



Second Demonstration and Field Trial



Deliverable D6.4

ICoSOLE identifier: ICoSOLE-D6.4-BBC-SecondFieldTrial-v09.doc

Deliverable number: D6.4

Main author of Deliverable: Dave Marston (BBC), Rik Bauwens (VRT), Nikki Peeters (VRT), Mike Matton (VRT)

Internal reviewer: Rik Bauwens (VRT)

Work package / task: WP6 / T 6.2

Document status: Final

Confidentiality: Public

Version	Date	Reason of change
1	2016-09-21	Document created
2	2016-09-28	Added sections for other trials
3	2016-10-03	Initial contribution from TaW
4	2016-10-05	Added iMinds content and more BBC stuff
5	2016-10-10	Additional VRT input
6	2016-10-13	Additional TaW input
7	2016-10-13	Ready for review
8	2016-11-02	Post-review version
9	2016-11-10	Final version

The work presented in this document was partially supported by the European Community under the 7th framework programme for R&D.

This document does not represent the opinion of the European Community, and the European Community is not responsible for any use that might be made of its content.

This document contains material, which is the copyright of certain ICoSOLE consortium parties, and may not be reproduced or copied without permission. All ICoSOLE consortium parties have agreed to full publication of this document. The commercial use of any information contained in this document may require a license from the proprietor of that information.

Neither the ICoSOLE consortium as a whole, nor a certain party of the ICoSOLE consortium warrant that the information contained in this document is capable of use, nor that use of the information is free from risk, and does not accept any liability for loss or damage suffered by any person using this information.

Table of Contents

List of Figures	iv
1 Executive Summary	1
2 Introduction.....	2
2.1 Purpose of this Document	2
2.2 Scope of this Document.....	2
2.3 Status of this Document.....	2
2.4 Related Documents	2
3 Field Trial Overviews	3
3.1 Committee of the Regions	3
3.2 Edinburgh Festival	4
3.3 Leffingeleuren	6
4 Committee of the Regions	7
4.1 Technical overview	7
4.2 User-generated content capture (Moments).....	7
4.3 Content selection (Trademark)	8
4.4 Playout (Wall)	10
5 Edinburgh Festival	11
5.1 Technical Overview	11
5.2 Professional Audio and Video Capture and Playout	12
5.3 User-Generated Content Capture, Analysis and Playout.....	14
6 Leffingeleuren.....	18
6.1 Preparation	18
6.2 Technical Overview	20
6.3 Professional Audio and Video Capture.....	20
6.4 User-Generated Content Capture, Analysis and Playout.....	22
7 Other Trials	25
7.1 Austrian Science Night 2016	25
7.2 MNM Start 2 DJ	25
8 Conclusions	27
9 References	28
10 Glossary	29

List of Figures

Figure 1: Interactive Wall at the ICoSOLE booth	3
Figure 2: BBC Blue Tent stage	4
Figure 3: BBC R&D Portakabin	5
Figure 4: Royal Mile	5
Figure 5: Wall of Moments system diagram	7
Figure 6: Overview of sent-in Moments	8
Figure 7: Moment feedback	8
Figure 8: Moment location metadata	8
Figure 9: Moments queue	9
Figure 10: Assess Moment	9
Figure 11: Technical setup at CoR Open Days	9
Figure 12: Interactive version of the Wall	10
Figure 13: Overall EdFest system diagram	11
Figure 14: One of the audience microphone rigs	12
Figure 15: 4k Camera	12
Figure 16: Payout webpage	14
Figure 17: Screenshot of the Wall of Moments from Edinburgh Festival.	15
Figure 18: Map of Moments from Edinburgh Festival, showing live and non-live UGC.	17
Figure 19: Leffingeleuren floor plan	19
Figure 20: Wall of Moments booth	19
Figure 21: Wall of Comments	19
Figure 22: Professional capture production chain	20
Figure 23: Live editing interface with three camera inputs	21
Figure 9: Live editing interface being demonstrated at IBC	22
Figure 24: Starring posts in Focus	23
Figure 12: The Wall of Moments at the Austrian Science Night.	25
Figure 26: Technical setup	26
Figure 27: Screenshot of modified Wall ('Screen')	26
Figure 26: The MNM Start 2 DJ event	26

1 Executive Summary

This deliverable covers the field trials and demonstrations that occurred in the final period of the project. While the original aim was to have a single all-encompassing final field trial, due to factors beyond the project partners' control this plan had to be revised into three smaller trials. While all the partners could not participate in all the trials, each trial had at least some of the partners working together.

The three main field trials (in chronological order) were the Committee of Regions, the Edinburgh Festival and the Leffingeuren Music Festival. Each trial provided different scenarios and opportunities for testing different components of the project.

The Committee of Regions focussed on a revised vision on content contribution by event participants, as well as the introduction of a content selection process for the editorial team. For this purpose, we built two new apps from the ground up: Moments (contribution) and Trademark (selection).

The Edinburgh Festival trial focussed on professional object-based audio capture and interactive playout and UGC capture and content. The BBC had a venue where they hosted shows that were broadcast on the radio and online. The project took advantage of this broadcast infrastructure to get audio feeds, and of course the shows for capture material. Much of the UGC capture was done around the city where there were plenty of street performances going on. There was a high-speed network between Edinburgh and Salford over which the content was DASH streamed for a live browser playout.

The Leffingeuren Music Festival trial focussed on more advanced content selection and enhancing the event experience, especially for people present at the event. More thorough analysis was done on incoming content, more content sources were added (Twitter and Instagram) and the Wall was personalized using beacons.

2 Introduction

2.1 Purpose of this Document

This document describes the field trials that have occurred within the last 12 months of the project.

2.2 Scope of this Document

This deliverable covers the field trials at the Edinburgh Festival, Leffingelueren and the Committee of Regions.

2.3 Status of this Document

Final version.

2.4 Related Documents

D2.4.2 "Evaluations v2" describes some of the results of tests that were performed at the field trials.

D4.2 "Algorithms for content analysis and for automatic content pre-selection"

D4.5 "Tools for semantic scene creation"

D4.6 "Audiovisual rendering engine and interface for linear broadcast production"

D4.7 "Audiovisual playout-engine for interactive playout"

3 Field Trial Overviews

Three major field trials took place in the last 12-month period of the project. The original plan was to have a single large field trial at the Wimbledon tennis championships, but due to the hosts of the initially planned field trial pulling out at the last minute, an alternative had to be quickly planned. However, there was not a single event when all the partners could attend at the same time due to other commitments, hence the division into separate trials. The three field trials were at the Committee of Regions, the Edinburgh Festival in Scotland (<http://www.eif.co.uk/>) and the Leffingeleuren music festival in Belgium (<http://leffingeleurenfestival.be>).

3.1 Committee of the Regions

The Open Days, organised by the European Regions and Cities, was held on October 12-15 2015. The event mainly focuses on European political events in the European Regions and is set in several venues around Brussels. This was an opportunity for ICoSOLE because of this spread-out nature of the event, as well as the different approach a political journalistic event brings, as opposed to a festival. Both VRT and JRS were present, and tested out a second iteration of the Wall of Moments incorporating feedback from previous demonstrations. This time, the focus was more on the editorial side: assessment of incoming content and publishing it to the Walls, which placed across the event. A small ICoSOLE booth was installed at the main event building, where visitors were able to try out the interactive Wall as well (Figure 1).

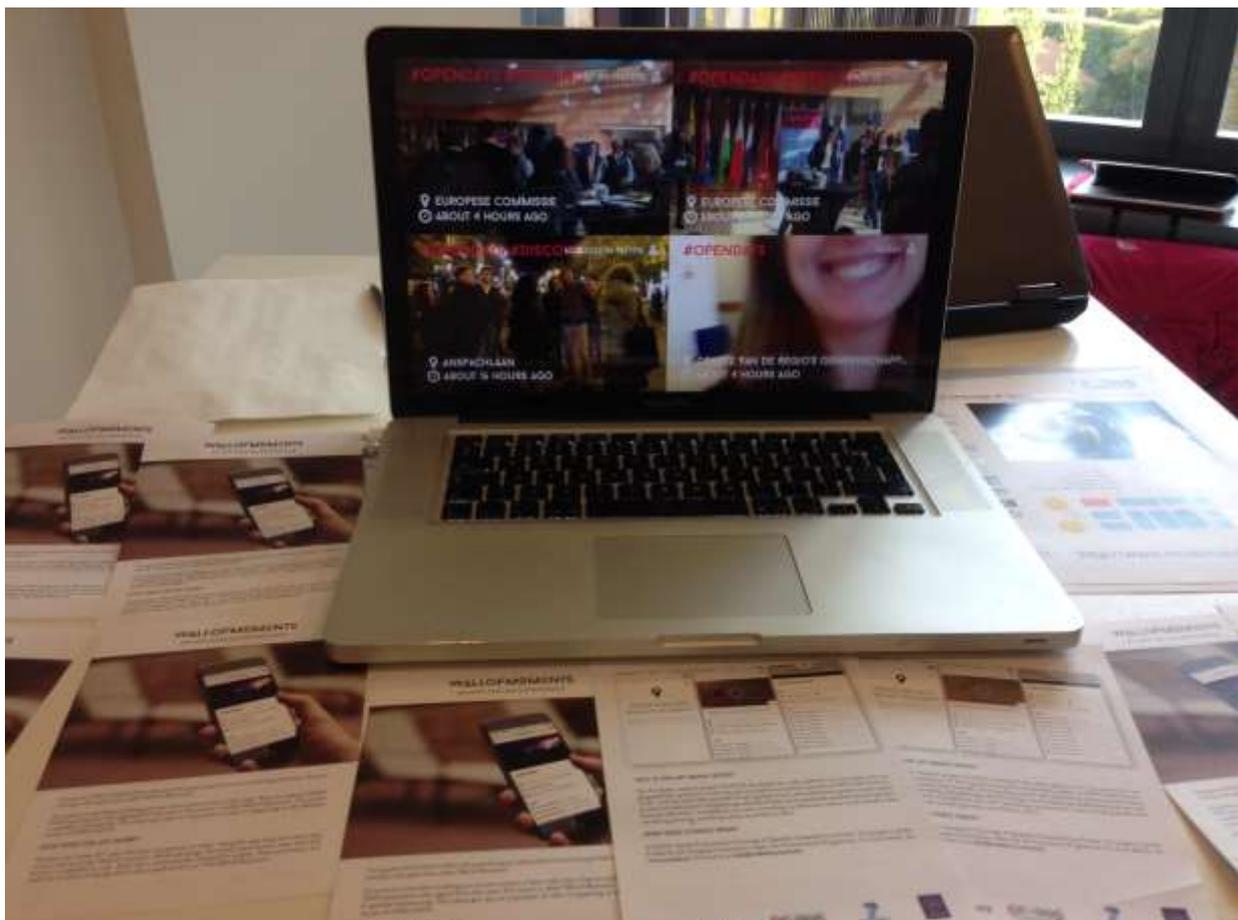


Figure 1: Interactive Wall at the ICoSOLE booth

3.2 Edinburgh Festival

The Edinburgh Festival takes over the Scottish capital for most of August (a month in which the occasional day of sunshine breaks the unremitting cloud and rain). Hundreds of venues, from large theatres to tiny rooms above pubs, as well as in public streets, host a wide variety of performers from comedians to musicians, ranging from famous household names to unknown artists trying to make the big break.

ICoSOLE's presence at the festival took advantage of the BBC's own venue as well as the multitude of street performers scattered all over the city centre (though particularly concentrated on the Royal Mile).

The BBC venue, based in the George Heriot school, called the Blue Tent, had a stage (see photo in Figure 2) with seating for an audience of around 150 people. Each day two or three shows were hosted there, many of them existing BBC Radio shows. These shows were also being broadcast on BBC Radio, so there was a full production setup there from BBC Scotland.



Figure 2: BBC Blue Tent stage

Several BBC R&D projects were doing trials at the Blue Tent (including the ICoSOLE field trial of course), who shared a portakabin (Figure 3) amongst the others in the production area at the side of the tent.



Figure 3: BBC R&D Portakabin

Outside of the BBC venue, user-generated content was captured out and about in the streets, particularly the Royal Mile (Figure 4) where all sorts of different street artists were performing.



Figure 4: Royal Mile

Team members from the BBC, JRS and BIT attended the festival between the 9th and 12th August, with the BBC concentrating their efforts in the BBC venue, and JRS and BIT roaming the streets capturing performances, while using one of the hotels as base and meeting point.

A team member from iMinds participated remotely in the field trial with the aim of validating the technical implementation of the Map of Moments playout client, while at the same time providing off-site testing support for the (live) UGC capture technology piloted at the festival by JRS and BIT.

3.3 Leffingeleuren

The Leffingeleuren music festival was on between the 9th and 11th September in the village of Leffinge in western Belgium near Ostend.

The presence of ICoSOLE at the Leffingeleuren festival was quite prominent, including the Wall of Moments booth and another installation of the Wall in a bar area. Beyond the test users, many visitors saw the Wall and were informed about the project and the technology being developed. Next to the Wall of Moments booth, the projector manufacturer Barco performed a trial of 360° video projection, with was fed by a camera installed by iMinds and sharing the network installed for the ICoSOLE trial. This increased the visibility of ICoSOLE at the event.

4 Committee of the Regions

4.1 Technical overview

During this field trial, a lot of emphasis was put on content selection and corresponding content analysis. All metadata was stored in Firebase¹, and data was uploaded to and accessed from Amazon S3². Both services are scalable out-of-the-box and offer interfaces in several programming languages (mostly JavaScript and Java were used). This made the overall implementation much simpler and gave some headroom to focus on the project's objectives, such as interactivity, analysis and layout.

Firestore was considered the central hub, to which all apps and services connect (Figure 5). The service also offers authorization and authentication, keeping user management simple to handle. When a user adds a new Moment using the Moments app, a new record is created in the database. After upload is complete, the URI to the file is also added. This triggers the analysis service (by JRS) to start processing the uploaded file. When analysis is complete, the result is pushed to the database, and because of the real-time nature of Firestore, the result is also rendered directly in the Moments app. The Trademark app shows a real-time queue of incoming Moments, which can be either approved or rejected by the editorial team. When approved, the Wall, which also connects to Firestore, picks up the Moment and adds it to the playlist queue.



Figure 5: Wall of Moments system diagram

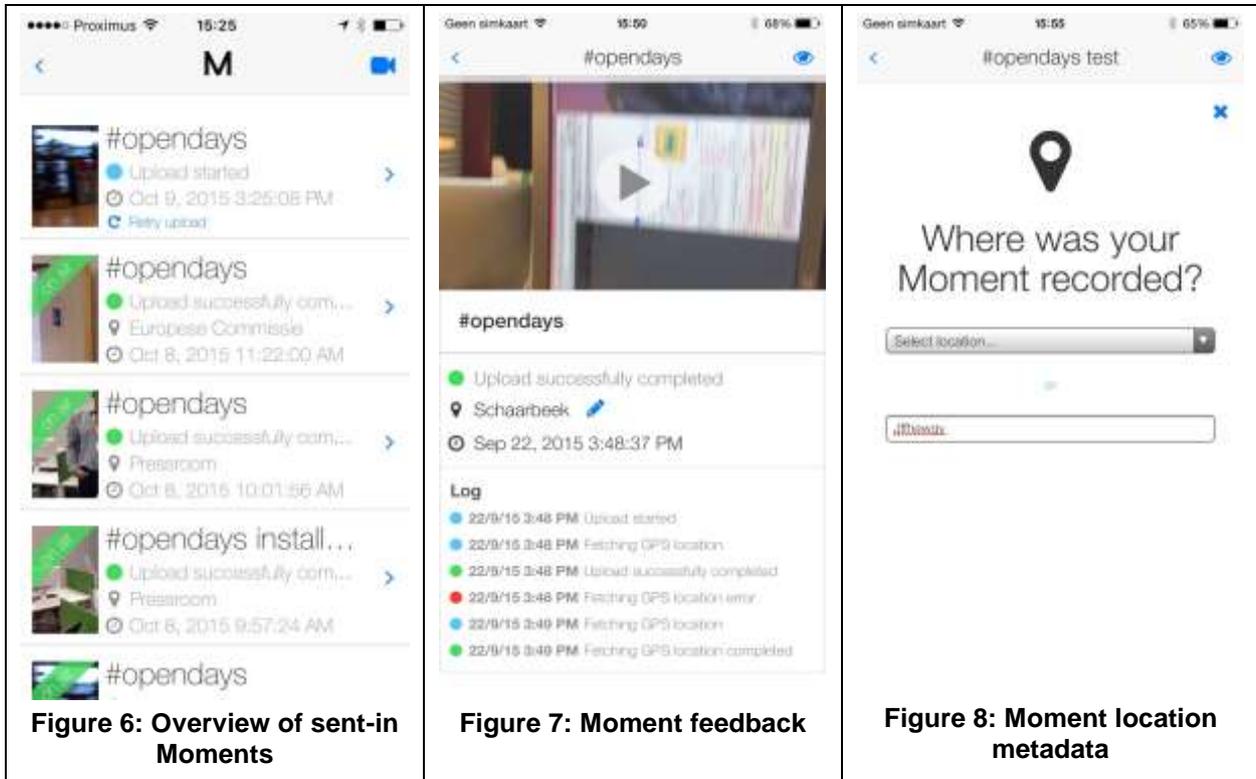
4.2 User-generated content capture (Moments)

Content was collected using a dedicated app: Moments. Available in the App Store (iOS) as well as the Play Store (Android), it was developed with a lean methodology (using web technology) with testing

¹ Firebase provides a realtime database and backend as a service. The service provides application developers an API that allows application data to be synchronized across clients and stored on Firebase's cloud. <https://firebase.google.com>

² Amazon S3 (Simple Storage Service) is an [online file storage web service](https://s3.amazonaws.com) offered by [Amazon Web Services](https://s3.amazonaws.com). <https://s3.amazonaws.com>

features iteratively. Users got a list of sent-in videos (Figure 6) and were able to upload new videos and received feedback on the status of their videos (Figure 7 and Figure 8, GPS/upload status, analysis results, Wall impressions). Analysis results yielded tips for capturing better content as well (e.g. try to find a better lit area).



Once the content is available on Firebase and Amazon S3, several backend processes are automatically started. The video analysis kicks into action, as well as a FFMPEG service, which converts the video into a web-safe format (x264 in MP4 container).

4.3 Content selection (Trademark)

When the FFMPEG service has converted the video file successfully, the Moment is added to the queue for review (Figure 9) [D4.5]. The ranking of the queue was quite basic, as it was the first version of the Trademark prototype: Moments were sorted ascending by time. It was possible, however, to filter content by video quality using a slider at the top (based on the analysis results). For example: show all unassessed Moments with an average quality above 80%.

A Moment could be either approved or rejected (Figure 10). This decision could be altered at a later time and is communicated to the user (along with the timestamp).



Figure 9: Moments queue



Figure 10: Assess Moment

As most of the services were cloud instances, the technical setup on-site was rather limited (Figure 10). The Trademark app was also deployed as a mobile app (using Cordova), which made it possible to do the assessment remotely.

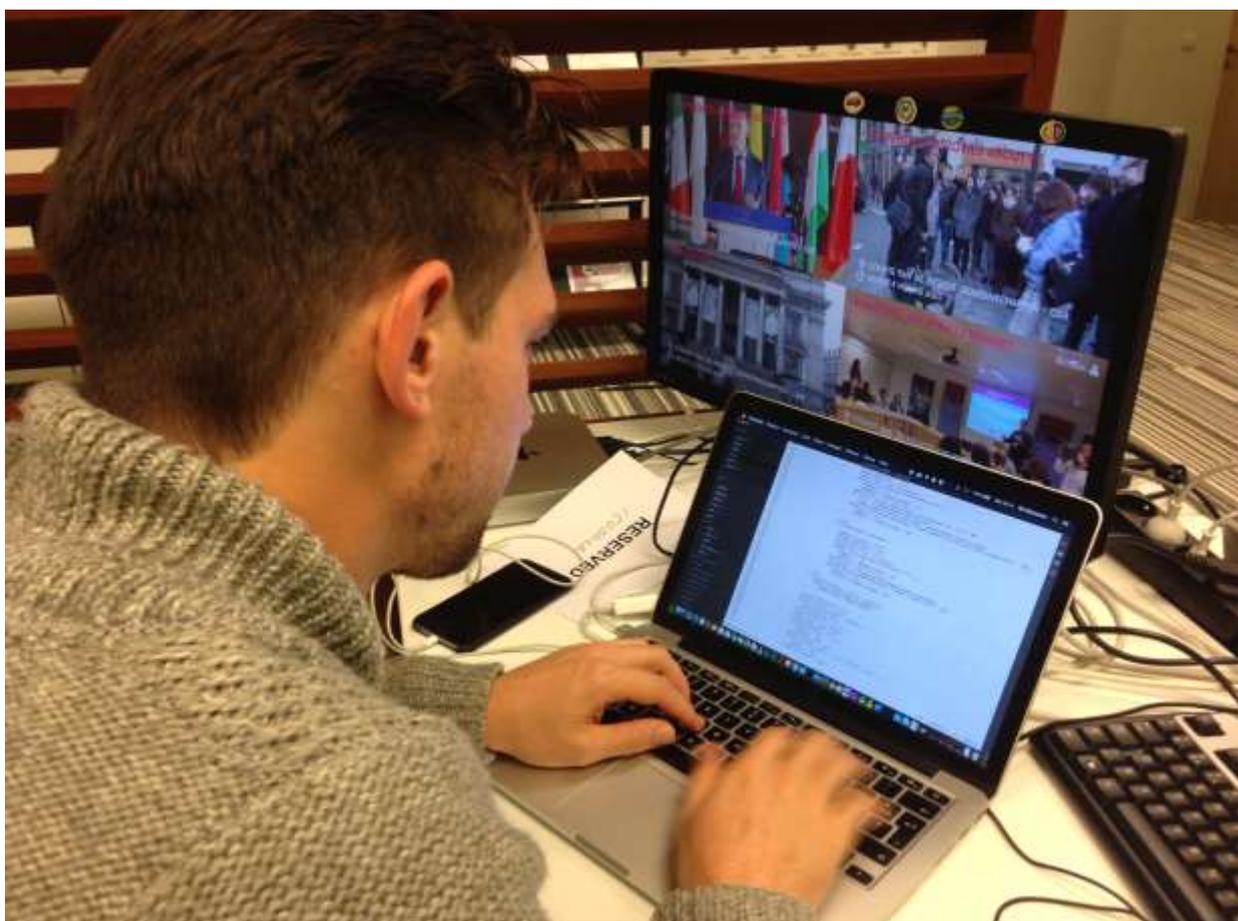


Figure 11: Technical setup at CoR Open Days

4.4 Playout (Wall)

Lastly, Moments were played out on screens installed on the event site. We developed the 'Wall' app, which connects to the same Firebase backend. The most recent approved Moments, along with a random Moment now and then are selected and rendered in one of four tiles of the Wall. This gives a vibrant overview over what's going on at the event at any time. Along with the videos, also location, user and time information is displayed. People at home were able to launch the app via a link we provided. Because four videos were streamed simultaneously, a decent computer (or a powerful Android tablet) was necessary to use the app. Also, we achieved best results using a Chrome browser (due to the HTML5 video implementation). This version was also interactive: every tile could be clicked to play full screen (Figure 12).

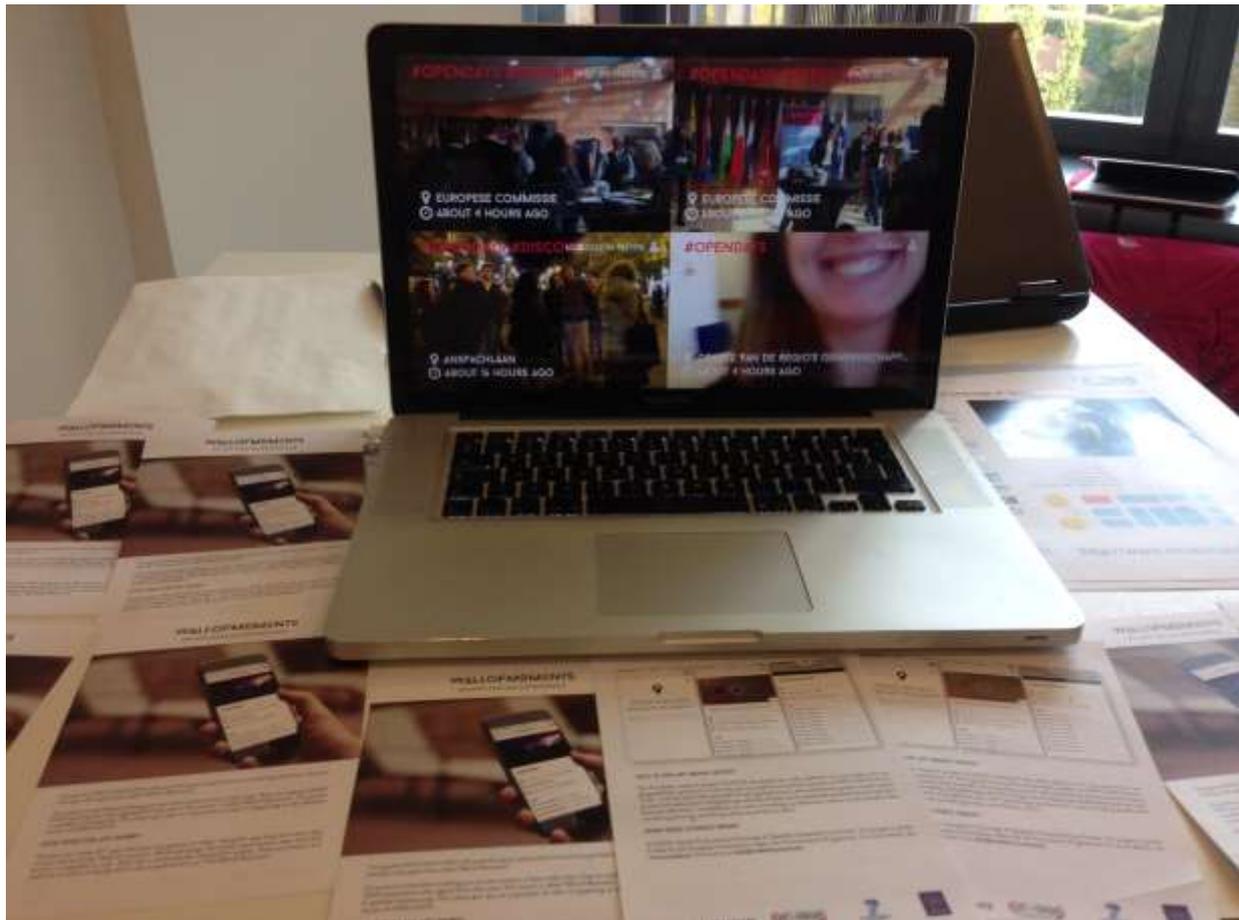


Figure 12: Interactive version of the Wall

5 Edinburgh Festival

5.1 Technical Overview

The field trial was split into two main parts: professional audio-visual capture at the BBC venue (i.e. fixed location), and user-generated capture material around the city (i.e. variety of locations), as at the BBC venue itself. Both these methods of capture were intended to be played-out over the same browser-based playout system.

To test the playout capabilities remotely, so mimicking the experience of the home user, the browser playout was tested in the BBC MCUK labs in Salford. A high-speed JANET link was connected between the BBC venue in Edinburgh and Salford, and this provided the connection for the professional content distribution. The UGC content made use of the Amazon EC2 cloud service, which was also received in MCUK. For the professional capture, distribution, storage and MPEG-DASH emission coding, the IP Studio system was used. The UGC captured live video also fed into the IP Studio system for storage and DASH playout.

The diagram in Figure 13 shows the general layout with the BBC Blue Tent containing the IP Studio capture machines for audio (z420-1) and video (z420-2). Primer was another BBC R&D project taking place, which had a video RTP that was used. The DASH encoding was done on these two machines and sent over the JANET link to MCUK. In MCUK the DASH was playout out via a browser. UGC was captured via mobile devices, which users were roaming around the city with filming interesting artists. Their videos were streamed up to the Amazon cloud where the video was analysed and processed before being streamed via RTP over the normal internet. This stream was received in MCUK and played out into an IP Studio capture machine (z220-2) for DASH encoding and storage. The UGC could then be played out on a browser alongside the professional capture material.

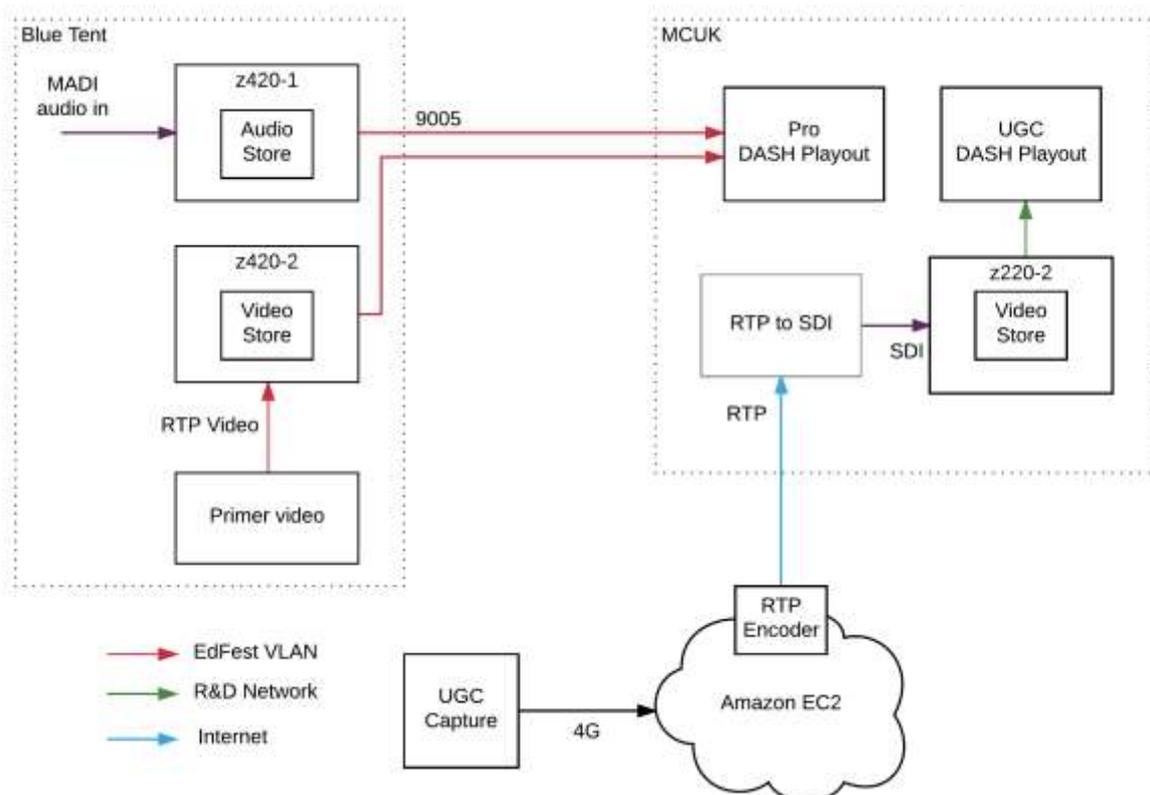


Figure 13: Overall EdFest system diagram

5.2 Professional Audio and Video Capture and Playback

5.2.1 Audio and Video Capture

The BBC's focus was on object-based audio capture and playback, but also combined standard video capture to ensure that aspect was covered in the trial. The BBC Blue Tent venue contained a stage with audience seating within the tent itself. The BBC was broadcasting the shows in the tent, so there were already microphones (Figure 14 shows one of the mic rigs to capture audience sounds) and video cameras installed for this purpose as well as all the production facilities that entails.

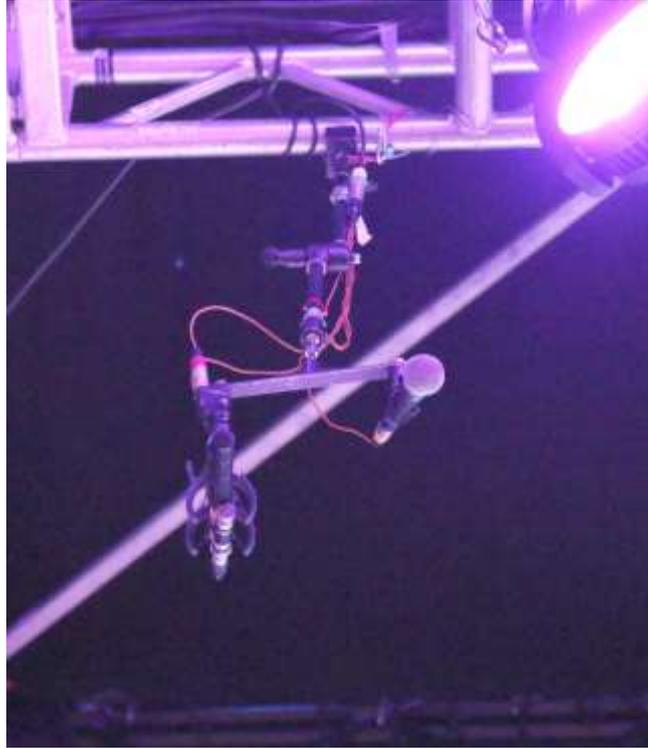


Figure 14: One of the audience microphone rigs

As well as ICoSOLE, BBC R&D also had some other projects occurring at the venue, including one called Primer (now called Near-live Production) which involved using 4k cameras for nearly-live video capture and editing. So for ICoSOLE, instead of using our own cameras, the Primer project provided downscaled (to HD) feed from one of their cameras (Figure 15). This was done via RTP from one of their IP Studio machines to our video capture IP Studio machine.



Figure 15: 4k Camera

For the audio capture, the microphones in the venue were connected to the outside broadcast (OB) van next to our portakabin. This van contained all the sound mixing and broadcast distribution equipment, as well as the sound engineers who mixed the live sound. From their mixing desk they provided a MADI connection (via coax) which contained individual feeds from the microphones (though some were mixed down for convenience) to allow us to generate object-based audio. The MADI was taken directly into an IP Studio machine, which could then record it in the sequence store and DASH-encode it.

5.2.2 Object-based Audio Production

The MADI feed contained several active channels (the number of which varied on the setup for each show), some of which were from stage mics, some from audience mics, and some from other feeds (sound effects, music and clips). A channel list was provided before each show, so it was known which channels were which. The channels were also post-fader, so their levels were generally correct with some equalisation applied for a good sound.

To turn these individual channels into audio objects metadata needed to be generated to attach to them. To generate this metadata a browser-based audio editor was developed. This gave simple controls for each channel to allow the azimuth, elevation and gain to be set; thus allowing an immersive sound scene to be generated. Knowledge of the rough locations of the microphones allowed positions to be set that could produce a realistic sound stage. This editor also included a binaural renderer to allow the user to listen to an immersive mix of the sound [D4.5].

The interactive and personalised aspect of the audio was to allow the consumer to adjust the levels of the stage and crowd sounds. This was achieved by grouping together the audio channels, and this grouping was carried in the metadata. The editor allowed for this grouping to be easily done by assigning names for each channel (so channels with the same name would be grouped together).

5.2.3 Video Production

As the camera used was a single locked-off 4k camera pointing at the whole stage, no further production was required for the video feed. The Primer project's IP Studio machine for that camera performed downscaling from the 4k to standard HD, and sent this H264 encoded video via RTP (the standard IP Studio method of streaming) to the ICoSOLE video capture machine.

5.2.4 Sequence Stores

Both the audio and video IP Studio capture machines contained sequence stores, which stored the incoming audio and video segments as they came in. The sequence stores the audio or video segments in such a way that they can be played out by a DASH player. A manifest file is also automatically generated for this purpose. These files are accessible via a RESTful interface, via port 9005 (one of IP Studio's specifically allocated ports).

5.2.5 DASH Streaming

The DASH encoded segments in the sequence stores (one each for audio and video), were accessible on the EdFest VLAN that was available over the JANET link to MCoUK. As the audio and video streams were from separate stores, so their DASH streams were separate, the DASH player needed to be able to synchronise both streams.

5.2.6 Audio and Video Playout

The playout was done via a browser (Chrome was used, and Firefox also appeared to work), so javascript code was written for the DASH playback for both audio and video [D4.7]. A web page was designed to present the video along with some interactive audio controls as shown in Figure 16.

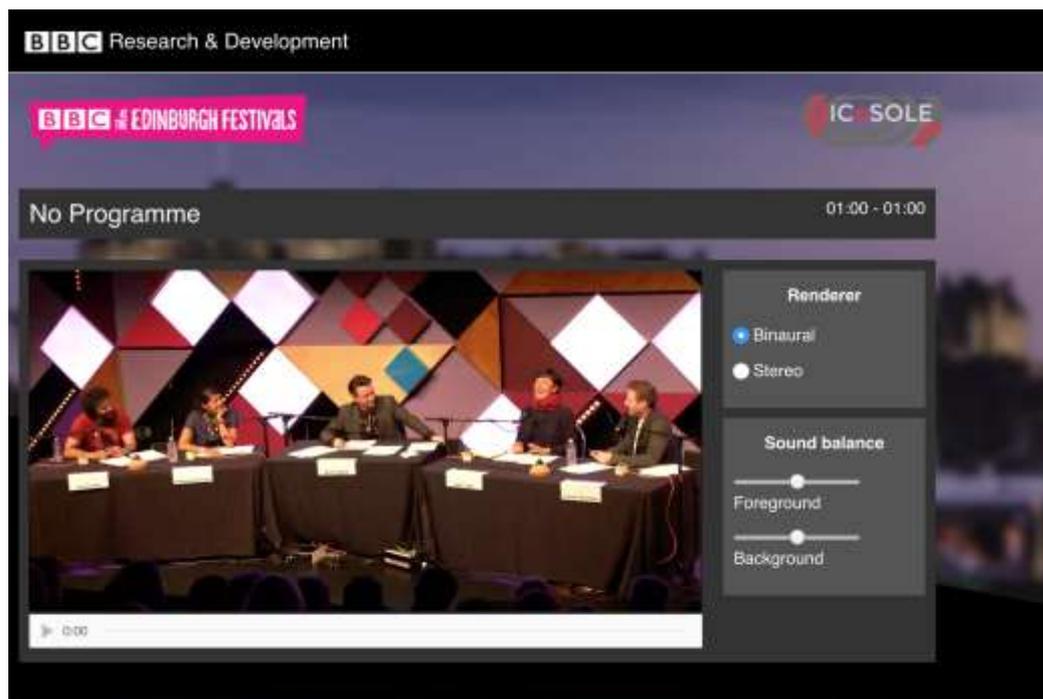


Figure 16: Playback webpage

The player included buttons to allow the user to switch between different types of renderer (binaural for immersive sound over headphones, or stereo for normal sound over speakers), and sliders to adjust the levels of different aspects of the sound (such as the stage and audience). The video playback has the conventional play/pause button and timeline.

Currently a single audio stream with greater than six channels cannot be decoded and played directly in the browser; but must be broken down into multiple streams of six channels or less. The Web Audio API allows sample-perfect scheduling (and consequently sample-perfect synchronisation) of these audio streams. However, there is no such method of synchronising playback of audio with video (as mentioned in Section 5.2.5).

A simple method of synchronisation was constructed: the current playback head of both the video and audio is polled, if the audio playback has drifted ahead, the video playback rate is slowed. Conversely, if the audio playback has drifted behind, the video playback rate is sped up. The playback rate adjustment is re-evaluated around twenty times per second in order to maintain synchronisation.

5.3 User-Generated Content Capture, Analysis and Playback

5.3.1 Moments Non-live Video Capture

Images and videos for the Wall of Moments have been captured by members of the ICoSOLE and Cognitus teams using an improved version of the capture app. This use served as a beta test before distributing the new version to test users at the Leffingeleuren trial. Only around 40 clips were collected, as the focus was on testing the new UGC live capture functionality.

The Wall of Moments has been made available on the web and as an app for mobile devices. A screenshot of the Wall with content from Edinburgh Festival is shown in Figure 17. Both the offline clips and the live streams have also been provided to the Map of Moments, which was first tested at this trial.



Figure 17: Screenshot of the Wall of Moments from Edinburgh Festival.

5.3.2 Moments Live Video Capture

The new live streaming processing and playout chain was tested at Edinburgh Festival. In contrast to previous trials, where the live streams were captured and processed locally, the capture and processing of live streams was performed entirely in the cloud. For the mobile phones used in the trials, prepaid 4G SIM cards were used. The test thus also brought insights into the feasibility of using the ICoSOLE live streaming app with a regular mobile data infrastructure. The network performance was overall satisfactory, except from few spots not well covered by the chosen operator, and some performances where several people streamed from their mobile phones nearby, thus impacting the available bandwidth. More than 100 UGC clips were recorded and sent through the system.

The automatic content analysis [D4.2] was performed in real-time on a machine in the cloud. The system is designed to assign the capture and processing machines dynamically. Thus additional machines can be started in the cloud, and will be assigned when new streams arrive. This initialisation step between the backend system and the mobile phone also includes the synchronisation of the capture app with the backend system, so that the streams are correctly time-stamped.

The content was either restreamed to BIT's DASH encoding service in order to obtain a DASH version of the content to be displayed on the Map of Moments or forwarded to BBC's IP studio infrastructure (see above). After the live stream ended, the highest quality level of the DASH segments was used to assemble a VOD version of the file. It turned out that some versions of Google Chrome have issues playing these concatenated files.

5.3.3 Live Streaming and Playout

The live stream was provided either as DASH stream (to be consumed via the Map of Moments or a standalone DASH player), or restreamed to IP Studio. In this case, NMOS compliant headers with identifiers and timestamps are inserted.

IP Studio captured the live stream of the UGC content after it had being through the video analysis processing in the cloud. Due to the restrictions on bridging different networks, a rather convoluted chain had to be implemented to perform this capture as shown in Figure 13, where the RTP stream from the cloud (on the normal internet) was played out using a RTP video player on a PC with an SDI output card. This SDI feed was connected to an IP Studio machine that was connected to the BBC R&D's network. This SDI connection prevented bridging of the normal internet with the R&D network. Not an ideal solution, but it at least got the live UGC video into the IP Studio machine situated in MCUK.

The IP Studio machine (z220-2) was configured to record the incoming video into a sequence store in the same way as was done for the professional live capture. By using the same method as the professional capture it meant the DASH playout for both the UGC and professional capture could be the same.

As with the professional playout, the DASH-ed UGC live content was played through a browser with a javascript DASH video player. This version did not require any interactive controls as it only contained video.

5.3.4 Map of Moments

An iMinds WP5 team member participated remotely in the Edinburgh Festival field trial. The goal of this remote participation was two-fold:

- Test and technically validate a preliminary version of the Map of Moments (MoM) demonstrator³.
- Provide feedback on the live UGC capture and streaming technology piloted on-site by JRS and BIT (see Section 5.3.2).

The MoM application had been tested and validated previously using existing UGC corpuses collected as part of earlier ICoSOLE field trials (e.g., the Austrian Science Night). The “issue” with such pre-existing content samples is that all individual content items are already available and no new UGC is ever added to them. As such, the Edinburgh Festival field trial represented a valid opportunity to test the MoM application’s reaction to UGC that was produced (and hence became available to the application) dynamically and incrementally as the real-world event progressed. It turned out that the initial implementation of the MoM application was largely able to cope with incrementally published UGC items; just a few smaller implementation issues had to be looked into and were subsequently amended.

Similarly, the Edinburgh Festival field trial marked the first opportunity to test the MoM application’s support for the streaming and playback of live UGC captured under realistic circumstances. As was the case with the incremental publishing of UGC contributions, the preliminary MoM implementation was found to be already quite robust with respect to the live streaming and playback of UGC; again only some minor implementation-related tweaks were required.

So, in summary, iMinds’ off-site involvement in the Edinburgh Festival field trial allowed the technical implementation of the MoM demonstrator to be polished with respect to supporting realistic, dynamic UGC posting patterns as well as live UGC streaming scenarios. The resulting revised implementation of the MoM demonstrator was then later demonstrated at the Leffingeleuren field trial (see Section 0).

Other than to improve the MoM implementation itself, the iMinds participation in the Edinburgh Festival field trial turned out to be valuable to on-site consortium partners JRS and BIT as well. In particular, the MoM experimentation involving live UGC footage from the Edinburgh Festival revealed a number of technical issues with the live UGC capturing technology piloted on-site by JRS and BIT. These issues were reported to the involved consortium partners so that the live UGC streaming implementation could be improved.

³ Please see deliverable [D5.4] for a description of the functional design and implementation of the Map of Moments application.

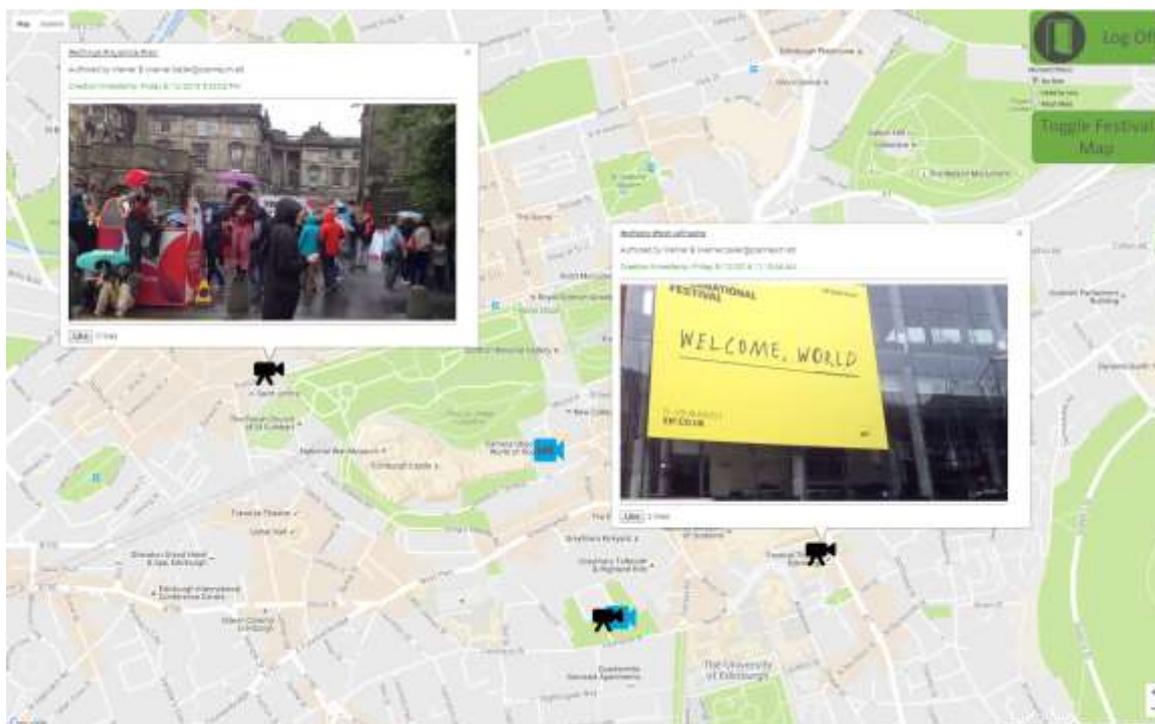


Figure 18: Map of Moments from Edinburgh Festival, showing live and non-live UGC.

6 Leffingeleuren

6.1 Preparation

Leffingeleuren is a family event. Mainly, inhabitants living around the village of Leffinge and volunteers organize the event in a professional manner. The event has existed for over 20 years and has gone through some transitions a couple of years ago. One of the results was that the event became smaller and only two of the locations require an entrance ticket. Because of this, the festival thrives on its friendly and easy-going atmosphere.

For the field trial, the following preparations had to be made:

1. A contract with the festival organizer: vzw De Zwerver
2. Separate contracts with each of the participating bands of the main stage
3. Gather around 40 test users registered in advance, and be prepared to welcome additional test users at the site
4. Construct a Wall of Moments booth in order to be publicly visible at the event location

6.1.1 *Contact with the festival organizer*

In the contract with the festival organizer, the following was agreed on:

- Permission to carry out field trials on vzw De Zwerver's festival grounds, with setup and teardown
- Permission to install cameras, microphones and visualization systems on the festival grounds, to use (and record) video and audio feeds provided by vzw De Zwerver, and to install cables for audio, video, power and WIFI-infrastructure
- The provision of a technical room by vzw De Zwerver, for the monitoring of test recordings and storage of equipment
- Room on the festival site to place an ICoSOLE booth
- Permission to place screens on several strategic locations across the festival terrain

6.1.2 *Contract with the bands*

For this field trial, receiving permission from the bands went well. As long as we focused on the fact that the recordings were only used for research purposes, bands agreed to sign the agreement. The first band even called us in advance, asking about the details of the project and whether they were allowed to already mention the project to their fans. Only one band wanted to make considerable changes to the contract, restricting audio rights and full entitlement to any footage. Having the permission of the bands playing in the afternoon enabled us to even let the cameras for the live stream be controlled by our partners at IBC in Amsterdam.

6.1.3 *Test users*

Similar to the field trial in Dranouter, finding test users proved to be challenging. After sending out our registration link, we first received a response from some media enthusiasts who already had a ticket, truly excited about testing the app. Then, we focused on targeting potential test users by emphasizing the exclusivity of the app, as part of a European research project. When we still did not reach our goal of 40 test users, we sent out more news about the free tickets.

We received about 60 registrations via the platform. However, after e-mailing these respondents with further instructions, several did not answer back. This consistent loss of respondents, even after notifying them about winning a free ticket, proved to be a challenging process. The fact that the festival only has 2 admittance rooms and is different than other festivals, addressing children and family atmosphere, could explain this. In the end, we were able to reach our 40 test users. Moreover, we even had about 60 test users more, who spontaneously downloaded the app at our booth.

During our communication campaign, we sent out 2 press releases and posted consistent updates on social media platforms. Several (online) newspapers and media websites picked up our press release.

There was an interview on a regional radio news bulletin and Studio Brussel also promoted the Wall of Moments by posting a contest on their website.

Via the registration platform, we could gather additional information (such as age, occupation, music interest). We handed out weekend tickets based on the motivations of the respondents. On social media, we posted the link to the Wall of Moments website to watch from home. Most feedback was gathered during the festival, when test users spontaneously mentioned tips and advice on the app at our booth. A survey was sent out directly after the festival.

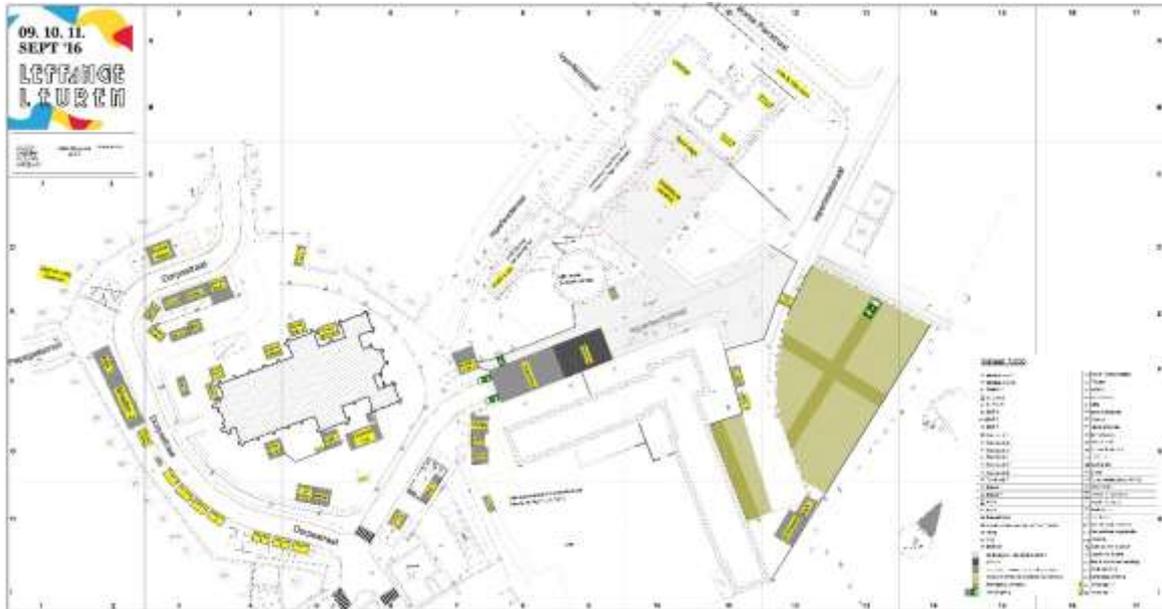


Figure 19: Leffingeleuren floor plan

6.1.4 The Wall of Moments booth

While our technical room was located at the main stage, our booth (Figure 20) was installed on the other side of the site. There was always a team present at the booth to help test users, inform and address passers-by about testing the app. For each new test user, we entered their e-mail address and type of smartphone. Test users were able to get access to our Wifi and borrow charged batteries. Each test user who registered in advance also received coupons for drinks. In order to maintain a link between the technical room and the public booth, 300m of optical fibre linked both locations together.

After the Dranouter field trial, we also created a new “Wall of Comments” book (Figure 21) for passers-by and test users to write down their instant feedback and impressions of the Wall of Moments.



Figure 20: Wall of Moments booth



Figure 21: Wall of Comments

6.2 Technical Overview

TaW's main goals of the field trial were threefold:

- To test the integration between TaW's live mixing engine and iMinds's multi-panoramic capture system's RESTful interface to control the cameras from TaW's live editing interface;
- To demonstrate that it is feasible to control the cameras remotely (with TaW's live editing interface being at a location remote from TaW's live mixing engine and iMinds's multi-panoramic capture system);
- To demonstrate the maximum possible production chain and integration between partners' systems.

Figure 22 unterhalb shows the production chain demonstrated at the field trial, from inputs (cameras, UGC) through the live mixing engine to the user interface to be mixed (with a back channel for camera control), to the outputs (SDI and DASH stream).

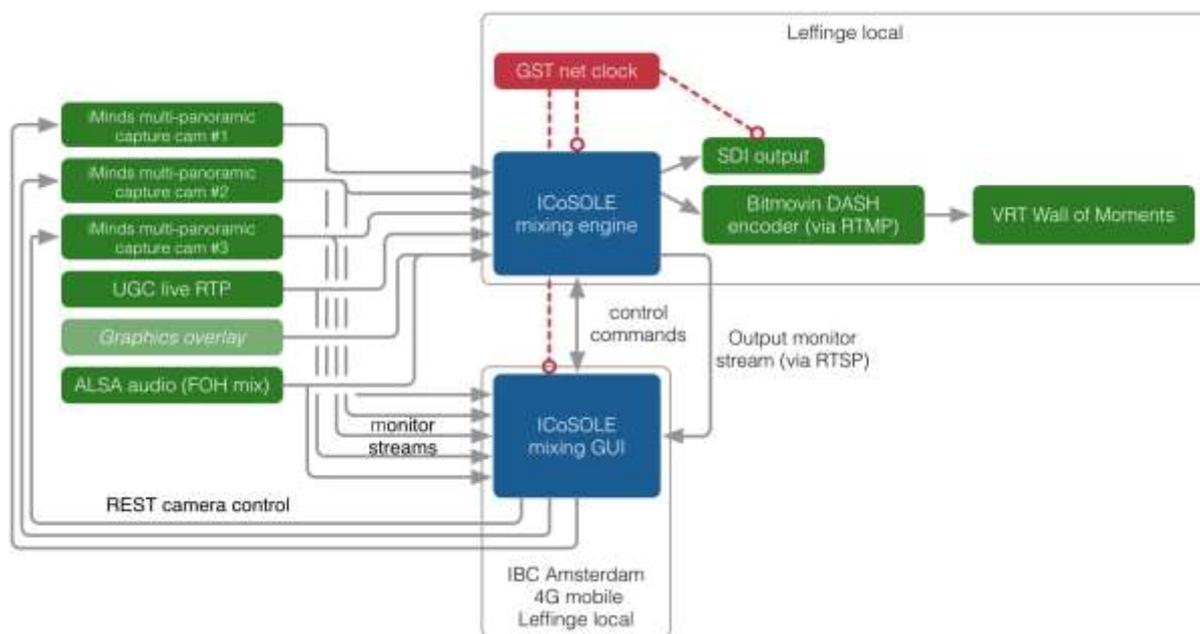


Figure 22: Professional capture production chain

6.3 Professional Audio and Video Capture

As in the previous field trial, a physical “screen” output from iMind’s multi-panoramic capture system was connected via HDMI to TaW’s live mixing engine [D4.6] and was then available as an input that could be selected and sent on air using the live editing interface. Since the first field trial iMinds had developed a new RESTful interface to allow for integration and to this end TaW added additional functionality to the live editing interface in order to:

- Map one or more virtual “cameras” to a given “screen” and allow the operator of the live editing interface to switch between the cameras;
- Allow the operator to pan, tilt, rotate and zoom the camera using the mouse in the live editing interface.

Specifically at the field trial, three physical “screen” outputs were attached to TaW’s live mixing engine and three “virtual” cameras were assigned to each of these “screens”, thus the operator could select between and control a total of nine cameras, with the restriction of a maximum of three visible in the live editing interface and/or on air at any point in time. For example, the first input in TaW’s engine could be switched between cameras 1, 2 or 3 from iMinds’s multi-panoramic capture system, the second input between cameras 4, 5 or 6 and so on. By dragging the mouse horizontally or vertically within the

camera's "view" in the live editing interface, the operator could control either the pan and tilt of the camera, or by holding the "alt" key at the same time, the rotation and zoom of the camera. The live editing interface is illustrated in Figure 23 unterhalb.

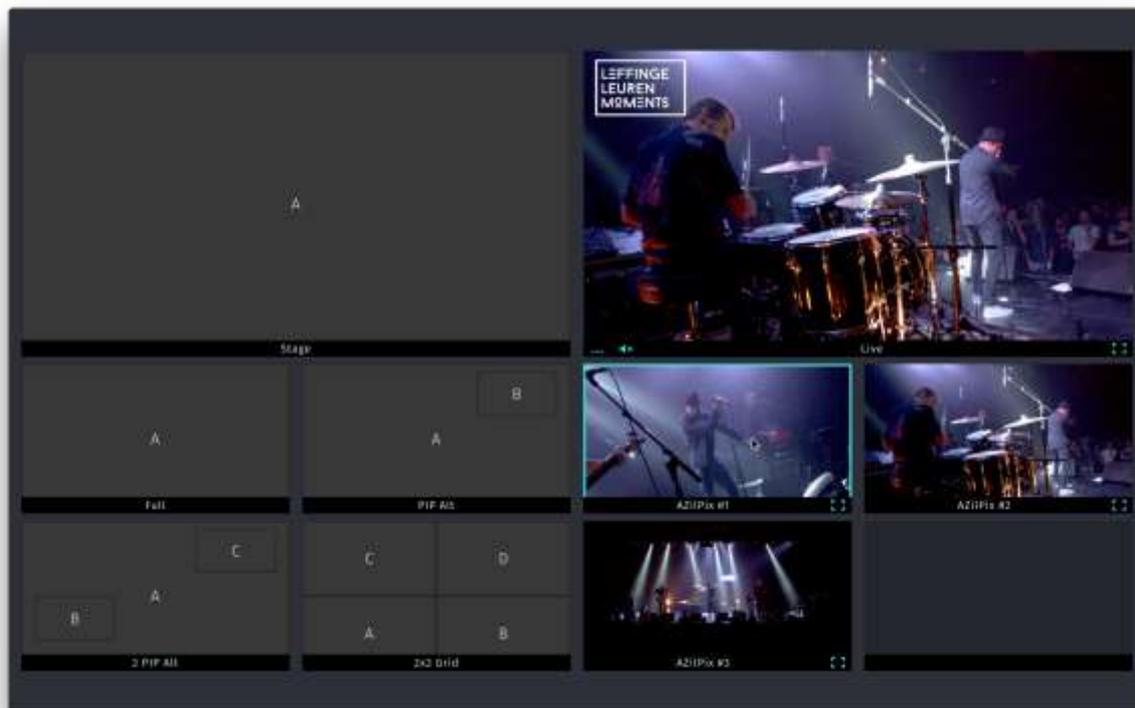


Figure 23: Live editing interface with three camera inputs

Two scenarios were tested during the field trial: with the live editing interface situated "locally" at the festival venue, and with the live editing interface situated "remotely" on the floor at IBC in Amsterdam, which overlapped part of the field trial. Specifically it was important to test:

- The robustness of the interfacing between the live editing interface and iMinds's multi-panoramic capture system;
- Test the round-trip latency involved in sending control commands from the live editing interface to iMinds's multi-panoramic capture system, those commands being accepted and implemented in the camera's view, then being ingested into TaW's live mixing engine and finally the preview being sent back to the live editing interface. A large amount of latency would make a difficult for the operator to correctly control a camera to achieve the desired framing;
- Test any additional latency introduced when the live editing interface was "remote" to see if it was feasible to remotely control the cameras in a live environment.



Figure 24: Live editing interface being demonstrated at IBC

In addition to the integration with iMinds's multi-panoramic capture system, the field trial was also used to demonstrate live streaming of UCG from mobile devices into TaW's live mixing engine, and the final output was streamed to BIT's cloud-based server to be DASHed so that any compatible web browser (desktop or mobile) could be used to view the live stream. Furthermore this live stream was displayed on VRT's Wall Of Moments during the festival. As such the full production chain demonstrated live integration of systems from 5 consortium partners (iMinds, TaW, JRS, BIT and VRT). Integration with BBC's IP Studio had already been tested and demonstrated at a previous field trial.

Finally the SDI output was recorded to professional-grade AJA KiPro recorders supplied by VRT for archival purposes.

6.4 User-Generated Content Capture, Analysis and Playout

The capture apps were those used in Edinburgh, with some small improvements. The backend systems for processing the online and offline had some additional features. The audio analysis modules, which were completed later than expected due to the shift of work from DTO to JRS were integrated into the system, providing audio quality scores along with the video quality scores. The system for automatic assessment of Moments in order to provide a ranking in Trademark has also been integrated as a service in the backend system.

In contrast to Edinburgh, a dedicated Wifi was used in Leffinge. This limited the places where live streaming could be performed to the areas well covered by the Wifi. Nearly 800 images and videos were sent in by the test users, of which about 60 were live streams.

6.4.1 Capture (Moments and Focus)

The major addition to the Moments app was interactivity. Aside from providing (automatic) feedback on sent-in content, we were able to contact test users and ask them to participate in polls. We also added the possibility to upload photos (in addition to videos). This was extremely useful for people with limited bandwidth available. Lastly, we added a live (Facebook-like) feed, which we used as the baseline of our storytelling.

We also used the Focus prototype to follow trending topics on social media and to forward posts to Trademark. Aside from following #hashtags and @mentions on Twitter and Instagram, it was also possible to intercept every Instagram post in a radius of 500 metres of the centre of Leffinge. When starred in Focus, the post was automatically transmitted to the Trademark queue.

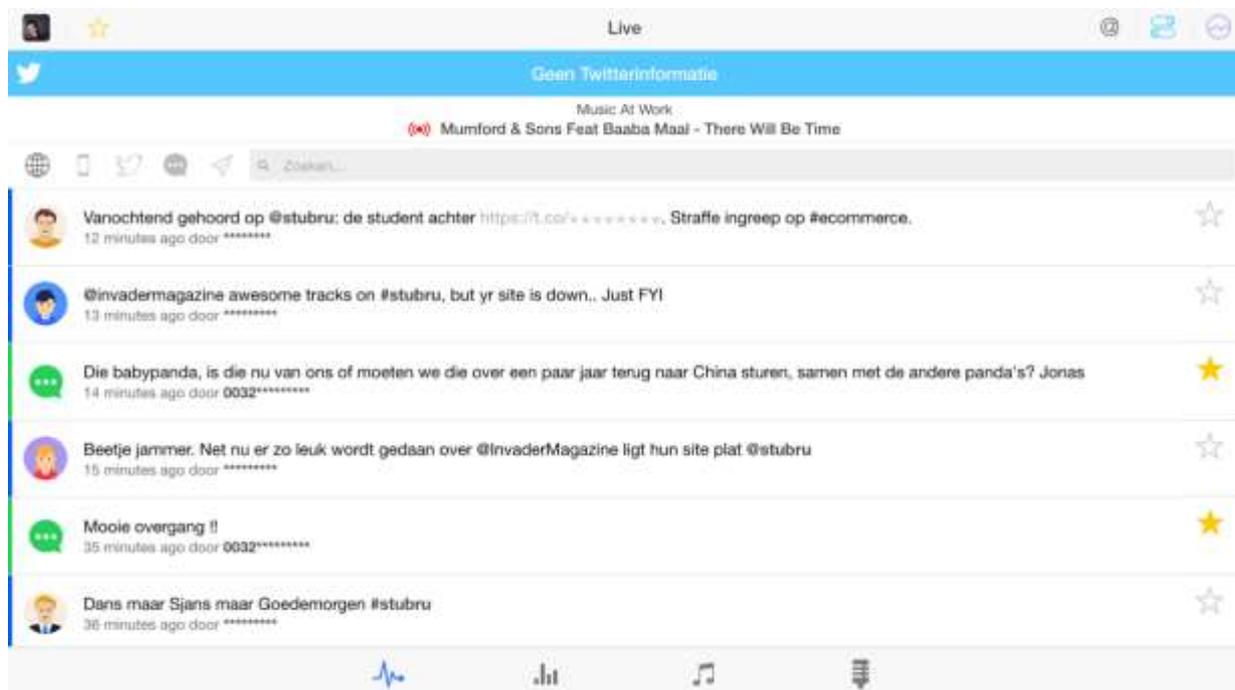


Figure 25: Starring posts in Focus

6.4.2 Storytelling (Trademark)

Note: more elaborate information on Trademark can be found in D4.5 Tools for semantic scene creation.

Content selection

The basis of this app hasn't changed: we use it to review incoming content (which could also be photos now, originating from social media too). Using backend services from JRS, we were able to do the ranking of the incoming content queue smarter, however. Moreover, behind this ranking system a learning algorithm was put. This way, future ranking was affected by present decisions on which content was approved or rejected.

Using the Microsoft Clarifai service, content was automatically tagged as well. This way, it was very easy to find every Moment containing something e.g. 'guitar' or 'beer'.

Interaction and steering

The link with test users was much more pronounced in this field trial, thanks to the advanced interaction possibilities in Trademark. To start, we had a map of the festival terrain (**Fehler! Verweisquelle konnte nicht gefunden werden.**), with dots indicating the location of test users present at the site. Using a rectangle, we selected a subgroup of these users (e.g. everyone in front of the main stage), and contacted them with a certain instructions (e.g. Send in some pictures of the stagedive!). This had the added value of only contacting people who qualify for a certain assignment. It makes no sense to bother people who are on the other side of the terrain or people who are not present at that time. We also used this segmented interaction to invite nearby people for an interview we later used in our documentary.

Aside from textual interaction, we implemented the option to do polls as well. This way, the communication with test users was richer and had better context. Results of these polls were posted to the feed in the Moments app. Lastly, we were able to request photos or videos as well. This made sure our communication was more organised, as every call was sorted per 'Shout'⁴.

Playout control

Finally, we used Trademark to configure and control the screens on the festival site remotely. This was especially useful because it was difficult to come physically near a screen during the festival (because of the crowd).

6.4.3 Playout

The Wall

We updated the Wall significantly for this field trial. Firstly, we switched from the matrix 2-by-2-view to a single-Moment-at-once view. This way, we were able to make text large and draw more attention to the screen. It was far less cluttered as well. As most Moments only last for a couple of seconds, this was a welcome addition to the experience. Secondly, we made the Wall context-aware by installing beacons nearby. Signals sent out by these beacons were picked up by smartphones nearby. If the Moments app was installed, this was communicated to the Firebase backend and picked up by the Wall. Avatars of nearby users popped onto the screen and Moments of nearby users were more prominently featured as well. Lastly, we integrated the Wall with the live stream produced by iMinds cameras, edited by Tools at Work and streamed by Bitmovin, by integrating the Bitdash player. Using Trademark, we had the option to play a live stream and display it on the Wall during a concert.

Map of Moments

The iMinds Map of Moments demonstrator was showcased at the Leffingeleuren field trial. In particular, it was hosted on an iPad device and on an iMac desktop at the VRT booth on the festival terrain. Festival visitors could leverage these devices to experiment with the MoM application. In practice, few visitors actually grasped this opportunity. As such, no relevant end-user feedback could be collected from this demonstration effort, unfortunately.

⁴ A Shout can be a request for photos, videos; or can be a poll (multiple choice, score, freeform).

7 Other Trials

7.1 Austrian Science Night 2016

On April 22, 2016, ICoSOLE successfully presented the “Wall of Moments” to several thousands of interested Austrians at the 7th Austrian Science Night (“Lange Nacht der Forschung”).

For the seventh time, almost all Austrian universities and research centres opened their doors to show insights about current projects, new findings and technologies. At more than 2000 partly interactive stations researchers showed experiments and explained entertainingly and easily understandable their work and research topics.

The Wall of Moments was installed at five stations in two provinces of Austria and showed short clips of the actions of other stations throughout the event. Every visitor could download the Moments App and contribute by recording own highlights/impressions and share personal Moments with other visitors all over Austria. About 80 clips were contributed, and the feedback about the Wall was very positive.



Figure 26: The Wall of Moments at the Austrian Science Night.

7.2 MNM Start 2 DJ

During the week of July 25th – 29th, VRT’s radio station MNM organised their yearly DJ contest, MNM Start 2 DJ. During that contest, several DJs battle for the opportunity to be MNM’s weekly DJ for a full year. MNM’s target audience ranges between 16 and 30, and a lot of this audience is reached via social media (Twitter, Instagram, Facebook, Snapchat). This yielded the idea to do more with social media – during the event– for this year’s edition.

To achieve this, we used the Focus app to follow what was being posted about the event on Twitter and Instagram, combined with SMS messages sent to MNM’s phone number. On Twitter, the audience was able to share photos and videos using the event’s hashtag. On Instagram, the same hashtag could be used, but adding the location of the event was enough to pick up the post in Focus. From the incoming feed, we could star relevant posts, which were added to the ‘approved’ list.

Secondly, we built an alternative version of the Wall ('Screen', Figure 28), to show the latest posts on screens installed at the event site. Only posts on the 'approved' list were shown. The screens were installed right next to the contesting DJ, which made them very visible for the audience. This yielded more interaction with the audience, because they had direct feedback from sending in content (in the form of being displayed on the screen).



Figure 27: Technical setup



Figure 28: Screenshot of modified Wall ('Screen')

The biggest takeaway from this test for us was a confirmation of a good UX for the Focus app. The event lasted for five consecutive days, and –due to agenda overlaps– we were only able to be present the first and last two days. Instead of suspending the test the second and third day, however, the MNM production crew took over the process of curating content and putting it on the screen flawlessly; and integrated this into their production workflow. This was for big part thanks to the fact that the Focus app was developed in close collaboration with VRT's other radio station, Studio Brussels, who did a lot of prototype testing in advance of this event.



Figure 29: The MNM Start 2 DJ event

8 Conclusions

The Edinburgh Festival provided a high-profile opportunity for a field trial with the benefit of having BBC production facilities to work alongside. Not only did this provide access to existing audio and video feeds, but gave visibility of the project to other parts of the BBC who may not otherwise be aware of it. The project also worked alongside other BBC R&D projects going on at the venue, including sharing streams and expertise. Despite some early teething problems with some of the captures, eventually some high quality capture and live playout was achieved. The development of some of the components throughout the trial proved largely successful and probably would not have happened under simulated lab conditions. The integration between live professional and UGC capture and playout proved to be successful, albeit with some workarounds due to networking limitations.

The RESTful interface provided by iMinds proved to be an effective way of controlling and switching between cameras using TaW's live editing interface. Even though it was not possible to view all nine cameras live, using three live inputs and switching each input between three cameras proved to be a decent compromise allowing control of all nine cameras. The latencies involved in having a remote live editing interface also proved to be surmountable, and did not cause any operational issues whilst mixing a live performance, including camera control. The ability to mix between nine virtual cameras and the flexibility in framing enabled by the pan/roll/tilt/zoom camera control made the final output significantly more professional and interesting. All in all the second field trial demonstrated significant integration between partners' systems, resulting in a relatively harmonious end-to-end production chain that stood up well to the rigors of a multi-day live production.

We did a fair amount of trialling with the Wall of Moments across several events. This yielded incremental updates to the capture app (Moments), the social app (Focus) the editorial app (Trademark) and the app (The Wall), which resulted in straightforward, simple-to-use and yet powerful apps. Working with a real event and radio production crews also proved to be very instructive experience. We experienced issues a real production crew would run into and were able to fix them on a regular basis. The development of the Focus app was also done in a very lean way, in close contact with Studio Brussels (a radio station owned by VRT), who have a large user base on social media. This made for a very user-friendly editorial tool, which proved its use during the MNM Start 2 DJ event. In conclusion, the shift to a more lean way of development and a lot of testing with potential clients in-between (events as well as day-to-day) yielded more useful and stable prototypes. These prototypes fit better into existing production workflows, and may well be the basis for a fully-fledged product sometime in the future, which is an aim of this project.

Overall, the field trials covered all the important components of the project, with many of them being integrated together at each trial. While there wasn't one major trial that combined all the components, the smaller trials allowed for tighter integration between the components that were tested at each trial. The different trials also offered different scenarios and content, which could test the components differently and gathered a greater variety of test material. All the trials proved to be successful providing realistic tests, good quality captured material and exposure of the project to a wider audience.

9 References

- [D2.4.2] "Evaluations v2"
- [D4.2] "Algorithms for content analysis and for automatic content pre-selection"
- [D4.5] "Tools for semantic scene creation"
- [D4.6] "Audiovisual rendering engine and interface for linear broadcast production"
- [D4.7] "Audiovisual playout-engine for interactive playout"

10 Glossary

Partner Acronyms

BBC	British Broadcasting Corporation, UK
BIT	Bitmovin GmbH, AT
DTO	Technicolor, DE
iMinds	iMinds Vzw, BE
JRS	JOANNEUM RESEARCH Forschungsgesellschaft mbH, AT
TaW	Tools at Work Hard+Soft Vertriebsgmbh, AT
VRT	De Vlaamse Radio en Televisieomroeporganisatie NV, BE

Acknowledgement: The research leading to these results has received funding from the European Union's Seventh Framework Programme (FP7/2007-2013) under grant agreement n° 610370.