost's future commercial activities.

Funded as part of open call 1 by FLAME
FLAME received funding from the European Union's H2020 research and innovation programme under grant agreement No 731677

Open Call

2

FLAME received funding from the European Union's H2020 research and innovation programme under grant agreement No 731677

Open Call 2

May 2019 – Mar 2020

FLAME's second open call provided awards of €70k for SMEs with trials to be performed in Bristol and Barcelona.

Cognitive Innovations: MoDASH

Mobile users accessing video resources sourced from the internet. Video is transcoded to different qualities and placed at the edge or in metro-DC according to usage.

D-Cube: IMRA

Mobile users watching a marathon streaming live video to local viewers. Edge services process the video using deep learning techniques to recognize athletes and create personalized streams for local viewers.

Gbanga, Millform Inc: ARENA

Multi-player localised augmented reality game with content and game state held in edge services.

IN2 Digital Innovations GmbH: DIGI-HI

Mobile users at an event streaming live and on-demand video to local viewers. Edge services ingest and store video content and create ephemeral local social networks through a recommender system.

Modio Computing PC: SmartMEC

Using deep-learning to augment the decision-making ability of the FLAME platform, predicting upcoming workflows and optimizing content placement.

FLAME received funding from the European Union's H2020 research and innovation programme under grant agreement No 731677

Mobile DASH Services using NFV small cell networks
Cognitive Innovations
Jun 2019 – Nov 2019

Summary & Objectives

MoDASH is considered a solution towards future multimedia services. In particular, we aim to future mobile DASH-enabled virtual reality (VR) applications. The low-latency and reliable distributed networking and computing solutions are addressed with our solution. FLAME platform provides an excellent infrastructure that can accommodate such an experiment with high degree of prototyping and deployment.

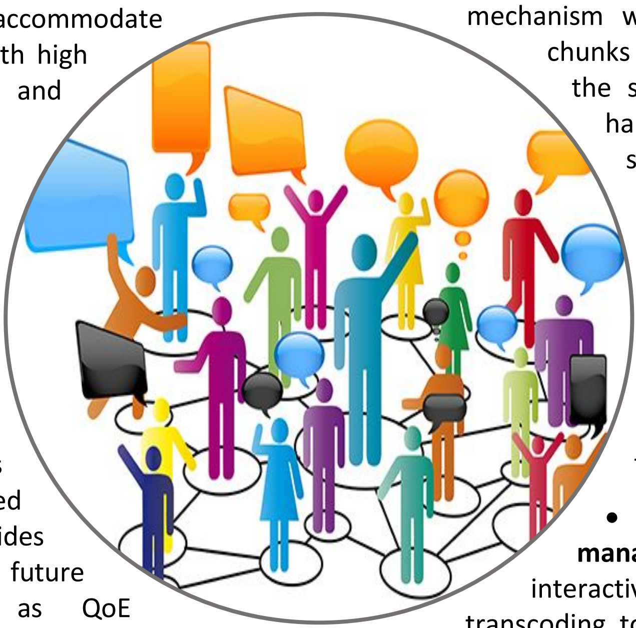

The overall infrastructure allows deploying a solution to the edge of the network orchestrating a Quality of Experience (QoE) for the end users over the wireless access network. The proposed QoE orchestration provides many 5G relevant future functionalities such as QoE parameters reporting and metrics estimation that are relevant to the VR user experience. The proposed QoE orchestration is deployed through the programmable FLAME-in-a-box framework. Future media service providers can support dynamic service request and routing capability. Such a solution can offer the potential to significantly reduce the overall costs while ensuring fast availability of services towards end users.

Overall, MoDASH aims to the QoE orchestration for mobility and interactivity provision in case of future multimedia services such as VR over wireless access networks. The trial will provide the management and orchestration of the following service functions through the EDGE and OMEGA servers:

- **Data offloading:** the data offloading mechanism will transfer and cache chunks of multimedia files at the street cabinet's storage hardware where an edge server will be deployed.
- **Computation offloading:** the mitigation of the streaming and the transcoding services, which are critical for the DASH delivery to the EDGE server.
- **QoE-aware resource management:** in case of an interactive video, video content transcoding to a large set of videos with different bitrates and resolutions is provided.

The proposed QoE orchestration follows a cross-layer design (through the CLMC) approach in order to monitor the QoE delivery and to manage the network resources in order to provide QoE to the connected end users.

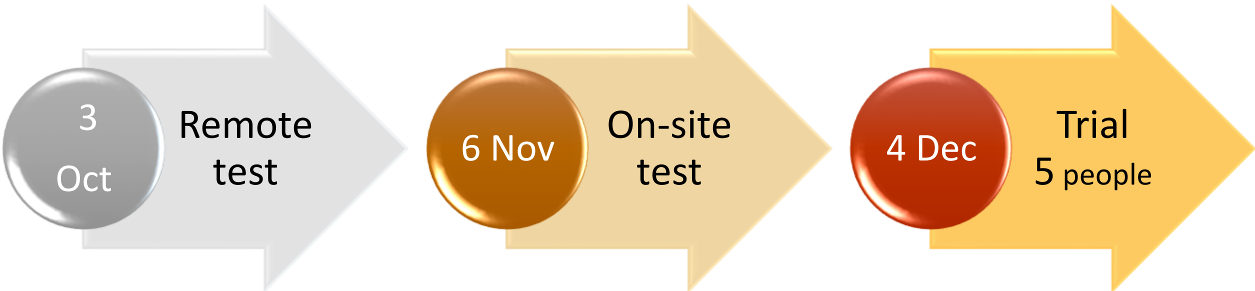

Funded as part of open call 2 by FLAME
FLAME received funding from the European Union's H2020 research and innovation programme under grant agreement No 731677

Key Results & Insights

MoDASH is considered a solution towards future mobile WebVR services over wireless mobile networks. The solution relies on the advanced of NFV to the edge networks and how this is deployed efficiently to the edge. Our remote tests dealt with Key Performance Indicators (KPIs) such as network latency, bitrate, initial playout delay, service response time, service availability, RAM utilization, CPU utilization, Disk storage capacity.

Having deployed MoDASH using FLAME platform to i2Cat infrastructure in Barcelona, we achieved the following key results with the corresponding useful insights:

- Network latency: the average latency of the network connection is low.
- Bitrate: the average bitrate of the user terminal is high
- Initial playout delay: the amount of time needed for a video to start playing after the user completes the request
- Service response time: the average amount of time a service takes to respond a user request.
- Service availability: the overall availability of the service to the end user.
- RAM utilization: the amount of RAM memory that is used on the EDGE and on the OMEGA.
- CPU utilization: the average CPU utilization on the EDGE and on the OMEGA.
- Disk storage capacity: the amount of disk space that remains free on the EDGE and on the OMEGA.

FLAME Insights

Benefits

FLAME platform is a media services platform facilitated by NFV technologies. Future networks are going to be completely softwarized with end-to-end virtualization, where such NFV enabled platforms provide the best way to deploy new services and upgrades the existing ones.

FLAME platform along with the experimental infrastructure give the opportunity to startup companies to deploy and test their solutions in a real world testbed. This is unique for startups to test their platforms and services before deploying into the market.

Support

The support of deploying the MoDASH service was excellent in the course of the trial tests and experiments regarding deploying the solution into Sandpit and then, into the infrastructure.

Recommendations

The platform could be deployed in a 5G-network infrastructure that provides advanced mobility services.

Next Steps

Cogninn exposes high expertise in the area of cognitive networking using software technologies. The MoDASH trial is towards cross-layer solutions to future mobile VR applications that Cogninn targets to incorporate in their platforms given the progress in the standardization in terms of QoE provision and measurements.

We are currrently working in a cloud based WebVR platform for future mobile wev services that is faciliated to the edge. The technologies deployed in FLAME through MoDASH has given useful insight of how our platform will be deployed in the future mobile Internet.

Funded as part of open call 2 by FLAME
FLAME received funding from the European Union's H2020 research and innovation programme under grant agreement No 731677

Summary & Objectives

IMRA, from D-Cube, is about a marathon application for outdoor running events designed to be deployed in 5G infrastructures. During such events, it is impossible for the viewers to monitor the entire race including their athlete(s) of interest. IMRA, is a crowdsourcing application were users are able to stream video content from the race, from their point of view. At the same time, all the available content is being processed in the FLAME platform in order to provide back to the users personalized content regarding the athlete that intrigues them. This solution hopes to provide to the viewers of running events with a personalized, based on their interests, coverage of the event, while providing high-level and sophisticated content that derives from Deep Learning (DL) and Machine Learning methodologies.

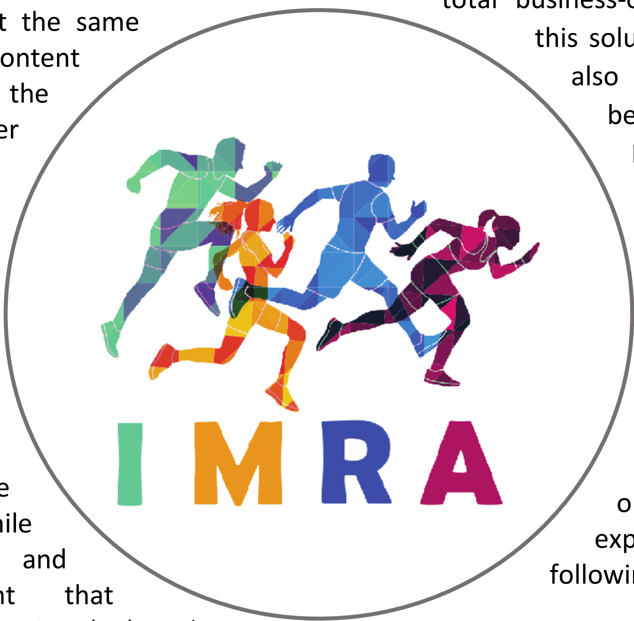

The designed experiment's aim is:

- to provide an intriguing application to marathon event's audience,
- to leverage the FLAME's platform to achieve the best possible QoE for the end-users,
- to experiment by deploying and benchmarking D-Cube's technologies on a new and disruptive 5G platform,
- to provide valuable feedback to the FLAME consortium regarding the feasibility of the deployment of Deep learning technologies on the FLAME platform.

D-Cube plans to exploit FLAME's platform and infrastructure to explore the benefits of a crowdsourced-based media content application for marathon and running events.

The experiment allows D-Cube to benchmark how this solution could be introduced on the market and its benefits in QoS and QoE aspects when compared to traditional cloud-based solutions. Even though, the primary purpose of this experiment is to perform a complex and total business-oriented assessment of this solution, a lot of effort has also been put into the benchmarking of the FLAME platform and infrastructure.

The core outcomes of the IMRA experiment are already envisioned through the objectives. Specifically, the main objectives for the IMRA experiment are the following:

- **O1:** To deploy and test an Immersive Marathon Runner Application (IMRA) that analyses multiple camera feeds in crowded areas
- **O2:** To use the FLAME infrastructure for experimenting with the offloading of computation at the FLAME edge
- **O3:** To enable the preparation of D-Cube offering in immersive sports events upon availability of the 5G networks
- **O4:** Reuse the FLAME offerings and provide feedback on the FLAME consortium

Expected Results & Insights

Considering the availability of 5G networks in the near future, this experiment has been a

Funded as part of open call 2 by FLAME
FLAME received funding from the European Union's H2020 research and innovation programme under grant agreement No 731677

great opportunity for D-Cube to gain valuable experience on the foreseen technologies and their possible constraints. Both technologies, 5G and deep learning (which D-Cube has an expertise on) are drivers of novel media networking applications and thus solutions that combine both are being explored during this experiment. The typical configuration of deep learning requires GPU processing that is difficult and costly to find in edge computing at least in the current configurations. This is was a great opportunity for D-Cube to work with an adaptive innovative 5G platform that has limited edge resources and attempt to offer high QoE to the end-users. Finally, D-Cube will perform experiments in streaming performance against Cloud deployment, in order to provide feedback to the FLAME consortium.

FLAME Insights

Benefits

FLAME provides features that ensure that services will have low latency and smooth scaling-out procedures:

- Trigger-based deployment of services
- Black-box routing based on FQDN
- Offloading at the edge

Support

D-Cube experimented in the performance enhancement of DL components by integrating optimization libraries that caused package components to have bigger size than usual. Support was required in order to handle deployment of those packages in the platform. FLAME support team also assisted in cases documentation was vague or missing. Multicast availability was also requested to be enabled in order to measure performance gains.

Recommendations

It was suggested to provide to experimenters a dedicated packaging server to minimize time between deployments for large components. Also, FLAME platform would be great to have GPU support, in order to provide DL services with low latency.

Next Steps

D-Cube gained experience in deploying state-of-the-art AI services in 5G topologies, where the resources are limited and costly. The derived know-how from this open call is of high value and will be used to enrich D-cube's flagship product “the Immersive Framework”, in order to support 5G topologies for immersive applications.

Funded as part of open call 2 by FLAME
FLAME received funding from the European Union’s H2020 research and innovation programme under grant agreement No 731677

ARENA | Gbanga, Millform Inc.

May 2019 – Dec 2019

Summary & Objectives

With the help of the FLAME infrastructure and its funding, Switzerland-based SME game studio Gbanga is creating a map-based and GPS-driven, multiplayer AR game platform called ARENA that transforms the city into an urban, digital playground. The goal is to develop a smartphone app for Android and iOS that offers different ways to create real-time game zones of different sizes within a city. So called “arenas” which are placed digitally over the cityscape in which you play virtual golf, go on treasure hunts or collect items. In the experiment players of all ages can together to play and connect and get a feel of how future, hyper-local entertainment can work. The strength of the platform is that it provides a toolbox to create new games and local entertainment based on mixed-reality technology. This allows scalability of content and various play experiences. The vision of the platform is that players and companies create arenas all over the world, like creating virtual playgrounds for friends or companies create custom-made team building adventures that fit that company’s culture and values. For example: players could collect virtual coffee beans in the street to get a real-world coffee or a city-wide treasure hunt could be sponsored by a jeweller brand in which a valuable ring is hidden. The experiment of the open call focuses on one game mode on the ARENA platform. The game is called “Chameleon Islands” and is an entertaining experience that turns your smartphone into a binocular. Charlie, his friends and family play hide-and-seek, and as a player you'll have to find them using the smartphone. The game takes place on the several islands, which are hidden in the real world like virtual mirages on Millennium Square and can be made visible with the location-based AR function of the mobile phone. ARENA as platform and the game mode “Chameleon Islands” are unique or rare in games and interactive experiences – and therefore the experiment is highly innovative. The likely potential of a positive reception by the audience has a strong industrial impact. The experiment also tackles transformational aspects in the public like on how the cityscape specifically and also how life in cities in general is experienced and can be more playful.

25 Nov	Remote test
16 Dec	On-site test
17 Dec	Trial 5 people

Funded as part of open call 2 by FLAME
FLAME received funding from the European Union’s H2020 research and innovation programme under grant agreement No 731677

ARENA | Gbanga, Millform Inc.

May 2019 – Dec 2019

Key Results & Insights

Thanks to the power of the FLAME infrastructure, the resilience of our ARENA platform was significantly increased. The possibility to work with the FLAME infrastructure and to collaborate with the consortium allowed to modularise the system and focus development on future-oriented infrastructure. Besides an improved player acceptance and a high quality of experience (QoE) thanks to shorter roundtrip times, we assume that latency and, due to nearest playout, cheat protection is stronger.

Expected FLAME Insights

Benefits

In general, the player always only connects to the local game servers in a FLAME edge node closest to the player. The FLAME infrastructure offers a unique feature that helps to scale interactive experiences like games in a fundamental way: to ability to trigger certain functions via defined behaviours. These rules can be defined in a so-called TOSCA file which allows to orchestrate the FLAME infrastructure dynamically. In the case of ARENA and our game mode “Chameleon Islands” we are able to define that as soon as a certain player base is reached, the load is distributed to additional local servers on servers closest to the player. So not only offers FLAME the ability to define such behaviours, it also allows to do this in an automated and a hyper-local way that saves development time and frees up time for game creation instead.

Support

Our project received guidance by mentors from the FLAME consortium who helped to define performance indicators and helped identify triggers and the proper configuration with TOSCA file as a crucial feature of FLAME to be used in our project. Additionally, documentation, the hosted ticketing service GitLab as well as OwnCloud helped us to have a trusted and central service that also secures data ownership thanks to the Open Source dogma of the chosen collaboration software. During the whole development process, the support was excellent and reported issues were answered and solved swiftly.

Recommendations

Millennium Square in Bristol offers a unique and exciting location to test a game on a modern infrastructure. For our game platform it allows interesting use cases from a game design perspective. A potential expansion of the FLAME infrastructure to a city-wide area would allow new game modes like modes where you capture flags all around the city, or even the country, the continent or the world. From a game developer’s perspective, the FLAME infrastructure could profit if technical standards more known to the game industry could be provided, such as real-time protocols (like UDP or WebSockets) as well as containerization platform Docker, which we were able to use but which meant additional integration time.

Next Steps

Thanks to the Open Call 2, a completely new game mode was created for the ARENA platform which not only means more content for our future product but also has allowed to integrate AR software within the platform itself which allows additional AR games to be released quicker, better and with an efficient and scaling infrastructure. We are most certain that thanks to the smart and targeted support of the open call the platform can be released in 2020.

Funded as part of open call 2 by FLAME
FLAME received funding from the European Union’s H2020 research and innovation programme under grant agreement No 731677

DIGI-HI | IN2 Digital Innovations GmbH

May 2019 – Nov 2019

Summary & Objectives

Say Hello is the result of DIGI-HI (Digital Companion for Localised Interactions). It is an application for festival attendees for creating and locally sharing mix-media stories using the FLAME infrastructure.

In recent years IN2 has been working closely with event organisers, such as festivals, building with them new solutions for increasing engagement between stakeholders. We thus developed new workflows for content and social media management, providing authoring tools for mixed-media storytelling and publishing.

The proposed experiment brought forward a more generic concept that IN2 is exploring for many years. We think that **location** and **time** are two important and currently not very well integrated, components of social interaction and community building. This experiment allowed us to test out our ideas about ad-hoc ephemeral and local social networks and gave us valuable feedback for providing a minimum viable product in this area.

We consider the deployment of 5G technologies as an enabler for such kinds of social networks and a unique opportunity that will divert the attention from global centralised platforms to open, local, ad-hoc communities.

In this experiment we investigated how content which is created and consumed in a geographically bounded way can lead to new experiences and interactions with content and with other people. Testing, validation and trials were carried out on Pere IV, a testbed street in Barcelona and helped us determine how well this type of FMI application is received by end-users and how it can benefit from a new infrastructure like FLAME.

Key Results & Insights

In this experiment we have created the *Say Hello* app and validated a proof of concept around social networks that are **ephemeral** (i.e. created only for a limited time), and **local** (i.e. participants are only included based on proximity). *Say Hello* has been built upon two technology pillars: the media management and publishing framework of IN2 and the FLAME Foundation Media Services (FMS), both deployed on top of the FLAME platform in Barcelona. We have used several FMS in order

Funded as part of open call 2 by FLAME
FLAME received funding from the European Union's H2020 research and innovation programme under grant agreement No 731677

to allow for the quick processing, streaming and sharing of content among the users:

- FMS-Transcoding: for transcoding the video files uploaded by users;
- FMS-Storage: for storing the content (i.e. image and video files) uploaded by users.

Say Hello allows end-users to be part of an ephemeral and local social network that has all four characteristics of personalisation (e.g. personal content and recommendations), interactivity (e.g. comments on stories shared), mobility (e.g. users can use the application as they are moving along the street) and localisation (e.g. only users in that specific location connect and have access).

The main achievements were the successful deployment of the app on the Barcelona infrastructure, recording a high Quality of Service of the underlying infrastructure during the trial, and the achievement of "real-time" scenarios using the FMS.

With the help of the CLMC component we have measured how well the application was performing and estimated from the data-points collected the Quality of Experience for the end-user.

We were able to confirm the usability and the applicability of FLAME in scenarios around a neighbourhood festivals and events to "share what is going on", "connect and explore" and "tune in".

As a direct outcome of the experiment we were able to modularise our media management and publishing platform further, optimising the containerized deployment of its services and evaluating its performance in local deployments with limited resources.

FLAME Insights

Benefits

It was straightforward to build upon the FLAME platform, re-use existing services FMS services and provision the *Say Hello* application in the FLAME infrastructure.

The FLAME tools (orchestration, service provision, management, monitoring) helped us to bootstrap very quickly an enterprise-grade testing environment.

Support

In all experiment phases, we relied heavily on the readily available documentation of FLAME components, API descriptions, roadmaps and guidance for experimenters. Overall the support from FLAME was excellent, all questions were promptly answered and resolved, both in the preparation as well in on-site testing and trials.

Recommendations

Overall the FLAME platform and the tools have performed very well. We were able to efficiently re-use media services and bind them into the Say Hello processing workflow. We would like to have had more actual time experimenting with the platform and been able to test additional service chain configurations.

Next Steps

As a result of the successful execution of the experiment, we plan to package *Say Hello* as a cloud service and deploy it in reference events to continue improving and adding features. Thus, we are looking for innovators and early adopters willing to use the app at their next event, conference or trade fair.

Funded as part of open call 2 by FLAME
FLAME received funding from the European Union's H2020 research and innovation programme under grant agreement No 731677

SmartMEC | Modio Computing PC

Jun 2019 – Mar 2020

Summary & Objectives

Modio is a European SME developing a commercial solution for computationally intelligent resource management of edge compute resources in Mobile Edge Computing (MEC) environments. Our value proposition lies in our technical innovation to cope with large variations of resources' demand, while avoiding either over or under allocating resources in MEC environments.

We have already validated the Proof-of-Concept of our Machine Learning (ML) autoscaling approach for the WebRTC service in a single computing OpenStack cloud in our SoftFIRE and Fed4FIRE+ experiments, using state-of-the-art ML methods. The purpose of our FLAME experiment is to add a new functionality to our commercial ML platform, Qiqbus (https://modio.io/qiqbus_modio/). This new functionality shall assist resource allocation in MEC environments by implementing a flavor of a Deep Reinforcement Learning (DRL) algorithm, called DeepRM, appropriate for distributed MEC environments.

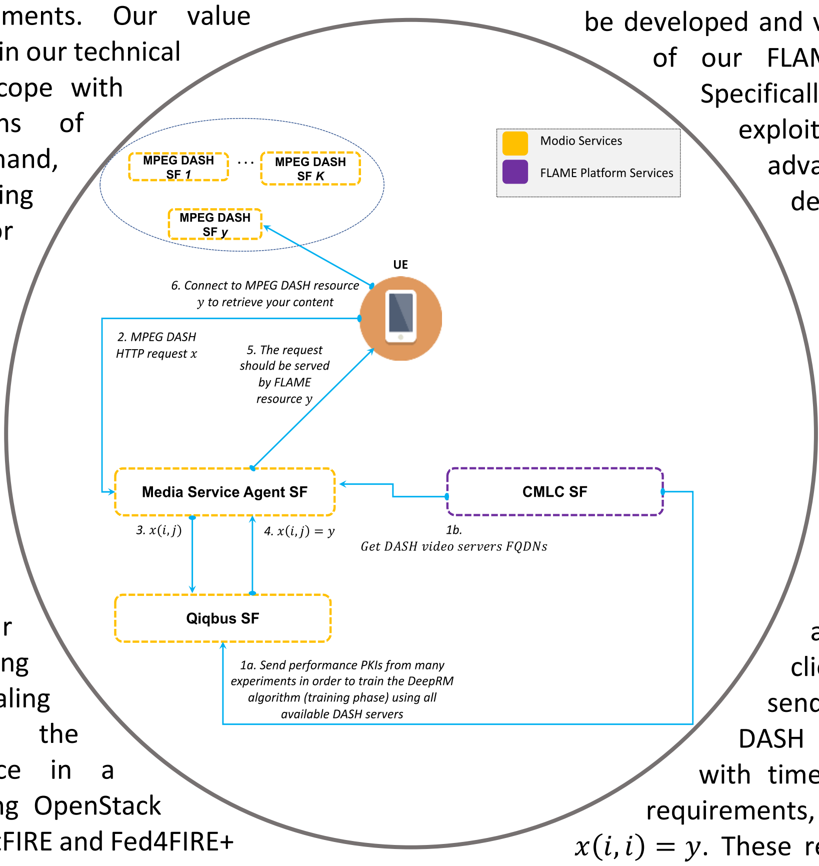

The sequence diagram displayed below, illustrates the main Service Functions (SFs) to be developed and validated as part of our FLAME experiment. Specifically, our aim is to exploit recent advances in DRL to develop a multi-machine scheduler engine (Qiqbus SF), which caters for deciding which FLAME Computing Node $y \in K$ (MPEG DASH SFs) should be assigned to a client (UE), who is sending an MPEG DASH HTTP request with time i and CPU j requirements, denoted as: $x(i,i) = y$. These requirements are known a priori by the client in our implementation, determined during prior experiments. Our novelty provides the unique feature that even if a Computing Node y may be initially assigned to a client, our Qiqbus SF may change the value of y, since decisions are performed on a per-segment basis and not per the whole MPEG DASH video file, as suggested

20th Nov — Remote test

15th Dec — On site test

1st Feb — Trial (3 people)

Funded as part of open call 2 by FLAME
FLAME received funding from the European Union's H2020 research and innovation programme under grant agreement No 731677

by our Mentor, leading to shorter decision times. In these cases, we will send the new value of y to the client through the Media Service SF, using an event-based mechanism.

To perform scheduling decisions on-the-fly, we aim to train DeepRM with training data collected by conducting sets of experiments within our replica in Bristol, employing various MEC congestion patterns. Test-NG shall be used for the purposes of training, while a number of Key Performance Indicators (PKIs) will be collected (through telegraph) from FLAME's CLMC SF. In parallel, we will collect all MPEG DASH video servers' FQDNs from CLMC SF. This way, Qiqbus SF will be aware of the servers that could serve an MPEG DASH request at any given point in time.

Expected Results & Insights

To evaluate the performance of our proposed solution, we plan to measure the mean delivery time of MPEG DASH clients when DeepRM dictates the Computing Node y (note that y might change per DASH segment), where a client should connect to retrieve its content. The performance of our approach shall be compared against the case when a client is connected to the Computing Node with the strongest Wi-Fi signal, i.e. the default behaviour of a FLAME deployment. Given that intuitively our approach avoids allocating client requests to overloaded Computing Nodes, we expect that our solution will lead to a lower mean delivery time compared to the default FLAME client session placement mechanism.

Expected FLAME Insights

Benefits

The MEC environment offered to Modio by the FLAME project is a unique environment, since no alternative exists in the available European testbeds. It consists of a real deployment in Millennium Square in Bristol, UK, which shall allow us to achieve a high TRL at the end of our experiment, especially when trials are involved to validate our software.

Support

The most valuable support was provided to us by our Mentor, who provided us with an overview of the FLAME project and links of the corresponding Gitlab resources, where a detailed list of documentation is maintained for every aspect of the FLAME services.

Recommendations

Our engineers, dedicated to our ongoing experiment, are acting as a software testing team for the FLAME software. Based on their experience with the platform and the supported services, any discovered bugs and suggestions for enhancement will be reported back to the FLAME project, while bug fixes shall be proposed into the FLAME's codebase.

Next Steps

Overall, our experiment's goal is to progress from a Proof-of-Concept of an intelligent resource management implementation, tailored to a single cloud, to an intelligent resource management implementation applicable to an MEC environment, which will be validated with real end-users in the city of Bristol, in the form of a small trial. **Our solution has a very good timing for the market. Currently, there are no alternative solutions in the market. Consequently, our value proposition (to be secure with a patent) is unique.** To capitalise our innovation, we intend that our solution will not be marketed as a standalone software product. Instead, we will market our solution through a licensing scheme, which has strong success record in major ML technology start-up acquisitions, including DeepMind (Google), Magic Pony (Twitter) and Perceptio (Apple).

Funded as part of open call 2 by FLAME
FLAME received funding from the European Union's H2020 research and innovation programme under grant agreement No 731677

Open Call

3

FLAME received funding from the European Union's H2020 research and innovation programme under grant agreement No 731677

Open Call 3
Jan 2020 – Oct 2020

FLAME's third open call provided awards of €50k for start-ups (BI2S, Hool, HTLab, INTELLIA), €70k for SMEs (IGLOR, medVC, Zubr) and €100k for industrial partners (Fincons, Consulservice) with trials to be performed at Bristol,.Barcelona, London and Buseto Palizzolo.

The projects were intended to run from Jan 2020 to Jun 2020 but once the impact of the Covid-19 pandemic became evident, some re-planning was performed in terms of both duration and scope.

Bi2S: Proactive Video Content Caching ICON

Mobile users streaming and uploading video content, with content proactively placed on edge nodes to reduce latency according to predictions given by trained neural networks.

Consulservice: OCTOPUS

Journalists performing outside broadcasts are assisted through real-time face recognition of the interviewee, while the local production unit can select the best video stream and enrich it with text such as the interviewee's name.

Fincons: FC 5 Live

Providing the best of the at-home experience to sports fans in a stadium. Delivering multi-camera live streams, live replay and downloads to mobile devices.

Hool: Chora

Mobile users play a multi-player augmented reality game through just a web browser, controlling a character and learning and sharing moves.

HTLab: Vision

Video processing deployed at the edge helps guide visually-impaired users through a tourist venue as well as detecting if an individual is lying on the ground (and so in need of assistance).

IGLOR: Kalliope

Making live experiences accessible and interactive for massive audiences by bringing services closer to the users.

IGLOR were unfortunately unable to complete their work because of Covid-19 issues. We hope to include a brochure in a revision of this document.

Zubr: Mesh Parade

Fast photogrammetry deployed at the edge enabling an innovative 3D sculpture-garden art experience in a city landscape.

Intellia ICT (HoloFLAME) and medVC.eu sp. z o.o. (medFLAME+)

Please see Open Call 4 for these extended projects' brochures.

FLAME received funding from the European Union's H2020 research and innovation programme under grant agreement No 731677

Proactive Video Content Caching ICON | Bi2S

Jan 2020 – Jun 2020

Summary & Objectives

ICON is a video distribution platform that provides proactive video content offloading to the EDGE of the network. We focus on the provisioning of various video resources to end users. Companies, local municipalities, various stakeholders and end users themselves may produce and upload video files and users are granted access to such content. ICON services are reachable to the end users through a simple web application offered by the ICON solution. ICON's aim is to enhance the FLAME infrastructure and minimize the latency in the provision of video content through a proactive caching mechanism. We have built our solution on top of FLAME's infrastructure, i.e., the core platform where the edge and access points are available in the city of Barcelona. Below we describe the key functionalities of our solution.

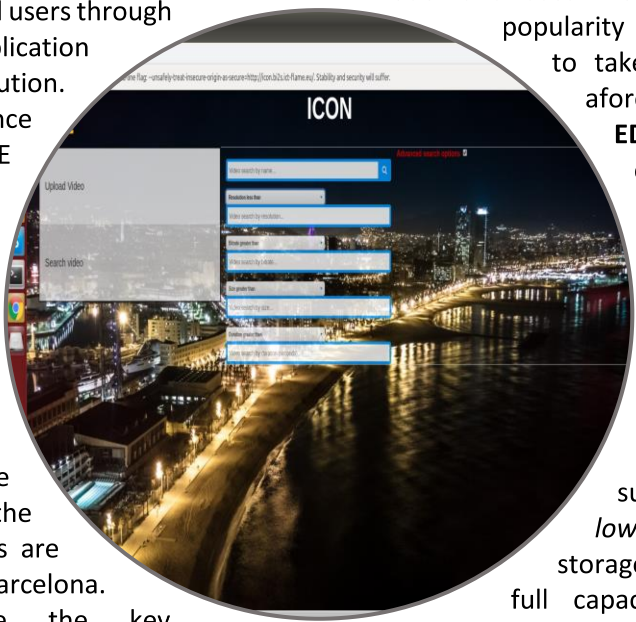

Neural network models for proactive content offloading. We employ two different machine learning (ML) models for enforcing proactive video offloading actions; a feed forward neural network (NN) and a deep neural network (DNN). Each model use inputs related to the user location of the uploaded videos, date of the upload, video popularity and video size. We train the models offline and we use the pre-trained models during the trial in order to offload content to the EDGE of the network. We evaluate each ML performance separately in order to obtain comparative information between the two proposed models.

Caching policy. We deploy a caching policy at the EDGE of the network capable of managing the amount of time each offloaded video stays on the EDGE before it is deleted. Such policy is enabled by the k-means machine learning model and uses the video size and video popularity characteristics in order to take action regarding the aforementioned options.

EDGE storage capacity configurations. We implement four different EDGE storage configurations in order to obtain information on how the NN models treat different levels of disk storage availability. We name such configurations *full*, *low*, *medium* and *high* storage availability options. In full capacity the EDGE server operates normally. In low availability option we tune ICON so that identifies the available EDGE remaining storage in low levels, specifically at 20%. In this sense we simulate a condition in which the EDGE server is overloaded with video files and the remaining storage is dropping in low levels. Similarly, with the medium availability option the ICON sees the EDGE as 50% full while the high storage availability option puts the EDGE at the 80% of its capacity.

15 May — Remote test

* — On-site test

6 June — Trial
4 on-street
+ 2 emulated

*Not possible due to Covid-19 situation

Funded as part of open call 3 by FLAME
FLAME received funding from the European Union's H2020 research and innovation programme under grant agreement No 731677

Video distribution platform. We develop and deploy a video distribution platform on top of the aforementioned functionalities in which users may upload and download video content. Moreover we enhance the platform by employing advanced search options where the users may perform video searches on the existing video base. Such search options include video characteristics as resolution, frame per second, bit rate, colour depth etc.

Key Results & Insights

The KPIs obtained during the trial are the following:

- **Service response time** which is the amount of time a service requires fulfilling a request.
- **Service reliability** which accounts for the percentage of time a service is able to execute the task that is assigned for.
- **User bandwidth** which is the upload/download speed of the users.
- **EDGE/OMEGA Download times** which reflect on the amount of time required for users to download video content when served by the EDGE vs when served by the OMEGA (main DC).
- **EDGE serving ratio** which is the percentage of user requests that are served by the EDGE server.

The evaluation of the performance of the ICON project is based on the aforementioned KPIs that are collected during the trial period. We observed that after a 15 minute period the machine learning algorithms managed to proactively offload the most of the content to the EDGE. As a result 88-90% of the users were served by the EDGE server and thus, experiencing very short download/upload times. We consider such results very satisfactory and we are confident they highlight the efficiency of the ICON solution.

FLAME Insights

Benefits

FLAME provided us with the opportunity to deploy a service function chain in a real-world NFV, SDN enabled infrastructure. With the provisioning of the EDGE server which is critical to our project, the excellent support and the FMS provided to us we were able to deploy and test the ICON solution. The combination of the aforementioned factors was very beneficial and we believe that they comprise the strength of the FLAME consortium.

Support

During the development and trial stages of the ICON project we received excellent support from the FLAME consortium. We have asked for technical guidance on how to use the FMS and how to overcome an Nginx web service problem. The responses we got in both cases were quick and very helpful as they solved the problems at hand.

Recommendations

FLAME could further improve its support to new experiments by expanding its FMS base. Such services could include machine learning and neural network related services that could be used by the experimenters as part of the FLAME platform.

Next Steps

The results we obtained from the experiment will help us to improve our understanding on the applicability of the proposed machine learning algorithms. To this end we plan on using the acquired knowledge on the development of the Bi2S video distribution platform, as well as on disseminating our work in multimedia conferences.

Funded as part of open call 3 by FLAME
FLAME received funding from the European Union's H2020 research and innovation programme under grant agreement No 731677

Summary & Objectives

The OCTOPUS project is targeted to broadcasters and content providers in order to help the journalists with out-of-band information regarding the persons who are interviewed.

The main idea of the OCTOPUS project is to avoid "epic failures" that can happen if the journalist does not recognize the person or in case the journalist does not remember exactly the name or the position of the person during the interview process.

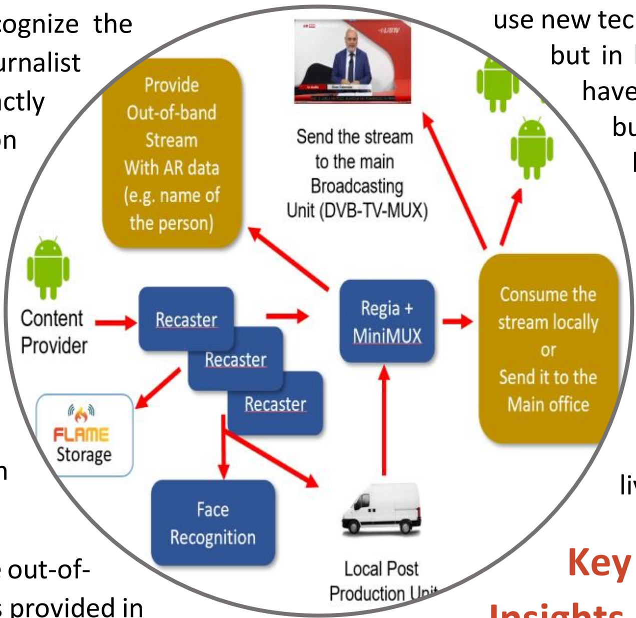

For this reason, the FLAME platform is used to assist the on-the-field work of the journalists and the local production unit, in two main aspects.

The first is related to the out-of-band information that is provided in real time regarding persons who could be interviewed. This information is presented directly to the journalist who can now, if desired, access to a complete description of the person in the interview process.

The second aspect is to support the local production unit, that must select the current flow that should be sent "on air". OCTOPUS uses a lot of the offloading functionalities of the FLAME infrastructure even to support the local production unit. Indeed, the face recognition function is running in real time on all the streams that are produced and then shown to the local production unit which can now also enrich the stream with information on who is interviewed (i.e. add text to the current stream such as the name of the person etc.).

The novelty of the OCTOPUS project is not to use new technologies by themselves but in how those technologies have been mixed together to build a product that can help and support the content production process for on-the-field journalists and technicians who must provide reactions in real time in a not controlled environment such as live events.

Key Results & Insights

Thanks to FLAME a novel architecture based on the edge computing paradigm has been developed. The results on a real testbed has proven that it is possible to achieve face recognition in less than 1 second with commodity hardware (~500ms for SD resolutions). It has also been demonstrated that N:1 bandwidth can be achieved when the streams are processed locally, i.e. send one

20 Sept — Remote test

28 Sept — On-site test

30 Sept — Trial, 2 people

Funded as part of open call 3 by FLAME
FLAME received funding from the European Union's H2020 research and innovation programme under grant agreement No 731677

stream instead of sending N streams to the central production office.

FLAME Insights

Benefits

The FLAME project offers a simple and unique environment for SME and companies that want to implement edge technologies in order to test new ideas or validate products in the field. The testbed environment offers a real platform that is ready to be used in order to speed up the development process.

Support

The support from the FLAME consortium was essential for the OCTOPUS project. Indeed, the FLAME environment is based on technologies such as TOSCA which are not always well known by SMEs and therefore the documentation provided by the FLAME consortium has been very helpful and valuable. The availability of a controlled environment such as the FLAME-in-a-box has also helped the development cycle of the experiment and has helped to understand the service function methodology since the very early beginning of the project.

The support on the real testbed has also been very responsive for any issue that has been encountered during the experiments.

Recommendations

The FLAME project offers a unique environment in order to implement edge computing in a simple fashion. The testbed in Buseto Palizzolo (Sicily) is a great opportunity for testing new ideas and validating products in a real environment. The first aspect that could be improved is the need of more resources in the FLAME sandpit in order to give the opportunity to run bigger and more complex experiments in the development stage before going in the real testbed.

Another aspect regarding the FLAME platform is the possibility to make it closer to the needs of the market with the integration of accelerated resources such as GPU. Indeed, the integration GPU hardware could give better performances for real time stream manipulation, face recognition and heavy computational tasks that are needed in real time experiments.

Next Steps

Thanks to the work done in the FLAME open call 3, a novel architecture has been developed. The edge computing facilities provided by FLAME proved that it's possible

i) to help the journalists and the local production unit with information on the person in the interview process;
ii) to speed up the content production process; and
iii) to lower the bandwidth consumption for the interconnection with the remote production unit.

The next steps are related to the improvement of the performance in the edge network with dedicated hardware and new policies on the collaboration between the edge and the central cloud in order to store as few faces as possible in the edge.

Funded as part of open call 3 by FLAME
FLAME received funding from the European Union's H2020 research and innovation programme under grant agreement No 731677

FC 5 Live | Fincons

Jan 2020 – Jun 2020

Summary & Objectives

As user media fruition is becoming more and more integrated, football clubs, and more generically sports teams, have greater opportunity to deliver their content with the help of **5G** technology at **live events**. Market research has highlighted that fans would like to have the best of the at-home experience available in the venue as well, overcoming the highest rated current limits to their in-stadium engagement, e.g., distance from an event scenario, the limited opportunity to watch replays, and the single view of the game.

Fincons has chosen the **FLAME** platform for experimenting with the **FC 5 Live** app, a media service having new functionalities such as direct live streaming with multi-camera choice, live replays and clip sharing over social networks.

With this industrial trial, Fincons leveraged the capabilities of the FLAME platform and its design patterns, taking advantage of features such as *low latency computation*, *adaptive media delivery*, *multicast http response and dynamic service routing, transcoding, trans-rating and adaptive streaming*, with the final objective to extend its own offering, following the current market strategy for the **Media** sector.

The **FC 5 Live** application aims at delivering **live high-quality video streams** to a very large number of end users in a stadium during sport events by the appropriate use of computing and network resources offered by the FLAME platform and infrastructure. The goals of the trial are to **extend**, **deploy** and **experiment** the Fincons **FC 5 Live** application thanks to the FLAME ICT infrastructure. In particular, the following functionalities of the FC5 Live app have been developed and tested: **Direct live streaming with multi-camera choice,** to watch the game from multiple angles; **Live Replay**, to enjoy again the relevant moments of the match; and **Clip download and sharing on Social Networks**.

The goal is to enable sport teams to provide the same services as a broadcaster, to enrich the customer experience by adding "premium" and "exclusive" content such as, e.g., interviews, highlights, press conference, live comment stream, running commentary and social sharing.

The "FC 5 Live" app has been deployed on the FLAME ICT Infrastructure, allowing users' mobile devices to connect with a versatile infrastructure and taking advantage of the features offered by FLAME, which allows switching from Cloud to Edge when necessary and leveraging the FLAME Design Patterns to optimize resources utilization for the video functionalities.

The FC5 Live App has been tested in the FLAME environment in Barcelona, thanks to the support of the FLAME ICT project partners based in Barcelona, IMI and i2CAT, in a remote trial with 10 users. In this experimentation, Fincons measured QoS and QoE parameters related to media services, which have been used to trigger appropriate reactions in FLAME platform to maintain the QoS/QoE target levels.

Funded as part of open call 3 by FLAME
FLAME received funding from the European Union's H2020 research and innovation programme under grant agreement No 731677

FC 5 Live | Fincons

Jan 2020 – Jun 2020

Key Results & Insights

Throughout the experimentation and the final trial, Fincons has identified KPIs and metrics to measure QoS and QoE parameters related to the media services developed within the app; on this basis, triggers have been defined to determine appropriate reactions in the FLAME platform which aim at maintaining the QoS and QoE target levels.

The system has been tested with increasing numbers of concurrent users; KPIs have been obtained for each functionality of the app, depending on their usage of the system and on different distributions of users on edge nodes. Moreover, useful insights have been gathered regarding QoS metrics in various system configurations, giving an overview of system performances in different load conditions.

FLAME Insights

Benefits

The technical capability of FLAME to cheaply broadcast content to multiple users in a transparent way for user perspective (opportunistic multicast) is essential for delivering media service based on live streaming in large venues, our target use case, where many people assist at the same time in the same location at a live event. The FC 5 Live app relies on the capability of FLAME to deploy services at the edge of the infrastructure, to have low-latency computing capabilities just one hop away from the users, and scale-out to the main Data Centre if needed.

Support

The FC 5 Live project received excellent support both from the Consortium mentor from Martel Innovate, who helped the team designing the system to take advantage of all relevant technical assets that the FLAME platform could offer, and from the local team from I2CAT and IMI in Barcelona, who helped Fincons in all the technical aspects that allowed to remotely run the trial and all the necessary tests.

Recommendations

The capability of FLAME to dynamically adapting the infrastructure depending on the actual usage of the system is very useful, as it allows to delegate infrastructural adaptations to the platform, focussing on the functionalities of the app.

Next Steps

With FC5 Live app, Fincons plans to extend its current media offering to leverage the football fans' experience during live events, offering functionalities that are not currently possible due to the lack of adequate 5G platforms. The FC5 Live app is expected to have a great potential both on "premium" clubs, but also on minor sports, thus contributing to the development of small (SME) clubs, and opening the opportunity to integrate with a range of added value services from third-party (commerce, mobility, etc..), with benefits for the entire value chain, and paving the way for a larger exploitation of the FLAME platform.

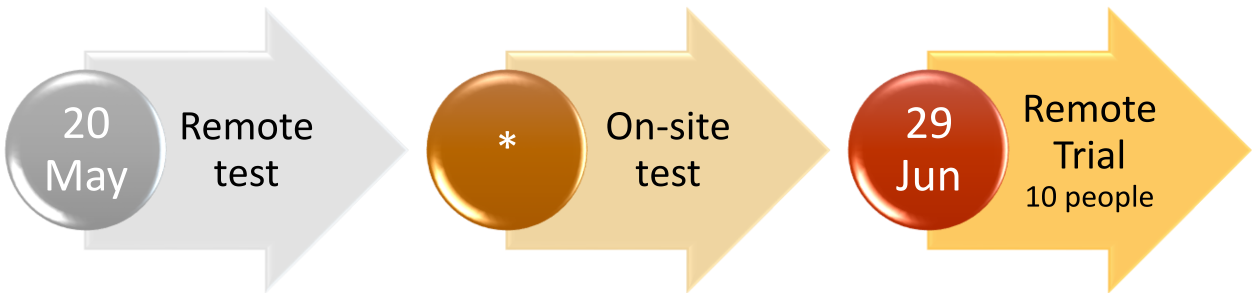

Funded as part of open call 3 by FLAME
FLAME received funding from the European Union's H2020 research and innovation programme under grant agreement No 731677

Chora | Hool

Feb 2020 – Sep 2020

Summary & Objectives

Chora is a multiplayer augmented reality experience played on your mobile through a web browser. Take control of a creature and try to express yourself through movement. Explore and meet others to learn, create and share moves.

***1 -* "Chora" Web App using FLAME & cloud service** Delivering a high-quality multiplayer augmented reality mobile experience with the FLAME infrastructure with a direct comparison to measure up against a cloud version of Chora will allow people to quickly understand the benefits of FLAME 5G. We will gain the knowledge of how to implement a multiplayer game in the web browser using FLAME.

***2 -* Scale service horizontally** We want to understand on a particle level how horizontal scaling can be beneficial to a multiplayer game or art service. To know where the benefits are specifically in controlling resources and offering a hyper-local service for an augmented-reality mobile application.

***3 -* Improved Quality of Service**

The impact of significant improvement to quality of service will be potentially huge for web-based games and art projects like Chora, as it will enable these services to be more accessible without the need to download an app. We are interested in latency reductions by placing service functions at the edge and how dramatic these can be.

Key Results & Insights

The KPIs obtained during the trial were the following:

1. A reduction in latency on the edge which results in better gameplay.
2. A reduction in latency of multiplayer gameplay while utilizing FLAME application comparative to cloud version at specific points.
3. The capture feedback from users.

15 Aug – Remote test

14 Sept – On-site test

1 Oct – Trial, 4 people

Funded as part of open call 3 by FLAME
FLAME received funding from the European Union's H2020 research and innovation programme under grant agreement No 731677

FLAME Insights

Benefits

What is unique in the FLAME platform?

FLAME provided a unique and innovative hyper-local infrastructure that was used by Hool to test our mobile web augmented-reality experience for city-wide use.

What could not have been done elsewhere?

FLAME provided us with the opportunity to deploy a service function chain in a real-world infrastructure. With the provisioning of the EDGE server, excellent support and the adaptive service states (placed and connected) we were able to deploy and test our project. The combination of the aforementioned factors was greatly beneficial, and we believe that they comprise the strength of the FLAME consortium.

Support

What support was required and provided?

Navid Solhjoo was our mentor throughout our project. Sebastian Robitzsch offered us technical support on the FLAME infrastructure. There was extensive documentation offered to us about the FLAME infrastructure online. We had live chat support with Sebastian and Stephen Philips with debugging a telegraf issue, HTTPS support and documentation and with Android app connection Manager issues. The responses we got in all cases were quick and extremely helpful as they solved the problems at hand.

Recommendations

What can FLAME improve?

Moving forward we believe FLAME could improve on offering more support and tools for web browser-based projects.

What cannot be done in the FLAME platform?

Switching a device's Wi-Fi connection from a webpage within a browser is not possible. This is for obvious security reasons. This poses a challenge for a horizontal scaling service that may lie behind an unknown SSID. It forces the user to use a new app rather than using a browser.

Next Steps

Summarise the benefits to you from the open call.

The benefits for Hool from the open call were that it enabled the creation of Chora and for us to experiment with web augmented-reality. It modularized our code base into different service functions. Moving forward, the work has inspired a specific type of gameplay and potential monetization opportunities within Chora relating to event-based user performances in different 5G locations.

Funded as part of open call 3 by FLAME
FLAME received funding from the European Union's H2020 research and innovation programme under grant agreement No 731677

Vision | HTLab

Jan 2020 – Oct 2020

Summary & Objectives

VISION has been realized by HTLab as an experiment of the FLAME platform. The mission of HTLab is targeted to all kinds of impairments and any application of technology to healthcare. VISION is a smart video surveillance framework for smart cities aimed at providing a social and shared video surveillance tool to help impaired people who could need to be monitored and guided safely within the city. Video sources are any kind of video transmitting devices, like IP cameras, action cams, video-cameras integrated in tablets and smartphones or on-board drones. To realize the project, Bristol has been chosen as replica site.

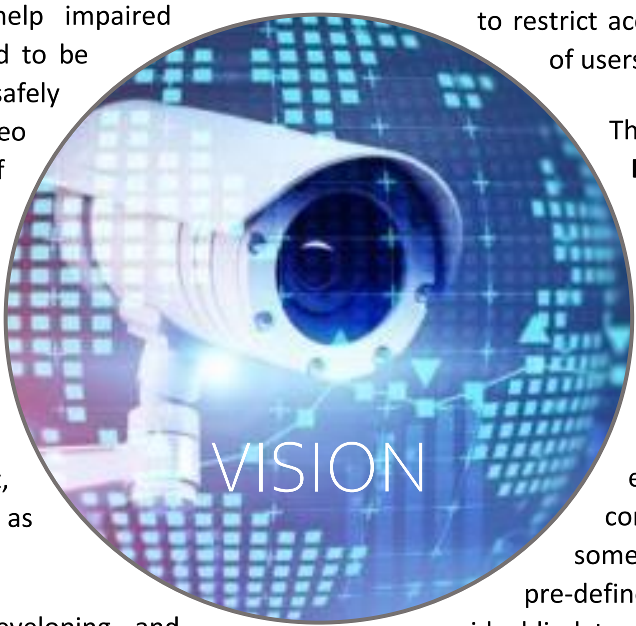

The advantage of developing and integrating the VISION project in the FLAME platform is that FLAME is the first platform that provides simultaneously a softwarized network-based 5G environment with a set of facilities consisting in the availability of Multi-Access Edge Computing (MEC) resources in proximity of users, and of some service design patterns that make service creation very flexible and easy.

VISION is composed of three architectural elements, specifically the DiMoVis platform, Tourist Eyes and the Lying Person Recognition element.

The **DiMoVis platform** is a smart, flexible and social video surveillance platform that, when deployed on the FLAME platform, also becomes highly scalable in terms of number of transmitting and receiving devices. Each user of the platform can easily share his own video source device (e.g. webcam, IP cam, smartphone or tablet) by connecting it to the network and register it to the DiMoVis platform through an Android app or a web portal. During registration, the owner of the video source device can specify one or more access profiles to restrict access to different groups of users.

The objective of **Tourist Eyes** is to provide blind tourists visiting a smart city with a framework for supporting their usual activities, in both outdoor and indoor environments. Thus, the goal of this service element is using 5G-compliant frameworks and some IP cameras installed on pre-defined touristic paths to guide blind tourist people to move in smart cities. Blind tourists, through voice commands issued by an appropriate smartphone app, request indications to reach some points-of-interest (POI) belonging to predefined paths, like restaurants, restrooms, museums, ATM machines. Then, blind tourists are localized by the system thanks to a wearable hat with a specific colour and receive back sound messages from the Tourist Eyes element through their smartphone device to follow the selected path.

The **Lying Person Recognition** element is able to automatically recognize if a person is lying on the ground in indoor or outdoor environments and send an alert to someone (e.g. a friend or a relative of that person, or to the local emergency or healthcare provider).

Funded as part of open call 3 by FLAME
FLAME received funding from the European Union's H2020 research and innovation programme under grant agreement No 731677

For example, this service can be activated by a relative of an old person if he/she wants to monitor the old person. In this case, video flows generated by all the IP cameras installed close to the places where the impaired person moves, are sent to the **Lying Person Recognition** element. Thanks to an artificial intelligence tool running as a virtual function on the FLAME platform as a virtual function, it can recognize people lying on the ground through an image recognition algorithm.

The VISION application framework deployed on the FLAME platform can be easily integrated with external devices, like the user-provided IP cameras and the mobile devices, which are connected to the FLAME platform by a WiFi link or through a VPN connection. In the experiment, in order to be compliant with the rest of the platform, we used some WiFi IP cameras transmitting HTTP-based flows natively, and some cameras whose flows are transcoded to this format with a software running on an attached Raspberry Pi board. Mobile devices, on the other hand, are Android smartphones or tablets where the VISION app is installed.

Key Results & Insights

Throughout the experimentation and the final trial, HTLab has identified KPIs and metrics to measure QoS and QoE parameters related to the media services developed within the app. On this basis, triggers have been defined to determine appropriate reactions in the FLAME platform which aim at maintaining the QoS and QoE target levels. KPIs have been decided for each functionality of the app, depending on their usage of the system and on different distributions of users on edge nodes.

Moreover, useful insights have been gathered regarding QoS metrics in various system configurations, giving an overview of system performance in different load conditions.

FLAME Insights

Benefits

FLAME provides the opportunity to deploy a service function chain in a real-world SDN/NFV-enabled infrastructure with MEC facilities. The capability of FLAME to deploy services at the edge of the infrastructure to have low-latency capabilities, to scale-out to the data centre, the excellent support and the various testing environments demonstrate the strength of the FLAME consortium.

Support

The VISION experiment received excellent support from the Consortium mentor and the whole FLAME technical staff, who helped the team solving several technical issues that arose along the way, especially during the deployment of the services.

Recommendations

The experience acquired throughout the development of the VISION platform suggests that it would be useful to support also shared resources (e.g., CPU, RAM) within the TOSCA definition by having the possibility to define services with shared resources: when a service is not exploiting the assigned resources at full capacity, other services could make use of the idle resources, as it is common practice in all modern virtualization environments.

Next Steps

The results we obtained from the experiment helped us identify the strengths and weaknesses of the VISION platform. To this end we plan on using the acquired knowledge to develop new functionalities, as well as disseminating our work through articles published on journal and conferences.

Funded as part of open call 3 by FLAME
FLAME received funding from the European Union's H2020 research and innovation programme under grant agreement No 731677

Summary & Objectives

For our 'MeshParade' project, Zubr set out to create an experience which allows anyone to be captured in 3D through a collaborative photogrammetry process, then left behind in an augmented reality 'sculpture garden' layer in the testbed environment.

The idea is that one person volunteers to be scanned, and must hold a pose while the other participants use the smartphone app to capture imagery of them; following instructions in augmented reality. The server continuously monitors the progress of the scan via the participants' devices, and automatically decides when the scan is complete.

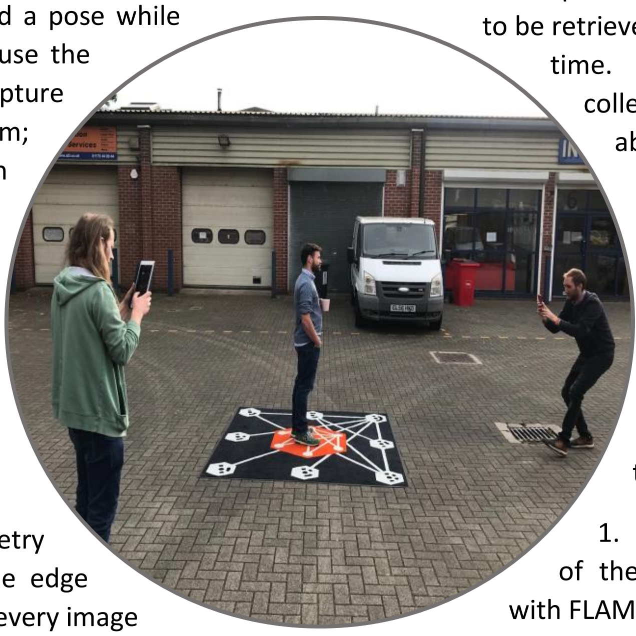

Our photogrammetry solution running on the edge receives and processes every image captured from each client device. A basic 3D model of the scan is churned out rapidly as a point cloud, and transmitted to all client devices, allowing participants to view the 'sculpture' as quickly as possible in AR. The engine then continues to process a higher-fidelity version of the 3D model, including textures which is transferred as soon as it is complete, replacing the previous one. The resulting user experience allows the participants to see the volunteer frozen in time, locked in their position in the augmented reality layer.

Each 3D scan is added to the testbed's cumulative 'sculpture garden' of AR 3D scans; to be retrieved by participants at any time. The more scans are collected, the further we are able to stress-test the latency of the system, as the models are not cached locally and must be downloaded afresh for each session.

For our service, we had the following objectives:

1. Validate compatibility of the MeshParade platform with FLAME
2. Verify if the service provides a good QoS - Are latency and download times reduced?
3. Verify if the service provides a good QoE - Are the models of a high quality?

Funded as part of open call 3 by FLAME
FLAME received funding from the European Union's H2020 research and innovation programme under grant agreement No 731677

MeshParade | Zubr

Jan 2020 – Sep 2020

Key Results & Insights

Our experiments and testing culminated in our final Trial run at the test bed in Millennium Square, on Weds 30th Sept. Six of the Zubr team were in attendance.

Overall, it was a successful trial. The resulting data allowed us to measure our KPIs, for example, including the distribution of the duration the "take photo" event – shown here across six different capture sessions.

This allowed us to establish that our KPIs relating to transfer speed were all within our desired range.

FLAME Insights

Benefits

We used 3 main aspects of the FLAME platform during this project: CLMC alerts; Edge nodes; TICK stack

Our primary use of the FLAME platform was using the CLMC to orchestrate the creation of proxy-storage servers at the edge to meet our KPI's. Using the alert specification we were able to reduce the data served by the DC by approximately 50%. This number could be larger as once the proxy storage cache thresholds are reached more data is served by the edge – our small trial gave only a glimpse at how this could be invaluable for a larger deployment.

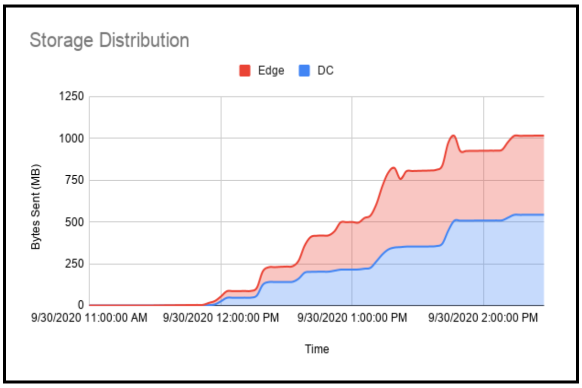

Support

The core Zubr team attended the launch briefing at the University of Bristol (5th Feb) to understand the FLAME platform and procedures. This provided the foundation for us to brief the wider team and proceed with the technical implementation for the project. Following this, due to COVID-19, all contact was remote via Gitlab, email support and video calls.

Recommendations

Early on in development we identified that the use of websockets would be very valuable. We were hoping to run the Lobby server (Websockets) and Photogrammetry server on the edge nodes but this is not possible in the FLAME architecture. Additionally, our application and computer vision apps in general benefit from having access to GPU's to perform repetitive tasks. If the FLAME platform supported GPUs we would be able to bring photogrammetry closer to real time performance.

Next Steps

We are excited by the solution we have developed to run on the FLAME architecture, and the results we gathered. We hope to continue building upon the project in future. For example, we would like to see a version of MeshParade that allows participants in different cities capture and share 3D scans together.

Funded as part of open call 3 by FLAME
FLAME received funding from the European Union's H2020 research and innovation programme under grant agreement No 731677

Open Call

4

FLAME received funding from the European Union's H2020 research and innovation programme under grant agreement No 731677

FLAME's fourth open call provided the opportunity for experimenters from any of the previous calls to return to FLAME and extend their previous work. The award was of €25k, with experiments to be based in Bristol, Barcelona or Buseto Palizzolo. Applicants were encouraged to consider Covid-19 contingency plans from the start.

D-Cube: IMRA 2

Mobile users are provided with enhanced views of marathon runners. AI systems provide athlete recognition from streaming video and video content is summarised and placed at the edge nodes where requests are expected.

Intellia ICT: HoloFLAME

Tourists visiting an agricultural museum are able to consume mixed- and augmented-reality content as well as upload and share their own content.

medVC.eu sp. z o.o.: medFLAME+

Synchronised video streams are provided to multiple mobile users to assist with remote medical education in the time of a pandemic.

FLAME received funding from the European Union's H2020 research and innovation programme under grant agreement No 731677

Summary & Objectives

IMRA 2.0 is an enhanced version of a marathon application for outdoor running events designed to be deployed in 5G infrastructures. During such events, it is impossible for the viewers to monitor the entire race including their athlete/s of interest. IMRA (Immersive Marathon Runner App) is a crowdsourcing application were users are able to stream video content from the race, from their point of view. At the same time, all the available content is being processed in the FLAME platform in order to provide back to the users personalized content regarding the athlete that intrigues them. This solution hopes to provide to the viewers of running events with a personalized, based on their interests, coverage of the event, while providing high-level and sophisticated content that derives from Deep Learning and Machine Learning methodologies.

The designed experiment's aim is:

- Performing module and deployment optimizations such as (i) adapting IMRA 2.0's modules in-real-time according to the number of requests, slightly compromising quality to maintain low latency; (ii) optimizing IMRA 2.0's data flow to serve multiple requests at the same time.
- Experimenting with different design patterns to improve the IMRA 2.0 design in terms of QoE and QoS (e.g. smart content placement based on predicted demand for reduction in start-up time for playout. etc.)
- Comparing FLAME's streaming performance against a cloud-based solution

The core outcomes of the IMRA experiment are already envisioned through the objectives. Specifically, the main objectives for the IMRA experiment are the following:

- O1: Develop IMRA 2.0 version that will be able to cope with multiple concurrent users
- O2: Experiment with different FLAME design patterns, compare to cloud-based solution and provide feedback
- O3: Dissemination

15 Jul — Flame-in-a-box

4 Sep — Test on Sandpit

8 Sep — Bristol Infra 5 UEs

Funded as part of open call 4 by FLAME
FLAME received funding from the European Union's H2020 research and innovation programme under grant agreement No 731677

IMRA 2.0 | D-Cube

Jun 2020 – Sep 2020

Key Results & Insights

D-Cube has experimented with the Geographical Scaling and the Content Placement design patterns. During the remote testing on Bristol platform, D-Cube deployed IMRA modules on all 5 towers, experimented on both scenarios and collected measurements that revealed the gain in QoE.

Geographical scaling

This feature ensures that all deployed resources are fully utilised, and that the user will be served by the closest endpoint. When the closest endpoint was activated during geographical scaling out, an initial playout delay of almost x2.5 times lower was achieved.

Content Placement based on prediction demand

IMRA 2.0's Smart Edge Storage module was notified by the Smart Summarization for newer versions of an athlete's summary and decided when to cache it locally based on previous requests. That way, in case of a future request, content is streamed from the closest node, resulting in lower initial playout time and higher bitrate.

FLAME benefits over Cloud

During experimentation, D-Cube also realized the following benefits of streaming from FLAME vs a cloud solution:

- 3.5 times lower latency
- 90% higher bitrate
- Higher user acceptance, as sensitive data always remain local

FLAME Insights

Benefits

D-Cube had access in a unique and real-world testbed to prepare IMRA for the market, upon the availability of 5G platforms. FLAME offered many important features with the most significant to be the adaptability of the platform and its ability to scale fast and efficiently to handle high loads.

Support

The technical support throughout the experiment was excellent, responding to queries on the documentation fast and efficiently. However, the most valuable support was provided by our mentor, Navid Solhjoo, who guided D-Cube in order to achieve the expected results. D-Cube extensively used all the available resources in Gitlab and OwnCloud to complete the experiments successfully.

Recommendations

D-Cube suggests the development of a graphical interface that could allow the initial configuration of resources and alerts at a higher level, to emphasize the ease of use of the platform and minimize the deployment time.

Next Steps

Experimenting with the deployment of cutting-edge AI services in a 5G platform with distributed resources at the edge allowed to gain insights for the challenges that D-Cube products would face and how to address them, upon 5G availability.

The derived know-how from this open call is of high value and will be used to enrich D-cube's flagship product "the Immersive Framework", in order to support 5G topologies for immersive applications.

Funded as part of open call 4 by FLAME
FLAME received funding from the European Union's H2020 research and innovation programme under grant agreement No 731677

HOLOFLAME | Intellia ICT

Jan 2020 – Oct 2020

Summary & Objectives

HOLOFLAME focused on conducting detailed experimentation and performance analysis of the media delivery technologies provided by the FLAME platform to support a wide range of Augmented Reality (AR) scenarios using commercial off-the-shelf (COTS) Mixed Reality (MR) hardware. The tests were executed over the Level7 infrastructure located at Buseto Palizzolo, Italy. Based on the interest expressed by the Buseto Palizzolo community, the set of selected scenarios were related to the field of Agritourism to support customised delivery of AR content.

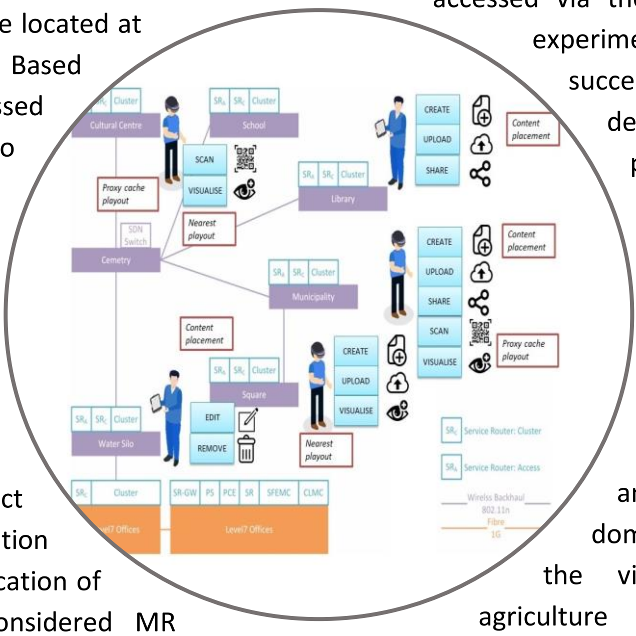

In particular, the project investigated the adaptation of an existing MR application of the company. The considered MR application retrieves and visualises AR/MR content about events including 3D maps, paths and routes, 360-degree videos, Points-of-Interest (POIs), etc. In the context of HOLOFLAME, the existing application has been extended with the capability of mobile crowdsensing by allowing users to enrich our database with new multimedia content that can be used by other participants and visualised in the MR application. By exploiting the unique service routing capabilities offered by FLAME, the users were able to build dynamic MR scenarios by uploading, customising and delivering AR content. The MR scenarios were decompiled in a series of service functions (for AR content, storage, execution, content delivery) that can be accessed via the network. During the experimentation period, we successfully validated 5 design patterns (nearest playout, proxy cache playout, content placement, scale geographically, application function offloading) in three different scenarios for the cultural heritage and the agricultural domain: (a) a use-case for the visitors of the local agriculture museum that have dynamic access to AR content for the artefacts; (b) a crowdsensing use-case that allows the users to upload and share their own content; (c) a Virtual Wall that offloads image processing functionality to the cloud. Tests that make use of the physical infrastructure in Sicily took place in June 2020 and September 2020.

25 May — Remote test

23 Jun — On-site test

29 Sep — Remote Trial, 3 people

Funded as part of open call 4 by FLAME
FLAME received funding from the European Union's H2020 research and innovation programme under grant agreement No 731677

HOLOFLAME | Intellia ICT

Jan 2020 – Oct 2020

Key Results & Insights

During the trials, a set of performance metrics were monitored to assess the quality of system performance. Such metrics included response times, download and processing times, storage capacities, network costs, battery consumption etc. The FLAME platform was successfully tested as a possible candidate to host AR applications. The ability to serve requests from the closest SF endpoint or dynamically cache content close to the demand was proved vital, as it reduces drastically the network cost and increases the quality of experience.

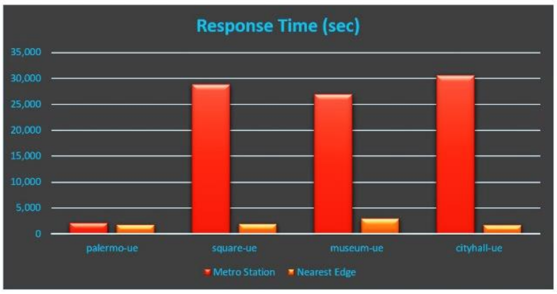

Figure 1: Downloading time for AR content from metro station and the nearest edge cluster

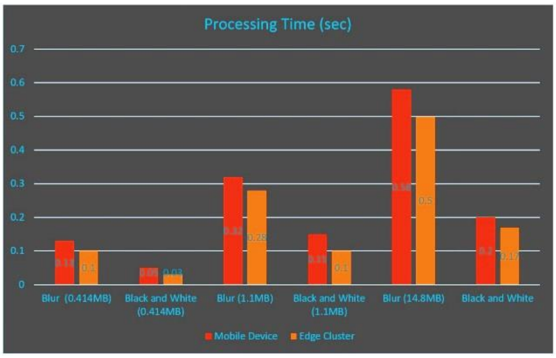

Figure 2: Response time for image processing

FLAME Insights

Benefits

A common issue that conventional commercial cloud-based infrastructures are suffering from is the fact that the user normally does not have real control on the physical machines. This issue becomes vital considering edge networking where a node usually serves the nearby users from a geographical point of view. In HOLOFLAME it is important to know *where* a service is executed to facilitate load balancing and that the overall workload is distributed fairly among the different nodes. Also, for an SME with limited resources, the reduction of development time is of utmost importance. FLAME provided an environment with unique opportunities for INTELLIA since it offers resources that are easily configurable through VM clusters that can be instantiated by the provider according to the specific needs of our experiments

Support

FLAME provided a full set of services to facilitate the HOLOFLAME experimental activities. Such services include the provision of tools for remote experimentation, documentation for the customization of the platform, continuous collaboration for service deployment and resolution of conflicts among experimenters, a set of paradigms and best practices to help experimenters get familiar with the platform and minimize the time to learn the platform. This last point is a major advantage for SMEs that want to design, implement, and deploy software solutions as fast as possible.

Recommendations

Availability of the platform was sometimes an issue, but this was somehow expected considering that HOLOFLAME 2.0 applications were hosted by an infrastructure that was an addition through the Open Call procedure.

Next Steps

Low-level access to physical resources and adaptation capabilities would not be feasible without FLAME. The company targets to pursue the development of a commercial Android application for touristic purposes based on the concept of crowdsensing and the use of dynamic AR content.

Funded as part of open call 4 by FLAME
FLAME received funding from the European Union's H2020 research and innovation programme under grant agreement No 731677

medFLAME+ | medVC.eu sp. z o.o.

Jan 2020 – Sep 2020

Summary & Objectives

medVC.eu is a company operating in the video-based telemedicine field. One of our products is the Interactive Medical University portal (https://edu.medvc.eu), which is a platform for improving the qualifications of doctors and medical students in various medical specialties. The subject of the experiment is an extension in the functionality for remote medical education, compared to the functionality we have tested in the open call 3 experiment. As the Covid-19 pandemic has disturbed the process of medical education in a severe manner, our services, which allow to teach medical students remotely, are of high need. The medical doctors, apart from using our real-time teaching tools, have requested a feature of synchronous live playout of recorded surgical sessions during a live lecture. To facilitate this, we have developed a module that enables synchronous playback of pre-recorded medical videos for the lecturer and the students. During the medFLAME+ experiment, we aimed to test various aspects of the QoS and QoE of this functionality.

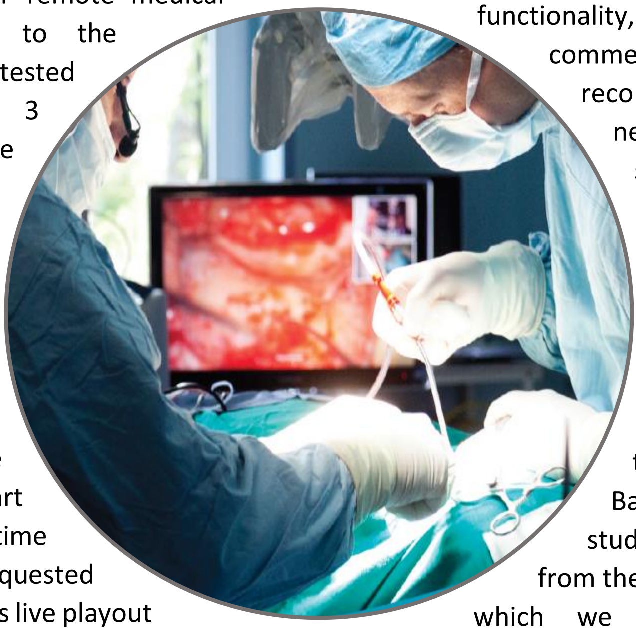

The idea of the medFLAME+ experiment was that during a live lecture conducted using web conferencing technology, the professor can choose a VoD that they want to show to the students. When the lecturer starts the playback, it begins synchronously on all the students' devices. The professor can pause and browse back and forth in the video, and the video players of the students follow those commands. If a student joins during a running session, they get synchronized. This is a crucial functionality, as the professor commenting live on the recorded surgical session needs to be sure that the students see exactly what the professor does, so they can focus their attention on the critical parts of the surgery.

The original plan was to perform trials in Barcelona with medical students and a professor from the Sant Pau Hospital, with which we collaborate regularly. Unfortunately, the pandemic situation did not allow that, so we decided to utilize on-site testing with the help of the infrastructure staff, complemented with methods for remote measurement of the QoS and QoE developed by us.

1 Sep — Remote test

16 Sep — On-site test

22 Sep — Trial 4 people

Funded as part of open call 4 by FLAME
FLAME received funding from the European Union's H2020 research and innovation programme under grant agreement No 731677

Key Results & Insights

The experiments where the participants of the lecture remained stationary resulted universally in very good marks - both streams (the one with the lecturer and the medical one) were received with at most minor issues. On the other hand, streams that involved mobility and handovers between the APs were very problematic, up to the point that sometimes it was not possible to take part in the lecture. This shows that while one could take part in the lecture while sitting in a cafe, it is impossible to do so while walking. Although the root cause of the problems does not lie within our solution - the problems are caused by severely delayed handovers between WiFi APs - in the end it is the QoE of our system that suffers. Therefore, it seems justified to have the system made more robust to circumvent exactly these problematic network conditions. It is obvious that it is the system that needs to be adjusted to the expected use cases and user scenarios, and not the other way around.

On the QoS level, the results show some sparing of DC bandwidth, however significant change occurred in a minority of the scenarios. Stream desynchronization in the scenarios where the streams were being viewed comfortably, was well within tolerable limits.

FLAME Insights

Benefits

FLAME's holistic approach to media experimentation saves time and lets the media service developer concentrate on the tests and research questions that need answering. The focus on edge-computing, specialized media processing modules, tools for monitoring the state of the platform, and possibility of small-scale pre-market deployment are the benefits making FLAME an extremely useful facility in the process of bringing media-based Internet technologies to the market.

Support

The dedicated and competent support was one of the strongest points of our experience with FLAME. Both from the platform and infrastructure side, it did not concern merely the technical aspects, but provided critical analysis of the experiment goals and paths of achieving them. Finally, it was only due to the help of infrastructure staff that we were able to perform the on-site testing despite the pandemic restrictions being in place.

Recommendations

Functionality-wise, the following issues could be improved from our point of view:

- the triggers and monitoring could have a temporal aspect, allowing to check if a condition has already been met for a specified time;
- request source detection;
- an issue that was specific for our scenario was the inability to use WebRTC in remote testing;
- another issue that turned up was the inability to use HTTPS in our scenario.

Next Steps

The feature of synchronous live playout of recorded surgical sessions during a live lecture is a needed addition to our product offering. The tests performed within the medFLAME+ experiment allowed us to evaluate the readiness of this functionality for production deployment. As it turns out, it works satisfactorily in stationary scenarios, but there is still some effort required to make the system work in mobile scenarios in an acceptable way.

These insights will allow us to plan the steps necessary to further develop this functionality.

Funded as part of open call 4 by FLAME
FLAME received funding from the European Union's H2020 research and innovation programme under grant agreement No 731677