

5

IoT Solutions for Large Open-Air Events

**Francesco Sottile¹, Jacopo Foglietti¹, Claudio Pastrone¹,
Maurizio A. Spirito¹, Antonio Defina¹, Markus Eisenhauer²,
Shreekantha Devasya², Arjen Schoneveld³,
Nathalie Frey³, Hors Pierre-Yves⁴, Paolo Remagnino⁵,
Mahdi Maktabdar Oghaz⁵, Karim Haddad⁶,
Charalampos S. Kouzinopoulos⁷, Georgios Stavropoulos⁷,
Patricio Munoz⁸, Sébastien Carra⁸, Peeter Kool⁹
and Peter Rosengren⁹**

¹LINKS Foundation, Italy

²Fraunhofer-Institut für Angewandte Informationstechnik, Germany

³DEXELS, the Netherlands

⁴OPTINVENT, France

⁵Kingston University, United Kingdom

⁶Brüel & Kjaer Sound & Vibration Measurement, Denmark

⁷CERTH/ITI, Greece

⁸Acoucity, France

⁹CNet Svenska AB, Sweden

Abstract

This chapter presents the main results of the MONICA project, one of the five large-scale pilot projects funded by the European Commission. MONICA focuses on the adoption of wearable IoT solutions for the management of safety and security in large open-air events as well as on the reduction of noise level for neighbours. The project addresses several challenges in eleven pilots of six major European cities using a large number of IoT wearables and sensors. The chapter first introduces all MONICA challenges in the context of large open-air events and then presents the corresponding adopted technical solutions, the defined IoT architecture and the perspective in integrating a wide range of heterogeneous sensors. On one side, the focus is on the solutions that have been adopted to improve the crowd management, crowd safety

and emergency responses by using wearables for both visitors and the security staff at the events, including also the adoption of video processing and data fusion algorithms to estimate the number of visitors and its distribution in the event area and to detect suspicious activity patterns. On the other hand, it describes how innovative Sound Level Meters (SLMs) can be deployed to monitor the sound propagation within the event area while reducing the noise impact on the neighbourhood.

5.1 Introduction

Sustainable urban development is recognised as a key challenge at the global level. The “Mapping Smart Cities in EU” report issued by ITRE (Industry, Research and Energy Committee) of the European Parliament points out that more than half of all European cities with more than 100,000 residents have implemented or planned measures to have “smarter” cities. The report defines six elements of Smart City characteristics: Smart Governance, Smart Economy, Smart Mobility, Smart Environment, Smart People, and Smart Living. Smart Living, the subject of MONICA, is defined as ICT that enables lifestyles, behaviour and consumption. It implies also healthy and safe living in a culturally vibrant city with diverse cultural facilities and incorporates good quality housing and accommodation.

On this background, the main focus of the MONICA project, one of the five large-scale IoT pilots funded by the European Commission, is on how to leverage on the use of multiple existing and novel IoT technologies to address real-world challenges in open-air event scenarios. More specifically, MONICA demonstrates a large-scale ecosystem that uses innovative wearables and IoT devices with closed-loop back-end services integrated into an interoperable, cloud-based platform capable of offering a multitude of simultaneous and targeted applications. In order to demonstrate the ability of deploying IoT solutions to address real challenges, three Smarter Living ecosystems with high societal relevance have been chosen: Security, which tackles the problem of managing public security and safety during big events; Acoustic, concerning the impact of noise propagation in neighbourhood areas while carrying out open-air festivals; IoT Platform, to demonstrate to developers and service providers the ability of the MONICA platform to cooperate with other Smart City IoT platforms, Open Data portals, and generic enablers. All the ecosystems will be demonstrated in the scope of large-scale city events; however, they have general applicability for any dynamic or fixed

large-scale application domain in smart cities such as open markets, fairs, cultural venues, sports events etc.

There are eleven main events targeted by MONICA and organised in six major European cities. The following provides a brief overview of the involved pilot cities, explaining their nature, characteristics and challenges.

Tivoli Gardens in Copenhagen. Tivoli Gardens in Copenhagen, Denmark, is a famous amusement park and pleasure garden that attracts around 4.5 million visitors annually. Tivoli organises several outdoor concerts during Friday evenings of the summer months. The number of complaints by the residents for noise pollution has increased over the years and the flow control in the venue is difficult and requires a lot of security personnel.

Kappa FuturFestival and Movida in Turin. Kappa FuturFestival in Turin, Italy, takes place once a year for two days. Various national and international artists for electronic dance music attract more than 45,000 participants for the occasion. The second pilot in Turin is the Movida in the San Salvario district: lots of bars, restaurants and liquor stores with thousands of visitors moving around the district. The city of Turin wants to address the noise and security situations in these events.

Dom and Port Anniversary in Hamburg. The city of Hamburg, Germany, attracts thousands of visitors by hosting several events every year. The focus of MONICA is both on the Hamburger DOM, the biggest public festival in Northern Germany held in the centre of Hamburg city three times a year, and on the Hamburg Port Anniversary, which comprises maritimes parades, historic sailboats and ships, music, fireworks and food. Both events attract a vast number of visitors and the major concern is crowd and crew management.

Fête des Lumières and Nuits Sonores in Lyon. In the city of Lyon, France, urban festivals are part of the local culture. For four nights in December, during Fête des Lumières a variety of different artists light up buildings, streets, squares and parks all over the city. Concerning music instead, Nuits Sonores festival has been known as one of the great European meetings in the realm of innovative music and creativity, taking over different locations of the city for six days and nights. City of Lyon has identified that the security and safety of participants are the most critical areas to be improved for these events, moreover, reducing complaints from residents about noise pollution is also an objective.

Pützchen's Markt and Rhein in Flammen in Bonn. The federal city of Bonn, Germany, hosts two large-scale events: Pützchen's Markt and Rhein in Flammen. Pützchen's Markt is one of the oldest fairs in Germany and is held every year on the second weekend of September. Huge number of visitors gathering in the narrow streets of the residential area makes the event more challenging for managing security and safety deployments. Rhein in Flammen is an open-air festival which takes place in the public park Rheinaue with three performing stages and fireworks on the banks of Rhein which makes the venue highly crowded. The main challenge that the city of Bonn is interested to handle is the staff management for both pilots. In addition, noise control is one of the important aspects to be addressed for Rhein in Flammen.

Headingley Stadium in Leeds. Headingley Stadium outside Leeds, England, is the home of Yorkshire County Cricket Club, Leeds Rhinos rugby league team and Yorkshire Carnegie rugby union team. It can accommodate 21,000 people and already has security arrangements with statically installed cameras across the stadium and body worn cameras by the security staff. The stadium authorities wish to increase the security by automated video analysis of the streamed images. In addition, they want to add staff monitoring and emergency evacuation planning.

Related MONICA use cases are described in detail in the next subsection.

5.1.1 Main MONICA Use Cases

Crowd and Capacity Monitoring

Crowd and capacity monitoring applications are useful to predict and handle emerging incidents in large open-air events. These applications are primarily based on data from crowd wristbands and CCTV cameras. Some features are reported as follows:

- Crowd count: to know the number of people being in a specific area.
- Crowd density: estimating crowd size and density shown as a heat map.
- Crowd flow: detecting anomalies such as fights, explosions or high-risk queues.
- Redirection: redirect visitors to safer areas using large screens or APPs.

Health, Security and Safety Incidents

Other applications support not only the detection of incidents but also the reporting and handling of them in order to provide a timely response. Relevant

applied devices include smart-glasses, staff wristbands, activity recognition wearables, CCTV cameras, mobile phones and environmental sensors (e.g. for the detection of strong wind). Some features are reported as follows:

- Notify staff/guards about type of incident and related location, for instance, through smart-glasses and staff wristbands.
- Report back to control room, for instance, through smart-glasses, staff wristbands.
- Detect person falling or fights, for instance, through CCTV, activity recognition wearables.
- Report incident using a mobile phone APP.
- Monitor wind speed, for instance, through environmental sensors, to detect high risk for tall amusement rides.

Missing Person

Knowing the exact position of staff members all the time is a high priority when incidents occur as they need assessment or assistance. Moreover, quickly finding a child (who is missing) or a friend is also relevant at some events. For the location of people, staff and crowd wristbands, tracker GPS devices, and mobile phones could be used. Some features are reported as follows:

- High-precision location (staff wristbands).
- Approximate location (crowd wristbands, trackers GPS devices, mobile phones).

Sound Monitoring and Control

Organisers of concerts want to give their performers and audiences the best music experience, but they also wish to comply with local regulations on environmental sound exposure. This produces a dilemma when you talk about outdoor concerts in residential areas. Since high sound pressure levels are necessary for optimal concert sound, there is a risk of the audience becoming disappointed and artists turning down invitations to perform if regulations say that volumes have to be turned down. And even if regulations are met, event organisers still have to deal with residents living next to the venue who complain about the noise coming from the concerts, and that it is affecting their quality of life.

To help solve these challenges, MONICA has developed and deployed an acoustic system providing sound zones at the venue. The system consists of

novel sound field control schemes which provide an optimised sound field in the audience area (called bright zone) while minimising the exposure to noise in neighbouring areas (called dark zones). Thus, at the front of the stage, the music can be louder for a better concert experience, whereas the sound level is reduced outside the concert area and in a selection of quiet zones. The sound levels are automatically controlled by the system, adjusting for changes in weather or size of audiences. Additionally, information about the sound levels is displayed and used for different purposes, i.e. moving to a quieter zone, checking that regulations are kept, studying health aspects etc.

The acoustic system can be divided into three main sub use cases (namely Sound Field Control System, Quiet Zones and Monitoring) that are presented in the following subsections.

Sound Field Control System. Based on an array of loudspeakers and integration with the existing sound system at the venue, MONICA has developed an Adaptive Sound Field Control System (ASFCS), providing an optimised sound field in front of the stage whereas reducing the sound level outside the zone. To accommodate for the variable weather conditions i.e. wind, temperature and humidity, which add complexity to the propagation of sound, the ASFCS employs an adaptive model that adjusts for changes in climate and audience configuration.

The data used for updating the propagation model come from various IoT-enabled devices at the venue such as SLMs, wind and temperature sensors. The IoT devices are integrated with the ASFCS using the MONICA cloud platform. As a result, the ASFCS can continually update the sound propagation model and the loudspeaker signals, using the data from the devices.

Silent Zone System. In addition to the ASFCS, a Silent Zone System provides local silent zones within and close to the loud event area. These quiet zones can be used by security and/or medical staff as well as by the public for any need.

Sound Monitoring. For awareness purposes and to comply with acoustic regulations in the city, IoT SLMs are used to measure sound levels at strategic locations and send the information to organisers, public services, the audience and the neighbourhood. Other IoT enabled devices such as smartphones and wristbands can also be used to measure non-calibrated sound levels.

5.2 MONICA IoT Platform and Data Modelling

5.2.1 MONICA Functional Architecture

The MONICA functional architecture has been iteratively defined considering the use cases and related functional and non-functional requirements. The MONICA functional architecture has been inspired by the reference High Level Architecture (HLA) developed by the Working Group 3 (WG3) of the AIOTI [1]. Thus, following the HLA, the MONICA architecture has been subdivided in different layers as depicted in Figure 5.1. The main difference with respect to the HLA is the definition of two additional layers namely *Device Layer* and *Edge Layer*. In particular, the *Device Layer* includes all the IoT wearables (e.g., crowd wristbands, staff wristbands and smart-glasses) and IoT sensors, which can be either fixed (e.g., SLMs, loudspeakers, cameras, environmental sensors) or mobile (e.g., wireless SLMs, cameras installed in a blimp). The MONICA architecture uses a distributed approach where different gateways (GWs) are deployed at the edge of the environment in which they operate to minimise the delay that both network and cloud modules may introduce. In fact, these GWs have to process in real-time big amount of raw data generated by a large number of IoT wearables and sensors. Thus, the MONICA GWs belong to the *Edge Layer* that incorporates Artificial Intelligence (AI) methods and techniques. In particular, the *Edge Layer* includes: Wristband-GW running localisation algorithms, Security Fusion Node executing video analytic algorithms, SLM-GW

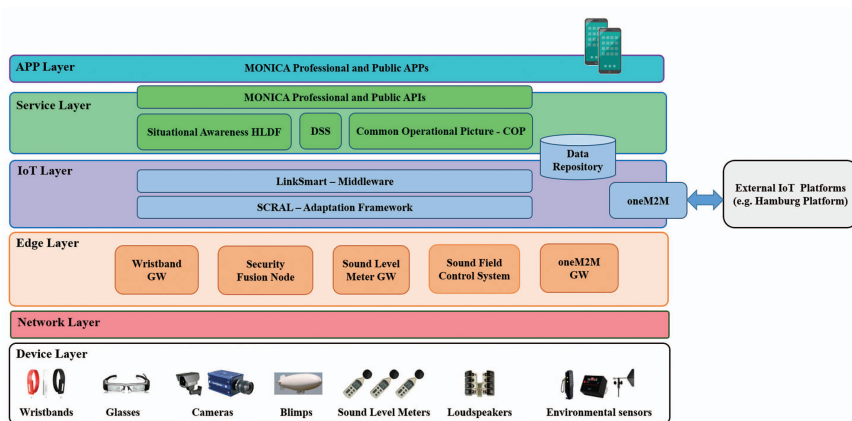


Figure 5.1 MONICA functional architecture.

processing sound data from SLMs, and oneM2M-GW collecting data from environmental sensors. All the functionalities offered by the GWs along with the states of IoT devices are semantically represented by the *IoT Layer* according to the OGC SensorThings standard [2]. In particular, the *IoT Layer* is composed of the following three sub-components:

- The *Adaptation Layer*, here represented by SCRAL, provides technology-independent management of IoT physical resources and uniform mapping of data into a standard representation that can be easily handled by the upper platform modules;
- *Middleware*, here represented by LinkSmart, offers data storage and directory services for resources registered in the IoT platform;
- *External IoT Platform Connectors* that handle the communication with external IoT platforms and the integration of data coming from outside (e.g. from the Hamburg Smart City platform) integrated using the oneM2M standard APIs.

Going upward, there is the *Services Layer* that implements the intelligence of the platform and integrates specific high-level processing modules providing technical solutions to meet the application requirements. The service modules are combined with knowledge base components and a Decision Support System (DSS), whose aim is to propose a set of intervention strategies to assist human operators in gathering context-sensitive information and decision making. Before the *APP Layer*, the MONICA system exposes specific Application Programming Interfaces (APIs) that provide service access points for MONICA application developers and external application developers that want to access MONICA functionalities and information streaming from the platform. Additionally, MONICA provides efficient deployment and monitoring tools that ease the platform deployment process as well as the operational monitoring of GWs, communication networks and system services.

5.2.2 IoT Middleware Platform

The IoT layer of the MONICA platform consists of an adaptation layer and a middleware. “Adaptation” means the chance of having transparent access to physical information without technology-specific knowledge, while a middleware offers storage and directory services. The Smart City Resource Adaptation Layer (SCRAL) is the adaptation layer adopted in MONICA, that provides technology independent management of physical resources and uniform mapping of data into standard representation. On the other hand,

the LinkSmart middleware, together with an OGC SensorThings compliant solution, is deployed as back-end infrastructure, giving support for storage utilities as well as event forwarding solutions and monitoring tools.

5.2.2.1 SCRAL

SCRAL is an interface layer that supports uniform and transparent access to physical resources for sensing and monitoring purposes. Its main function is to ease the integration of devices providing the platform with a standard way to communicate with them. The integration process is implemented by the Field Access Level, a SCRAL sublayer in charge of the physical communication with resources and external platforms. In particular, this sublayer includes specific adapters for each type of protocol and network configuration that can be used either as a web server or client, depending on the type of API available for a given resource. Data gathered from the access level are afterwards delivered to a Core module, a virtualisation layer adopting an Open Geospatial Consortium (OGC)-based Semantic Framework, and finally stored into GOST, a Go implementation of the OGC SensorThings API (see Section 5.2.3 for more details about the standard and data model). The API is only accessible by administrators having authenticated access via a VPN connection. However GOST can also be configured to read Observations from an event broker by subscribing to a pre-defined topic. Further details about Event Broker and LinkSmart services are presented in the following section.

5.2.2.2 LinkSmart

LinkSmart is an open source IoT platform providing support for data storage and management, service and resource directory, stream mining and machine learning. LinkSmart services follow a microservice pattern. Hence, they can be easily integrated and orchestrated based on a concrete use case. In the context of the MONICA project, LinkSmart provides:

- Service orchestration with Service Catalog.
- Brokering solutions with Message Queue Telemetry Transport (MQTT) protocol.
- Authorization proxy service with Border Gateway.

LinkSmart Service Catalog¹ provides a directory of all the services running in the MONICA ecosystem. The Service Catalog contains entries of everything

¹<https://docs.linksmart.eu/display/SC>




Figure 5.2 Service catalog data model.

that is meant to be discovered by other applications and services. A discoverable entry “service” has the data model as shown in Figure 5.2. Each service has a unique “id” and a user defined “name”. The detailed description about the service is added in the optional “description” field. Each registered service has an expiry time beyond which the entry in the registry shall be removed for the service. This is called Time to Live (TTL) and service entry has to be updated before TTL seconds. “Docs” links to external resources such as OpenAPI specifications or wiki pages containing the documentation about the service. This enables direct service to service interaction if valid semantics are followed. “APIs” represent service endpoints, subscription or publish topics and the protocols to be used for communication such as MQTT, HTTP or SOAP. Service entry also contains the details about when it is “created”, when the entry is recently “updated” and when the entry shall be “expired”. The expiration value is calculated based on TTL. Service Catalog provides both HTTP and MQTT based interfaces. HTTP API exposed by the Service Catalog provides the CRUD functions for editing and managing the catalog. Service Catalog also uses MQTT for service registration and de-registration by interacting through a pre-defined MQTT topic.

Event broker provides a message bus for efficient asynchronous communication of sensor data streams implementing the publish/subscribe communication pattern. MQTT is indeed considered as the ideal protocol for low bandwidth, high latency and non-reliable connections which exactly describes a typical IoT scenario. MQTT broker is responsible for maintaining the clients and forwarding the messages across the clients. In the MONICA context different services and resources communicate to each other through MQTT, for example, to get notified when a new device is registered to the platform or also to generate new statistics feeding the application layer. In addition, being the broker a service itself, it can be discovered with the help of the Service Catalog. Concerning security, a MQTT broker can be

vulnerable at different levels, such as network, transport or protocol level. In the network level, using a Virtual Private Network (VPN) for communication between the clients could save a broker from external attacks. In other cases, the LinkSmart Border Gateway is used.

LinkSmart Border Gateway² provides a single point of entry for accessing all the services running in IoT autonomous systems (IoT-AS) such as the MONICA cloud. It provides the functionality to intercept all the MQTT and HTTP REST API requests and performs authentication and authorisation for these requests with the help of identity providers conforming to the OpenID Connection Protocol. In the context of MONICA, Keycloak³ is used as the Identity Provider. Authorisation rules are maintained as user or group attributes in Keycloak and can be defined by administrators based on path/topic and methods, e.g. to limit permissions for certain users to perform only HTTP GET requests on a certain REST API or to allow only MQTT subscriptions on certain topics. The rules are provided by the OpenID Connect provider as a custom claim in the access token to the Border Gateway. To improve performance, Border Gateway caches the access tokens during the span of their lifetime in a Redis database.

5.2.3 Semantic Framework

A Semantic Framework is a semantic data model that implements a distributed knowledge management infrastructure. More specifically, the MONICA Semantic Framework is a conceptual data model focused on the semantic virtualisation of both IoT resources and services employed in the project. This model supports various information domains and adopts the SensorThings API standard [2]: an OGC standard providing an open and unified framework to interconnect IoT sensing devices, data and applications over the web. It is an open standard addressing both the syntactic and semantic interoperability of the IoT; thus, it supports the exchange of information in a shared meaning among entities of a different nature belonging to the same system. It follows REST principles, the JSON encoding, the OASIS OData protocol, the URL conventions and also provides a MQTT extension allowing users and devices to publish and subscribe updates from devices. The UML of the OGC data model is reported in Figure 5.3.

²<https://docs.linksmart.eu/display/BGW>

³<https://www.keycloak.org/>

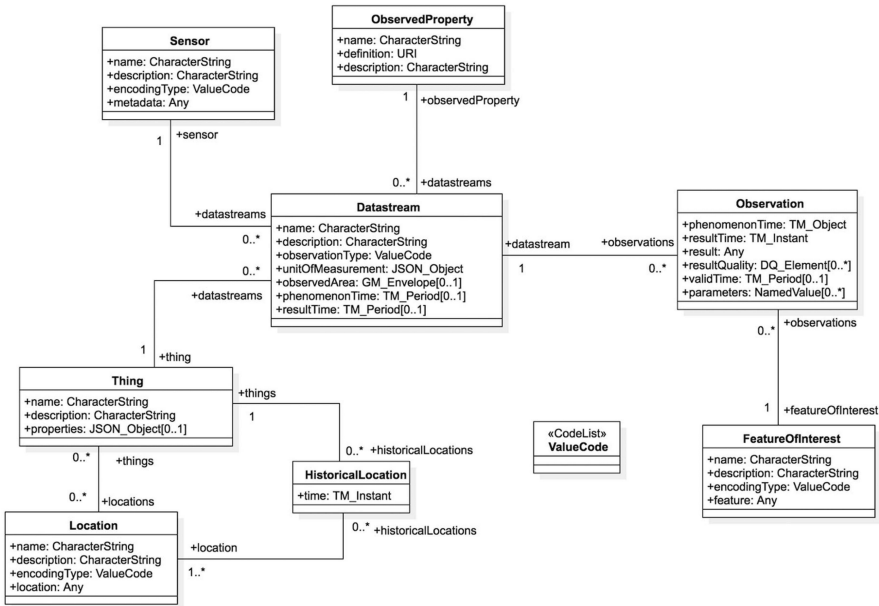


Figure 5.3 OGC SensorThings data model.

The metadata generated by the SCRAL concerning “Things”, “Sensors” etc. as well as the concrete “Observations” (i.e. the sensor data) are permanently stored in the backend using an open source PostgreSQL database. In addition, the OGC server API provides endpoints for maintaining the metadata and writing and retrieving the data over the web in JSON format. The OGC server can also subscribe to different MQTT topics in order to fetch observations from sensors and to store them. Most of the sensor data collected during the MONICA pilot events are stored in the OGC server. Some of the exceptions include raw camera measurements which are processed locally. In case of multimedia data, only aggregated information is stored in the OGC server. The OGC server also acts as the intermediate storage for all the transient data created by different high-level services. Finally, the collected data can be analysed and visualised (live or post-event) using the open source platform Grafana [2, 3] in combination with the Grafana OGC SensorThings Plugin [4]. This offers powerful capabilities for data analysis and data presentation.

5.3 Security and Safety Solutions for Large Open-Air Events

5.3.1 Introduction

This section focuses on security and safety solutions based on wearables as well as video analytics. Different types of wearables have been developed within MONICA. One is a low-cost wristband providing coarse localisation based on a sub-GHz radio and used for crowd monitoring. Another type of wristband has been developed for the security staff. This is a more expensive device than the crowd wristband as it is equipped with a display supporting staff communication with the Command Centre (CC) and adopts UWB technology providing a more accurate localisation. A third type of wearable has been developed by MONICA proving tracking capabilities based on GNSS and LoRa communication; thus, getting rid of a cumbersome anchors deployment so it is able to cover very large areas with reduced cost. Finally, a smart glasses application has been developed to support the security staff to quickly share information about situations during an event.

In addition, MONICA has developed solutions for crowd management and monitoring using existing camera infrastructure, where Artificial Intelligence (AI) modules implementing deep learning algorithms have been used to model complex crowd behaviours and characteristics. Several crowd analysis algorithms have been developed including crowd counting and density estimation, crowd flow analysis, gate counting and crowd anomaly detection.

5.3.2 Crowd Wristband

The crowd wristband is a low-cost wearable device based on a bi-directional sub-GHz radio technology (compliant with the standard ETSI EN 300 220 V3.2.1) able to achieve a radio link larger than 100 m. Within the MONICA project, the crowd wristbands have been used for monitoring the density and location of a very large number of people at open-air events. The collected crowd monitoring data can be used to create heat maps of the crowd; thus, showing crowd size and densities. Moreover, the crowd wristband integrates an RFID interface to support access control and cashless payments. Two LEDs are also available on the crowd wristband, which can be used for entertainment purposes and for crowd control.

The first large-scale implementation of the crowd wristband was deployed by DEXELS in 2014 during two weekends of the Tomorrowland festival in



Figure 5.4 LED show with crowd wristbands in front of the main stage at Tomorrowland 2014.

Boom, Belgium. Each weekend a total number of 125,000 crowd wristbands were active for three days. The system supported “scanning” all 125,000 wristbands in 8 minutes. Within the MONICA project, the wristbands infrastructure has been integrated with the IoT platform and it was demonstrated first in the Tivoli Garden (in Copenhagen, the 26th of April 2019) and then in Lyon in occasion of the Woodstower festival (30th–31st August 2019), where more than 6,200 wristbands were distributed to the audience.

In total, 60 base stations (BSs) were used in Tomorrowland (see image in Figure 5.4), 15 BSs in Tivoli and 15 BSs in Woodstower, spread over the festival area, connected using Ethernet cables and powered by Power over Ethernet (PoE).

5.3.2.1 Technology overview

The crowd wristband can appear in multiple bracelet incarnations. It could be a leather bracelet, a textile wristband or a silicon wearable as in MONICA (see Figure 5.5). The current version of the crowd wristband contains the following components: radio chip and MCU (SoC solution), CR2032 battery, two bright RGB LEDs, RFID/NFC chip, one button, clock, and two antennas (HF RFID and UHF).

Together these parts cooperate to form the crowd wristband solution. The software running on the MCU manages the operation of the wristband. It controls the radio communication, the LEDs and the behaviour of the button press. The MCU can wake up from its deep sleep mode in several ways. On wakeup, the MCU starts its normal operation by listening to radio messages




Figure 5.5 Crowd wristband.

transmitted from a BS. The BS messages synchronise the wristband clocks and send commands either to a particular wristband or to all wristbands. The commands instruct the MCU to light up the LEDs. Each wristband can be addressed separately by means of a unique ID. Depending on the application, this unique ID could be associated with personal details of the visitor wearing the wristband. However, for crowd monitoring purposes, the wristbands in MONICA were anonymously distributed to the audience of the festivals. The button can be used for several user inputs. One example is “Friend-Connect” or “Contact Sharing” as demonstrated during IoT-Week 2019 in Aarhus. In particular, two people can connect and exchange personal information. The only thing that needs to be done to connect is holding two wristbands in close proximity of each other and pressing the button until the LEDs flash green in both wristbands. Upon this action, the two wristbands’ identifiers are included in a “connect” message and sent to the MONICA platform that in turn sends it to an external cloud platform related to the IoT-Week APP. It should be mentioned that this “Contact Sharing” functionality was supported by two preliminary steps where the user first registers himself/herself using the APP and then associates this with his/her own MONICA wristband. Therefore, all the personal information was managed by the IoT-Week APP and the related cloud platform.

5.3.2.2 Crowd wristband TDMA protocol

The crowd wristband protocol uses parallel time slots to allow for higher throughput of wristband messages. Four radio channels (TRX, RX1, RX2, RX3) are used simultaneously as shown in Figure 5.6, where TRX is called “Pilot Channel”. Wristbands messages are sent using a TDMA protocol. Moreover, a CSMA phase in the TDMA scheme is used for sending “urgent” messages like “wristband connects” and button pushes.

The wristband protocol relies on tight clock-synchronisation to support the TDMA protocol. The clocks on the wristband are synchronised by pilot messages sent by each BS. Depending on a unique ID every BS sends its pilot




Figure 5.6 Crowd wristband TDMA protocol.

message in a predefined time slot. The length of a time slot, for both pilots and wristband messages, is set to 3 ms. A maximum number of 16 BSs time slots initiate a new “communication window” as shown in Figure 5.6 (see red interval). The wristband protocol supports a maximum of 16 BSs per “Pilot Channel” (TRX). Hence, the pilot-phase of the messaging windows always takes 50 ms. A wristband uses the pilot message to synchronise its local clock. Besides clock synchronisation, a pilot message contains the “wristband ID range” parameters and optional LED commands. The “wristband ID range” is used to define the logical range of wristband IDs (WUIDs) that need to be polled. The range is defined by a start- and an end-value. The WUID of a wristband is first masked to fall within the polling range before its reporting time slot/channel is determined.

The remainder of the messaging window is used to send a wristband “reporting message”. Each wristband is assigned its unique timeslot and channel (1 out of 4) to send this message. The reporting window size can be adjusted from 0 ms to 200 ms, resulting in a total communication window of 50–250 ms.

A total number of 6 Pilot Channels can be defined in an event area, resulting in a maximum infrastructure of 96 BSs. The default operation mode is 250 ms, resulting in a maximum reporting throughput of almost 6,400 wristbands per second in an event area (thanks to the 4-fold parallelism and 6 different Pilot Channels). A wristband determines the “strongest” Pilot Channel by scanning other channels every 30 s. By keeping a list of strongest channels, it decides whether it is time to switch to another Pilot Channel. This effectively implements a channel handover procedure for the wristbands allowing for larger areas that can be supported by the protocol. This still limits the maximum area size that can be covered by the wristbands; exploring ways to bypass this limitation is part of future work.

The 50 ms mode is used for low-latency LED operation, allowing a new LED command to be sent every 50 ms. This comes at the expense

of not being able to report wristband messages during 50 ms of operation. Moreover, urgent message mode for “wristband connect” and “button push” messages can be sent using a CSMA protocol providing a responsive and correct implementation of the so called “Friend-Connect” feature.

5.3.2.3 Infrastructure for crowd wristbands

The deployment scheme of the crowd wristband infrastructure is depicted in Figure 5.7. As it can be observed, the crowd wristbands need a dedicated infrastructure of BSs that communicate with each other and with the wristbands. The maximum safe range between a wristband and a BS is 75 m. This implies that a wristband must always be at maximum 75 m away from a BS in order to have coverage. This characteristic can be used to design and setup the BS infrastructure for a specific venue. Since there is a limit to the number of BSs (96) there is a limit to the area that can be covered. Hence, currently the spatial scalability is limited. The number of wristbands that is currently supported is limited by the 3 bytes that are used to identify a wristband. There is no inherent limitation to the number of wristbands in the protocol itself. The BS radio is controlled by an ARM based PC board running Linux and the BS software. The BSs themselves are joined together in a software cluster. A unique redundant-communication protocol has been developed that enables the use of multiple physical communication layers between the BSs.

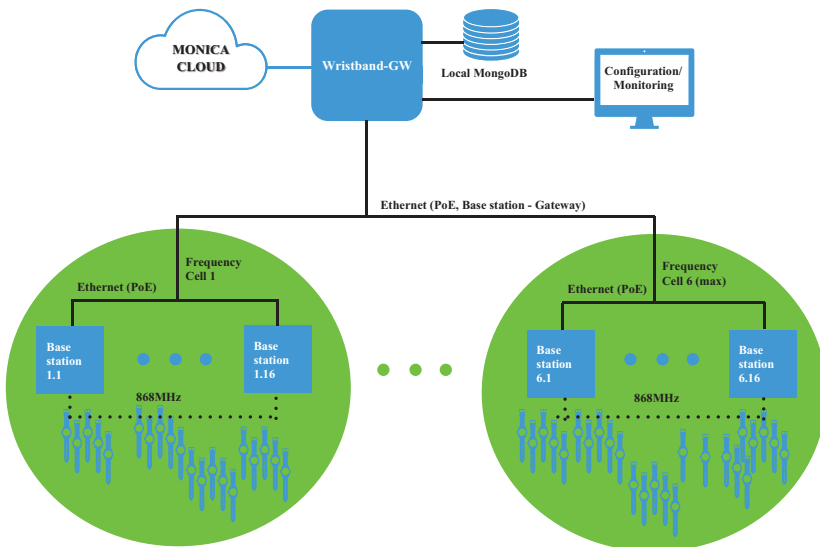


Figure 5.7 Crowd wristbands infrastructure deployment.

TCP/IP based communication, both Ethernet and Wi-Fi, as well as several low-bandwidth wireless communication technologies are supported between BSs and the GW. Altogether this creates a highly fault-tolerant communication channel. If for example the Ethernet or Wi-Fi infrastructure fails, the messages are still sent using the alternative available wireless infrastructure, making the system independent of the festival's infrastructure. A typical communication use case is a message that originates from a wristband, being received by one or more BSs and further transported to the GW.

The GW is deployed locally on the festival premises and it sends wristbands messages to the MONICA cloud. This setup enables the mobile APPs that are running on the visitor's smartphones to interact with the cloud system. The messages that are received from the wristband are collected by the GW that performs a real-time triangulation to drive the crowd monitoring system, heat map visualisation and individual wristband tracking. Since Received Signal Strength Indicator (RSSI) is used as range measurement, the typical localisation accuracy of crowd wristbands is 15–20 m. A management console is available for operators to control the entire system.

5.3.2.4 Services enabled by crowd wristbands

The location of the visitors can be used to calculate a crowd density. The resolution of this discrete density field is typically $5\text{ m} \times 5\text{ m}$. This is a useful feature for crowd monitoring, i.e. knowing the number of visitors in various event areas at any instant. This could also be used to detect high-risk queues (or at least high-risk densities) based on the maximum capacity of these areas. The locations collected by the crowd wristbands are used to create a current overview of the crowd distribution in the event area. These crowd services run in the MONICA cloud where also the DSS implements algorithms to detect over capacity or high-risk queues. The crowd density can be visualised in a dashboard application running in the CC of the event. In addition to the DSS, this information can be used by CC staff to detect hot spots in crowd densities.

5.3.3 Staff Wristband

The staff wristband (see image in Figure 5.8) comes with more features than the crowd wristband. The staff wristband provides a more accurate tracking capability. While the typical accuracy of crowd wristbands is 10–20 m, the accuracy of the staff wristband is less than 50 cm thanks to the adoption of Ultra-Wide Band (UWB) radio technology, compliant to the standard ETSI EN 302 065-2 (UWB Location Tracking).



Figure 5.8 UWB-based staff wristband.

In addition, the staff wristband has an integrated BLE radio that can be used to communicate with a smartphone or smart glasses and also features a LCD screen that can be used to notify or instruct the user.

The fundamental building blocks of the staff wristband system are anchors and tags. Anchors are fixed location UWB nodes, containing at least one so called *master anchor* that is responsible for collecting all the data (wirelessly or wired) from the other anchors. Anchors send/receive messages to/from mobile tags. These messages are used in the localisation process as well as for communicating so called user payloads. These payloads can include e.g. data from sensors integrated into the tags.

The system uses UWB-based geometrical localisation. In particular, the ranging measurements use a Two Way Ranging (TWR) Time of Arrival (TWR ToA) method. The TWR does not require any synchronisation of the clocks at all, however this comes at the expense of having to communicate at least three messages between tag/anchor before the range can be determined. This means that with TWR, less tags can be tracked in a certain amount of time compared to a Time Differential of Arrival TDoA method. Using TWR, the system is able to support 1,200 location updates per second. Hence, 1,200 tags can be ranged running at an update rate of 1 Hz.

The next step in the process is localisation that calculates the position of the tag based on the distances between the tag and the (visible) anchors. The position is calculated using a lateration algorithm. It is worth remarking that the ranging accuracy achieved by the Decawave UWB chip is ± 10 cm in Line-of-Sight (LoS) conditions and ± 30 cm in Non-Line-of-Sight (NLoS) conditions; thus, there is always additive (white) noise present in the measured distances. As a consequence, an exact (closed form) solution of the lateration problem is not possible. One has to rely on an optimisation procedure to calculate the location. Typically, a Non-Linear Least Square (NLLS) method is used. In case of tracking a moving object, additional methods are used. Jitter in the calculated track is usually mitigated using some smoothing or filtering method. In MONICA it was used either an Extended Kalman Filter (EKF) or an Extended Finite Impulse Response (EFIR) in combination with an NLoS detection and mitigation method. These methods result in low jitter while still having acceptable latencies (< 500 ms in case of a 20 Hz update rate).

Figure 5.9 shows the deployment scheme of the UWB staff wristband, where a cluster of UWB anchors are managed by a master anchor that in turn sends the ranging measurements to the GW that runs the lateration algorithm.

5.3.3.1 Technology overview

The UWB staff wristband integrates the following of components: 1.3" 176×176 display, Bluetooth LE, USB and wireless charging, DecaWave DW1000 UWB radio, ARM Cortex M4, and a 400 mAh battery. Furthermore, several sensors and actuators are available: light sensor, IR proximity, pressure sensor, temperature sensor, humidity sensor, microphone, 9 axis IMU (accelerometer, gyroscope, magnetometer), 2 buttons and haptic feedback.

5.3.3.2 Services enabled by staff wristbands

Security staff localisation. The staff wristbands allow for the localisation of staff members. As such, they help to implement several use cases. Besides staff localisation, the wristbands can be used to notify staff members by sending text messages that are displayed on the LCD screen of the wristband. Moreover, the buttons can be used to send notifications to the CC.

Health/security incidents. By leveraging the IMU of the staff wristband, certain health incidents can be detected. A wrist-worn accelerometer is used for recognition of abnormal activities related to stewards and security staff visitors in crowded environments.

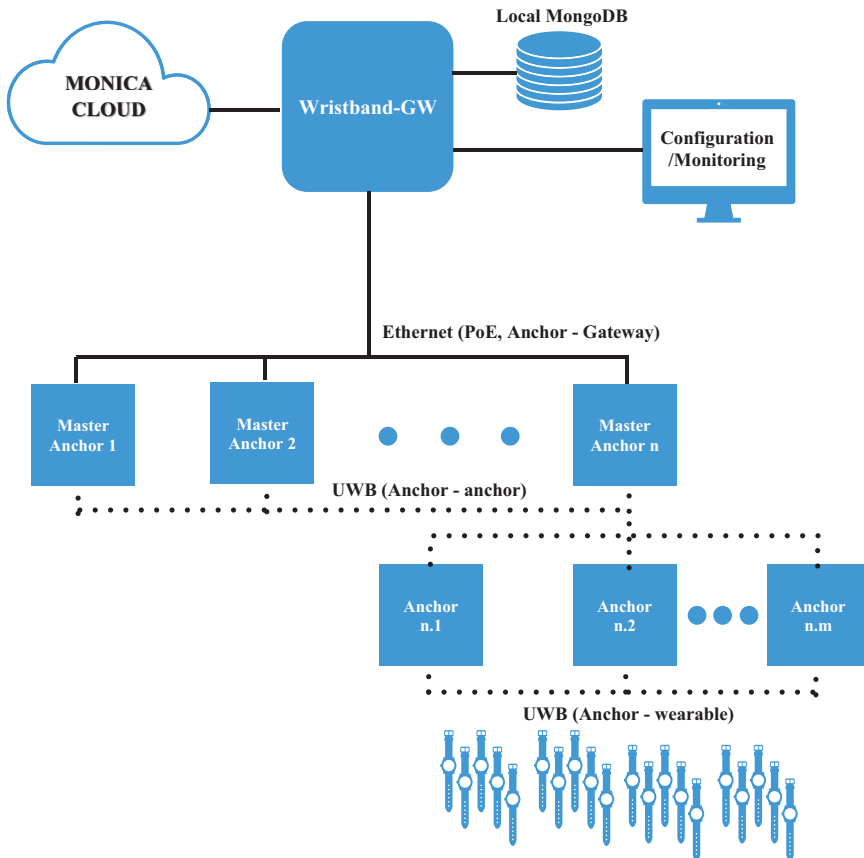


Figure 5.9 UWB staff wristbands infrastructure deployment.

5.3.4 Staff Tracking Device

Tracking devices based on the Global Navigation Satellite System (GNSS) have also been developed in MONICA for staff members to be used in large open-air events. Contrary to UWB staff wristbands (presented in 5.3.3), these tracking devices do not need the installation of UWB anchors for localisation purposes as they rely on signals from satellites. The devices record latitude and longitude data from the GNSS module and periodically send them to BSs using Long Range (LoRa) technology [5]. The LoRa BSs or gateways are integrated with LoRa signal transceivers and forward these data to the MONICA IoT platform. The rest of this section presents both the GNSS tracking devices and the LoRa Gateways.

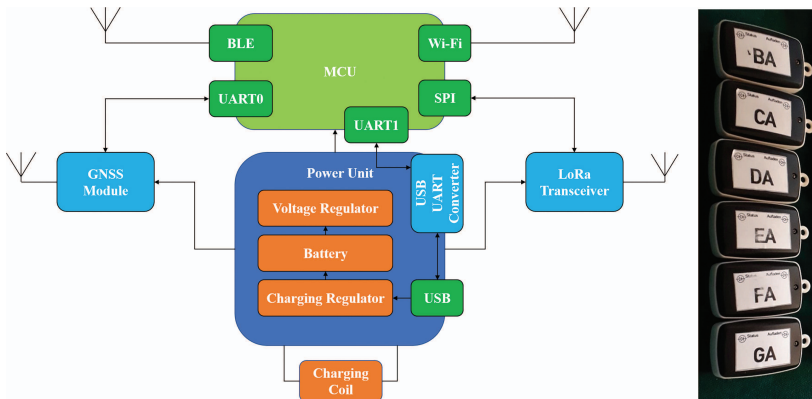


Figure 5.10 Modular diagram of HW components in trackers (left) and enclosed tracking devices (right).

The trackers are encapsulated in a small car key sized form factor. The hardware components are described in Figure 5.10. They contain an 1100 mAh battery which is wirelessly chargeable and lasts for more than 22 hours of active usage. An MCU is responsible for executing the tracker logic for collecting and sending the location information. The MCU is connected to Bluetooth Low Energy (BLE) and Wi-Fi transceivers. But these modules are deactivated in normal operations and used only for special purposes such as software update. The MCU is connected via a Serial Peripheral Interface (SPI) to the LoRa transceiver.

The GNSS module is connected to the MCU via a Universal Asynchronous Receiver Transmitter (UART). The GNSS module calculates the position every second with an accuracy of up to two meters. The localisation accuracy depends on many factors such as presence of high-rise buildings around the tracker. As soon as the MCU is powered on, it sends a registration request to the IoT Platform. The MCU is usually in deep sleep mode in order to save battery. It wakes up at a fixed interval to send the last recorded geo location coordinates. The trackers also have a programmable LED to show the status of the trackers so that the users will know whenever the status of the tracker (e.g. turned off, sending data, sending emergency beacons). The tracker has a button which is used for turning the device on or off and to initiate emergency beacons.

The LoRa gateway is a Raspberry PI device enclosed in a waterproof box. The gateway is integrated with a LoRa transceiver. One or more gateways are strategically installed in the venue in such a way that cover as much area as possible. One gateway can cover an area spanning thousands of meters. The



Figure 5.11 Signal coverage of two LoRa gateways (Red and Green) in the Rhein in Flammen event.

gateways are PoE. The data transferred from the tracker to the LoRa gateway include the device name, latitude, longitude, battery level, Signal to Noise Ratio (SNR) and accuracy of the geo-location. The gateway reads this data and forwards the data received from the trackers to the IoT Platform over HTTP.

Around 45 GNSS trackers have been successfully used by volunteers of the fire brigade, police, public order office and emergency response in pilot events at Bonn in 2019 (Rhein in Flammen and Pützchens markt). There were no concrete failure scenarios except minor signal breaks. The signal coverage was sufficient to contain the entire event venue just with two gateways as shown in Figure 5.11. Message duplication was one of the problems while reading LoRa channels. The reason behind duplication was the spreading of the messages across different time slots. This was filtered in the DSS before sending to the Common Operational Picture (COP) dashboard so that duplicate events are not generated.

Future work with respect to the GNSS tracker is an addition of an energy efficient user interface (e.g. an LCD) showing the messages from the CC. As of now, a simple LED is used blinking different patterns to give the user feedback. This needs to be improved with a descriptive display showing the messages and status.



Figure 5.12 ORA-2 smart glasses.

5.3.5 Smart Glasses

The ORA-2 smart glasses enable hands free mobile computing and Augmented Reality (AR) applications such as remote maintenance, logistics, remote training, situation awareness, and much more. It can run applications as a standalone wearable computer and can connect to a network via Wi-Fi and to any smart device via Bluetooth.

The ORA-2, image shown in Figure 5.12, features a disruptive transparent retinal projection technology. The virtual screen of the ORA-2 has two configurations allowing both “augmented reality” and “glance” modes. This “Flip-Vu” feature allows the image to be either directly in the wearer’s field of view or just below.

The ORA-2 is equipped with a dual core processor GPU, camera, microphone, sound, inertial sensors, Wi-Fi, Bluetooth, GPS, ambient light sensor, and a high capacity rechargeable battery, and for a better ergonomic, a small Bluetooth remote controller has been added. The patented ORA-2 smart glasses hardware platform comes complete with its own flexible Android SDK in order to develop apps and fine tune the user experience.

An application, called MonicOra, has been developed by the Optinvent team for the MONICA project, which is based on a simple architecture so that the security staff can use it after few learning instructions. The goal of this equipment is to quickly share information about situations during an event. The security staff using the smart glasses can send predefined messages, audios, photos and videos to the CC, which can be visualised in the COP dashboard. In turn the CC (through the COP dashboard) can respond using the same way of communication, to one, or some/all smart glasses. The MonicOra APP fits many security use cases involving policemen, security staff, firemen, stadium stewards by changing few configurations with a quick

update. Controlled by a little Bluetooth remote controller in hand, the smart glasses can easily manage all the applications using only 3 buttons.

One important aspect of this device is the battery lifetime. Since display, LED, Wi-Fi, GPS and a lots of other components are working together, it has been learned how to increase the activity in two ways. First, a power bank can be provided to be connected when the APP is running. Second, the users have been instructed to follow the correct usage. This equipment has to be in a sleep mode most of the time and be active very quickly to communicate with the COP. The most important thing is that the security staff work is supported by the smart glasses but are aware of the environment as much as possible at the same time.

Due to the technology, the AR glasses, which have the littlest optical engine size, provide a good experience in middle and low luminosity environments. This system has to be enhanced a lot if the amount of luminosity is high. In fact, all the MONICA demonstrations have been performed during low luminosity ambient so that the camera could be more sensitive and efficient.

5.3.5.1 Services enabled by smart glasses

The security staff equipped with ORA-2 glasses is completely aware of his environment because of the transparent lightguide (in sleep mode, no light disturbs the eye) and could easily report any situation by sending predefined messages, recording audio, picture or video (in active mode). Since the smart glasses are constantly tracked by GPS, the COP dashboard can send back the right information to the closest wearer(s) to find, for instance, a criminal sending an ID (for police officers), a plan with security exits (for stadium stewards), or audio and written translation for searching a lost child etc. With this equipment, combined with the other devices involved in the MONICA project, either a city or an event organiser can better manage all the deployed security forces in a more efficient way to prevent, take care, respond and quickly solve many situations.

5.3.6 Crowd Heat Map Module

5.3.6.1 High level data fusion anomaly detection module overview

A High-Level Data Fusion and Anomaly Detection (HLDF-AD) module represents the first stage of heterogeneous data aggregation and intelligent algorithm aiming to provide more complex and refined information as

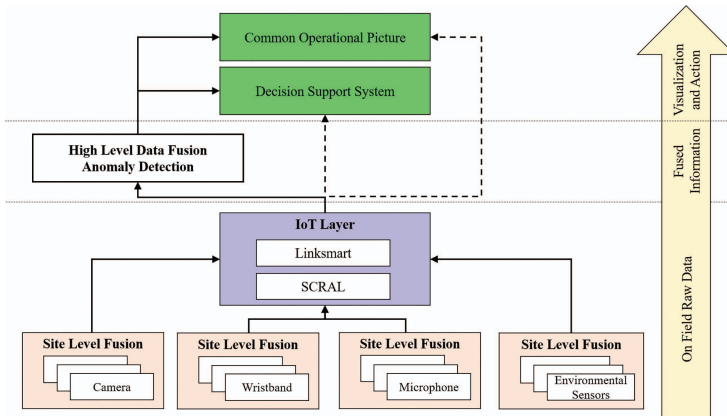


Figure 5.13 Simplified version of the MONICA architecture.

Table 5.1 Raw data available for HLDF-AD module

On Field Raw Data	Input Source	Main Fields
Crowd density global	Cameras	Geo Spatial density map that identifies people distribution
Gate people counting	Cameras	Entrance-Exit people events
People flow	Cameras	People movement trend
Wristband positions	Wristbands	Single person localisation
Sound level	SLMs	Geo spatial Sound Heat Map

requested by the DSS and COP modules. More specific, such a module acquires raw data broadcasted by MONICA IoT wearables and sensors deployed in the field and tries to fuse them in order to generate added value information and feed the upper level’s components with real status information of the monitored area and eventually notifying anomaly conditions. Figure 5.13 shows a simplified version of the MONICA architecture just to show the data flow and highlight the role of the HLDF-AD module. From a general perspective, the MONICA ecosystem foresees a first stage of onsite elaboration (if possible) for each kind of sensors involved. The HLDF-AD module can combine them in order to complete or certify different input sources data and extend meaning towards higher-level modules.

Table 5.1 lists the most important raw data available for the HLDF-AD module. It shall be taken into account that these table reflects a complete set of MONICA sensors that might not be available for all monitored sites as it happened during some pilot executions.

Based on the inputs reported in Table 5.1, the HLDF-AD module should be able to produce both detection in terms of instantaneous on-site

observation and prediction information as an extension of historical data collection and trend analysis on that. For the second stage, a huge amount of historical information is required.

5.3.6.2 Crowd Heat Map Algorithm overview

Among different types of output information, the most significant example is represented by the Crowd Heat Map estimation, which is an estimation of 2D East-North people geographic distribution with respect to a reference position. The HLDF-AD module combines raw data coming from cameras and wristbands in order to calculate people distribution in a monitored geographic area and estimate the total amount of people attending a specific event. In particular, the input information data exploited by such algorithm are Crowd Density Global, Gate People Counting, and Wristband Positions. In principle, Crowd Heat Map and Crowd Density Global could be the same information. However, both information can be combined into a unique geospatial matrix that describes a more refined instantaneous people distribution. It is worth remarking that Crowd Density Global might be affected by issues depending on camera configuration. For instance, the number of cameras or the related coverage could be too low with respect to the event area, and a wrong setup of them (e.g. too high, too low, very low environment visibility) might have impact on the results. Moreover, the possibility for the HLDF-AD module to perform cross-check between camera results and other data sources allows to enhance the confidence level of the output result.

The Crowd Density Global is the most important input from HLDF-AD perspective in view of Crowd Heat Map provisioning. The HLDF-AD module acquires them and, after a clean-up procedure, follows some logic steps considering also historical information and Gate People Counting. This helps to eventually correct the total amount of people based on the principle that the total number of people at time T_1 equals the total number of people at time T_0 plus the number of people entered/exited to/from the monitored area. It is assumed that the monitored area is fenced, where people can enter or exit just passing through the gates monitored by the MONICA cameras. For instance, if at time T_0 there are 100 people in the monitored area and it has been estimated that 10 people have entered between T_1 and T_0 , then at T_1 the HLDF-AD module assumes that there are 110 people. Potential differences raised by the Crowd Density Global have been cleaned by the HLDF-AD algorithm that estimates a new refined Crowd Heat Map. The next step considers Wristband Positions data. The HLDF-AD module calculates another Crowd Heat Map considering the Wristband Positions, i.e. counting

the number of positions inside each geospatial cell of the matrix as 2D East-North distance with respect to a reference position. The new computed Crowd Heat Map can be used as the first output in case of missing cameras or in order to refine the confidence level of the Crowd Density Global computed at the first stage. This data process shall be carried out by means of a rough estimation of the percentage of people who take and use the wristbands with respect to the total amount of people attending the event. Moreover, it should be assumed that the distribution of the people using the wristband is equal to the distribution of the people not wearing the wristbands. Furthermore, it must be taken into account that the percentage of people with a wristband should be significant with respect to the real total amount of people in order to obtain reliable results (at least 20%).

5.3.6.3 Crowd Heat Map at the Woodstower event

During the Woodstower festival (in Lyon, 30th–31st August 2019) the Crowd Heat Map service performed by the HLDF-AD was tested. Since this pilot did not have any cameras installed, the Crowd Heat Map was estimated taking as input only the wristbands' positions from the Wristband-GW. In this pilot, the MONICA partners anonymously distributed up to 6,200 wristbands to the audience of the musical event. The position of each wristband was estimated by the Wristband-GW and transmitted every 4 minutes to the MONICA cloud. The HLDF-AD module stored temporarily each wristbands' position, and every 4 minutes calculated the Crowd Heat Map. This service was registered into the Service Catalogue of the MONICA platform. The size of the monitored area was 300 m by 200 m. The representation was performed on a 10 m × 10 m cell size.

Figure 5.14 shows an example of the estimated Crowd Heat Map based on crowd wristband positions. It can be seen that, the crowd was mainly concentrated in front of the four stages.

5.3.7 Crowd Management and Monitoring Using CCTV Cameras

With the help of one of the most prevalent IoT technologies such as IP surveillance cameras, the MONICA project offers intelligent and autonomous solutions for crowd management problems based on cameras. These solutions require minimum human intervention and can be scaled up indefinitely with no compromise on reliability. MONICA solutions for crowd management and

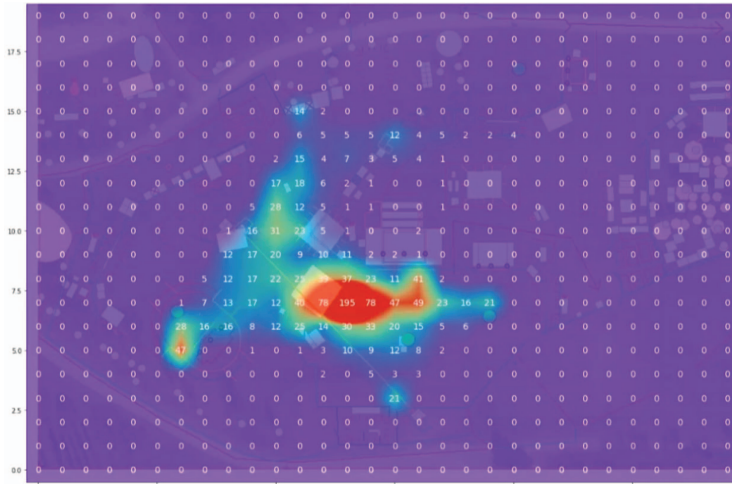


Figure 5.14 Example of a crowd heat map estimated during the Woodstower event.

monitoring can be adopted to the majority of the existing surveillance infrastructure which significantly reduces the costs yet improves the performance and productivity. Unlike classical labour-intensive monitoring systems which suffer from issues like scalability, reliability and cost-effectiveness, intelligent surveillance systems can be scaled up indefinitely with no compromise on reliability. In the last decade, the computer vision community has pushed on crowd behaviour analysis and has made a lot of progress in this field. The emergence of deep learning and Convolutional Neural Network (CNN) in the last decade has boosted the performance of image classification techniques and has started having a positive impact on crowd behaviour analysis. CNNs have gained ground in crowd monitoring and behaviour analysis. In MONICA, deep learning has been used to model complex crowd behaviours and characteristics. Several crowd analysis algorithms including crowd counting, localisation, density estimation, crowd flow analysis, gate counting, crowd anomaly detection, fight detection, and object detection have been developed in MONICA and deployed in several pilot plans across Europe. The following subsections describe some of these algorithms in more detail.

5.3.7.1 Crowd counting and density estimation

Crowd counting and density estimation are of great importance in computer vision due to its essential role in a wide range of surveillance applications including crowd management and public security. However, drastic scale




Figure 5.15 Left: CNN for crowd counting and density estimation. Right: the original image, the Ground truth and the predicted density map.

variations, the clutter background scene, and severe occlusions make it challenging to generate high-quality crowd density maps. MONICA offers crowd counting and density estimation algorithms based on CNN. The proposed solution aimed to address a wide variety of crowd density levels by incorporating a high-level prior into the deep convolutional neural network. The high-level prior learns to classify the count into various groups whose class labels are based on the discretised number of people present in the scene. The count class label allows us to estimate coarse count of people in the given regardless of scale variations thereby enabling the network to learn more discriminative global features. The high-level prior is jointly learned along with density map estimation using a cascade of CNN networks as shown in the following Figure 5.15. The two tasks (crowd count classification and density estimation) share an initial set of convolutional layers which is followed by two parallel sets of networks that learn high dimensional feature maps relevant to high-level prior and density estimation, respectively. The global features learned by the high-level prior are concatenated with the feature maps obtained from the second set of convolutional layers and further processed by a set of fractionally strided convolutional layers to produce high resolution density maps.

5.3.7.2 Crowd flow analysis

Crowd flow is another informative metric in crowd management and behaviour analysis. There is a critical capacity where flow begins to decrease as the crowd's density increases. A dense crowd with high flow magnetite poses a serious safety threat and might lead to a human stampede. In visual surveillance, optical flow algorithms have become an important component of crowded scene analysis. The application of optical flow allows crowd motion dynamics of hundreds of individuals to be measured without the need to detect and track them explicitly, which is an unsolved problem for dense crowds. MONICA uses Flownet CNN which first produces representations of




Figure 5.16 CNN for crowd flow estimation.

the two temporally consecutive images separately, and then combines them together in the ‘correlation layer’ and learn the higher representation together. The representations then will be up-sampled to reproduce the flow map in original image scale. The correlation layer is used to perform multiplicative patch comparisons between two feature maps. Figure 5.16 shows the FlowNet architecture used to measure crowd flow and dynamic.

5.3.7.3 Fight detection

Fight and public violence are prevalent phenomena in large public crowds. Aside from the possible injuries and fatalities, violence can cause damage to public assets and disturb public peace and tranquillity. An automated fight detection model could be a great addition to the existing surveillance systems deployed in public large events such as concerts, stadiums, and amusement parks. The proposed fight detector in MONICA uses a modular architecture, with distinguished components, which were responsible for the fundamental operations, such as video capturing and event detection. Particularly, a computer vision algorithm constitutes the core of the system and it captures mid-term motion patterns by isolating regions of interest in each frame of a video stream and forming point trajectories through the optical flow. After that, the enclosed trajectories over the last frames extract a motion histogram for each region.

5.3.7.4 Pipelining in the security fusion node

The Security Fusion Node (SFN) sits as a separate component to the Video Processing Pipeline and the video algorithms; however, in practice they often run on the same processing node. The SFN acts as an interface between the outputs of the algorithms and the higher-level MONICA services. By having awareness of all the cameras (via a camera registration process), the

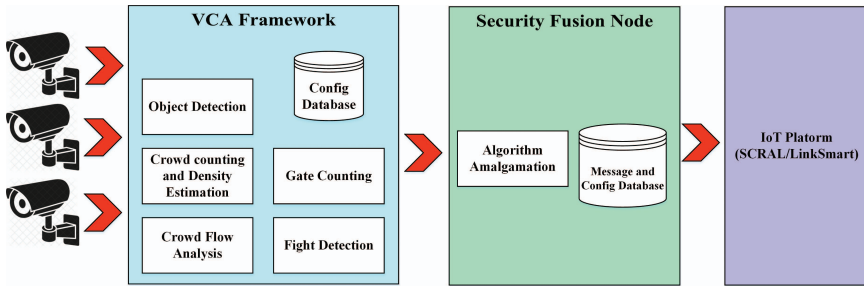


Figure 5.17 High-level overview of where the SFN is located within the MONICA architecture.

SFN is broadly speaking a lightweight multithreaded REST API responsible for forwarding messages from the edge layer up to the cloud and adding additional information where required. The SFN is a kind of central gateway concept that facilitates the platform for information fusion. As such, this node takes as input the output from the algorithms running on the local cameras, whether they be IP cameras, smart glasses or camera equipment mounted on a blimp. Figure 5.17 shows the SFN architecture relative to MONICA.

5.4 Sound Monitoring Solutions for Large Open-Air Events

5.4.1 Introduction

The MONICA project aims at applying IoT technologies for the management of large-scale outdoor cultural events. These types of events come often with various disturbances to the neighbourhood. One of the most noticeable is noise. Moreover, among the different pilot tests in the MONICA project, many of them concern music, where the acoustic experience is prominent for the participants. For both reasons, noise impact on the neighbourhood and quality of sound for the audience, sound monitoring is an important piece of the management of such outdoor events. Additionally, the noise level is often regulated by local legislations, specifying generally a threshold that should not be exceeded otherwise exposing possibly the event organisers to penalties. It is then important to be able to set thresholds and alerts in the sound management system.

A sound monitoring system is made of IoT SLMs that are deployed at different places, in the neighbourhood and in the venue. These units send




Figure 5.18 Sound level meter box.

different acoustic data to a cloud platform, where they are recorded, possibly further processed and sent to an interface for visualisation and alerts. The visualisation concerns outputs of individual sensors and higher-level quantities based on multiple SLMs. These elements are described in the following sub-sections.

5.4.2 Sound Level Meter

A Sound Level Meter (SLM) is generally composed of a high-quality microphone, connected to a board for signal conditioning, processing, local storage, and a display. For the MONICA project, the SLM is put inside a water-proof box, that includes as well a wireless router acting as a mobile Wi-Fi hotspot and a power bank for power supply to the SLM and the router. From the box, the microphone sticks out, protected by a windscreen against wind and rain. An image of the system is shown in Figure 5.18.

The SLM sends data through the router inside the box, and therefore relies on the telecommunication network (3G or 4G). The SLM can also connect to any existing Wi-Fi network in the vicinity.

The types of data that are sent depend on the requirements for the pilot test. For all cases, two different types of levels, the so called L_{Aeq} and L_{Ceq} , and spectrum data, are measured every second. L_{Aeq} and L_{Ceq} provide averaged levels, commonly used for environmental acoustic measurements. The spectrum decomposes the noise level into 1/3 octave frequency filters, from 12.5 Hz to 20 kHz. This represents a vector of 33 values. Sound levels and spectrum data occupy 116 bytes and are sent every second.

For advanced acoustic processing, audio recordings could be necessary. In this case, the amount of transferred data is much higher, since the sampling rate of the recordings could be relatively high. To reduce the required bandwidth, the data are compressed (but without loss) before sent out on the network. For the MONICA project, the required bandwidth is 1 Mbit/sec (since the compression of the audio recordings is dynamic, it varies, but it is generally less than 1 Mbit/sec).

The sound data are first sent to a dedicated cloud, called SLM-GW, for data management, storage and data reduction. The data are then accessible to the MONICA cloud through a RESTful interface. The data reduction tasks in the SLM-GW consists of processing the acoustic levels or the audio recordings, to provide averaged values and advanced data to the MONICA cloud. Two of the advanced calculations performed in the SLM-GW are the Annoyance Likelihood Index and the Contribution analysis. The Annoyance Likelihood Index is a metric between 0 and 10 indicating the level of noise annoyance [6] for the neighbourhood, 10 meaning a maximum annoyance. It is based on the comparison between initial measurements of background noise (before the event) using the levels L_{Ceq} , and the same type of levels during the event. This metric is available for every minute to the MONICA cloud. The purpose of the contribution analysis is to estimate the noise contribution of the venue in the neighbourhood. The noise level measured by a SLM located in the neighbourhood is often a mix of contributions: from the venue, but also from acoustic sources nearby such as cars. Therefore, to recover the contribution from only the event, specific processing needs to be performed. Two approaches are considered: one based on spectral data (implemented in the DSS module, see Section 5.4.3), the other is based on audio recordings (implemented in the SLM-GW). As mentioned earlier, the required transfer data rate between the SLMs and the cloud for audio recordings is relatively high.

The Sound Monitoring System has been tested during different pilot tests in Europe. To illustrate its implementation, we discuss experiences of two different pilot tests in 2019: KappaFutur Festival and Rhein in Flammen.

KappaFutur Festival. This event is located in Torino. It takes place every year in Parco Dora and is dedicated to electronic and techno music. It gathers about 50,000 participants during two days around 4 stages. Focusing on noise impact, low frequencies are pointed out by the neighbourhood as the main reason of annoyance, together with excessive levels during some music performances [6]. The noise monitoring system setup for 2019 edition in July

consisted of 8 SLMs, of which four inside the venue, in front of each stage, and four outside, on balconies of dwellings.

The requested acoustic data are the same as for the previous case, so the data sent by each SLM was 116 bytes every second. For the SLMs in the dwellings, very small data losses have been noted. For the SLMs inside the venue, variations were seen in the quality of the communication network, causing some significant losses (not much the first day, but mostly the second day of the event). One explanation could be that the high number of people gathered in a relative limited area could generate situations where the cell phone network is overloaded.

Two additional SLMs have been set up, sending audio recordings, that require 1 Mbit/s data transfer rate. These SLMs were located outside of the venue, and the purpose was to try out the contribution analysis module based on audio recordings on one hand, and to support another system in the MONICA project called the Adaptive Sound Field Control (ASFC), on the other hand. For these two SLMs, the data loss rate was high because of insufficient network bandwidth outside the venue to support the required transfer data rate.

Rhein in Flammen. This event is located in Bonn, Germany. Here also audio recordings for six SLMs were used, relying on the telecommunication network. Based on local experiences concerning reliability and bandwidth, the local network operator was selected accordingly.

The bandwidth requirement for sending audio recordings is 1 Mbit/s/SLM which is equivalent to 429 MB/hour/SLM but the observed averaged data transfer was 316 MB/hour/SLM, thanks to the audio lossless compression performed in the SLMs. Very small amount of data loss has been noted. It demonstrates the feasibility of sending high sampling rate audio recordings, at the condition of a good telecommunication network support. However, one should be aware of the additional cost due to the high volume of transferred data.

5.4.3 DSS for Sound Monitoring

The DSS is the intelligence layer of the MONICA platform. It is used for a variety of applications, including sound monitoring and control. As an input, it uses data provided by the HLDF-AD and GOST components, such as sound heat maps, A- and C-weighted sound levels, sound spectra as well as location information from IoT devices such as wristbands. The output of

the DSS is subsequently forwarded to the COP dashboard in order to be visualised in a clear and consistent way.

In a real-world scenario, information can be affected by uncertainty. In these cases, it is often not easy or desirable to group all pieces of information into crisp sets (sets with a binary membership where objects are either members of the set or not) based on precise parameters. Instead, fuzzy sets can be used for information classification based on imprecise criteria. For this reason, the DSS is developed as a Fuzzy Logic Device, producing deterministic output from deterministic input starting from a set of rules relating linguistic variables to one another using fuzzy logic. For the mapping to be performed, a subset of deterministic values are converted into fuzzy values, and vice-versa [7].

The fuzzy set theory was first introduced in [8] as an extension of classical fuzzy models. A fuzzy set is a collection of objects that do not have explicitly defined criteria of membership. Instead of a binary membership (i.e. 0 or 1) to a set, each object has a grade of membership $[0, 1]$ instead, that indicates its degree of truth in a subjective way. In other words, its degree of membership indicates how that object “fits” into the set.

In complex sound management scenarios, there are two distinct types of problem knowledges that can usually be inferred for a situation, *objective knowledge* and *subjective knowledge* [9]. Objective knowledge refers to quantitative variables that can be used to accurately represent information. Subjective knowledge on the other hand corresponds to usually not-quantifiable knowledge in the form of verbal statements. Obtaining objective knowledge for a specific event, area or situation though is not easy and can often be inaccurate or not applicable; many times gaining subjective knowledge through the interaction with experts is better suited to capture the imprecise modes of reasoning that is essential for the ability of people to make decisions in an uncertain environment [10].

The DSS uses both types of information to create fuzzy antecedent-consequent “IF-THEN” rules that can be used to propose different courses of actions. One such example that describes a straightforward course of action is – *If the sound volume is too loud, lower the volume of the loudspeaker*. In this example, the notion of *too loud* is a linguistic variable that forms the antecedent part of the rule and can have a different meaning in different settings. The notion of *lower the volume of the loudspeaker* represents a crisp course of action and is the consequent part of the rule.

The DSS consists of four main modules: the fuzzification, knowledge base, decision making and defuzzification modules.

Input from HLDF-AD and GOST is received as crisp numerical values. A fuzzification module is then used, responsible for the mapping of such values to fuzzy sets. The resulting sets can be used in turn to activate all relevant rules by calculating their membership functions. The knowledge base is the process model of the system. It consists of a database that contains the rule structure for each different venue and event that the MONICA project is employed, constructed based on expert input. The fuzzy rules together with the input from the fuzzification module are combined to generate output by the decision-making module. The defuzzification module aggregates the rule consequents and selects the highest-rated one to produce a crisp output as the outcome of the Fuzzy Logic Device (FLD). This outcome is forwarded to the COP for visualisation. This allows stakeholders to have a real time assessment of the event with regards to local regulations.

Within the MONICA project, work has been carried out, as a collaboration between CERTH and Acoucité, to implement a DSS for sound monitoring for two pilots in Lyon. The French regulation regarding sound in festivals⁴ was then taken as reference. The main requirements of the decree are:

- The sound pressure level (15 minutes average) should not exceed 102 dBA and 118 dBC at any place accessible to public in the audience area.
- For thresholds in the neighbouring area, the values are dependent on the background noise levels existing at the measurement point. A measurement campaign to determine the background noise levels was thus required. The requirement is that sound pressure level (15 minutes average) in the neighbouring area should not exceed the background noise level at the measurement points by more than the “emergence” values. Emergence values are defined in terms of Overall A-weighted Sound Pressure Level (L_{Aeq}) and in terms of octave bands from 125 Hz to 4 kHz. Emergence values depend on the period (day/night) and the duration of the event.

IoT SLMs of Noise Monitoring System (see Section 5.4.2) is set up to provide every second with the following acoustic data: A-weighted, C-weighted and 1/3 octave frequency bands sound pressure levels (1 second average).

⁴Decree no. 2017–1244 of 7th August 2017, in force since October 31st 2018.




Figure 5.19 Woodstower 2019, Left: IoT SLMs in audience area (highlighted in red), Right: Municipalities where SLMs outside the venue have been deployed.

This data cannot be directly used for assessing the event with regards to the regulatory limit thresholds. Thus, further computations are required at DSS. These computations are of two types:

- Obtaining 15 minutes average Sound Pressure Level from 1 second average Sound Pressure Level;
- Obtaining octave band spectrum from 1/3 octave band spectrum.

Once those computations are done, relevant data can be shown on the Noise Monitoring System Display and alerts can be set for monitoring threshold exceedances.

An example of the application of DSS for Noise Monitoring System at the Woodstower festival in Lyon, France during August of 2019 is presented here after.

5.4.3.1 DSS for sound monitoring at Woodstower festival 2019

A total of eight IoT SLMs have been deployed for the Woodstower festival, of which four in the audience area and four in the residential area located in municipalities around the venue of the festival (see Figure 5.19). The devices in the audience area have been located near the sound engineer's console at each of the four stages in the venue.

The IoT SLMs provided overall sound pressure levels and 1/3 octave band spectra and the connection was established using a Wi-Fi network specifically deployed for the MONICA project.

Positioning of the SLMs was planned to cover municipalities located to the north and south of the venue in order to include critical areas independently of meteorological conditions (i.e. wind direction varies from one year to another).

The devices sent overall sound pressure levels and 1/3 octave band spectra. Data were sent from IoT SLMs to the MONICA cloud using 4G connection through the MiFi device inside the IoT SLM box.



Figure 5.20 Sound levels in audience – Woodsfloor stage of the Woodstower festival.

Sound Monitoring Display for levels in the audience, as shown in the COP dashboard, is presented in Figure 5.20. Levels at the most critical point in the audience are estimated based on results from preliminary acoustic measurements done by Acoucité.

5.4.4 Sound Heat Map Module

The Sound Heat Map provides an estimate of the sound pressure level (SPL) at other positions than the one being measured by the SLMs. This is done using a forward sound propagation model, based on the existing sound propagation model “Nord2000”. The model is built in 2D and requires the position of sound sources, positions of reflecting surfaces (walls) and a computation area (surface and grid size).


The computation model uses SPLs measured by the IoT SLM located in front of the stage to dynamically adapt the power of each source included in the model. Then, the map is calculated and becomes available in the MONICA cloud via a REST API.

It is nowadays common practice to use directive loudspeaker systems in large-scale concerts. Thanks to progresses made in the past few decades in

terms of loudspeaker design, today systems allow getting directive stages even in low frequency sounds (sub-woofers). Thus, when computing a sound map of a venue, the possibility of controlling the directivity of sound sources is essential for getting accurate results in terms of spatial distribution of sound pressure levels within the audience.

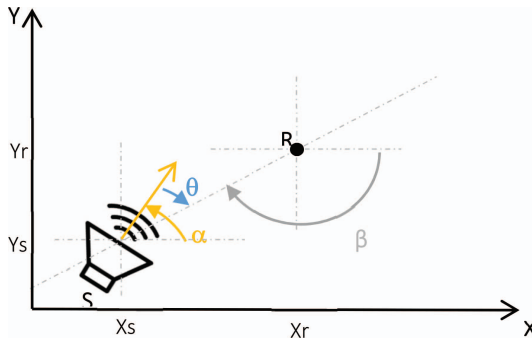
The directivity can be expressed as a cardioid function. The basic equation of a cardioid pattern is:

$$D(\theta) = \frac{1}{2}(1 + \cos(\theta))$$



Cardioid polar pattern

where D is the directivity factor (from 0 to 1) between the source and the receiver, and θ is the angle with respect to the source's radiation axis. In a projected system (conic projection, ex: lambert 93), the angles are calculated from the Cartesian coordinates of the source and the receiver (X_s, Y_s, X_r, Y_r) as shown in the figure below:



Cartesian coordinates of the source and the receiver.

$$D(\theta) = \frac{1}{2}(1 + \cos(\theta)) \quad \text{with}$$

$$\theta = (\beta - \pi) - \alpha \quad \text{and}$$

$$\beta = a \tan 2 \left(\frac{Y_s - Y_r}{X_s - X_r} \right)$$

The two inputs are the azimuth of the source direction (loudspeaker) and the azimuth of the path between the receiver and the source. Both are defined relatively to the X axis of the Cartesian extent. The angle β is calculated from the X and Y Cartesian coordinates. The computed directivity factors for cardioid source are used in the computation model for each source-receiver pair.

5.4.4.1 Sound Heat Map in Woodstower festival

Woodstower is a music festival taking place in a massive park in the north-eastern part of the urban area of Lyon (France) at the end of August. During four days, musicians performed on four stages, namely: Mainstage, Wood-floor, Scène Saint Denis and Chapiteau. The event goers can walk around the venue and attend various performances.

The main challenge for the event manager is to respect the new French regulation (reported in Section 5.4.3). As explained in Section 5.4.2, IoT SLMs are an excellent way to control the levels at a long-term monitoring point (usually near the sound engineer's console). However, they do not allow getting information about the sound level distribution within the whole venue area. The Sound Heat Map gives that opportunity.

A computation model for the Sound Heat Map was built based on information provided by the Woodstower organisers such as venue layout and sound systems configuration, shown in see Figure 5.21.

The event manager chose only loudspeakers and subwoofers with a cardioid directivity pattern. The code for the Sound Heat Map computation should then be able to:

- Use cardioid sound sources.
- Calculate a Sound Heat Map for each of the four stages and to display the resulting global Sound Heat Map (i.e. the summation of the four stages) on the COP dashboard.
- Display the global level in dBA and dBC based on an energy summation of all the band frequencies.

Each stage has been modeled using a simplified model of the sound system (three cardioid sound sources per stage). The obtained directivity is shown in Figure 5.22.

Implementing the directivity pattern in the initial implementation of the heat map did not affect the calculation performance and allows a better identification of sound levels in different areas inside the festival site.




Figure 5.21 Sound system configuration for the main stage of Woodstower festival 2019 (d&B audiotechnik).




Figure 5.22 (Left) directivity of main stage at a frequency of 40 Hz. (Right) computed Sound Heat Map with three stages active (Scène Saint Denis, Chapiteau and Woodsfloor).

5.5 MONICA APP Layer

5.5.1 MONICA APIs

All applications and apps that need to access MONICA functionalities need to use the available MONICA APIs. There are three main MONICA APIs targeting different needs: Professional API, Public API, MessageHub.

5.5.1.1 Professional API

The Professional API is intended for components delivering data and services to the professional staff at the event. The main users of the API are the Common Operational Picture UI and the different apps that are part of MONICA.

The Professional API is based on the following technologies:

- MQTT interface for receiving messages for updating the COP status.
- Authorisation model based on Keycloak REST API.
- Odata-based API for retrieving resources from the IoT DB.
- Integrates SQL database with an OGC Sensorthings database for storage and management of time series of observations.

The Professional API provides the following main functionalities:

- Incident classification and management.
- Division of a geographical area into zones and subzones.
- Mapping of incidents, sensors, facilities and people/groups/crowds to zones.
- Fast retrieval of current status of the situational objects of interest.

5.5.1.2 Public API

The Public API is intended to be used by public apps and is basically a subset of the Professional API. The content that is accessible depends on what the event organiser makes public. Typically, the following items are available in the API:

- Public points of interest, for instance, position of medical services, toilets etc.
- MONICA collected data, which is made public. For instance, people count in different areas.
- Feedback collection, i.e. allowing the public apps to collect simple feedback of the event.

5.5.1.3 MessageHub

In addition to the rest-based API there is also a message hub implemented for pushing information to the apps and user interfaces. The main reason for supporting a push-based approach is to improve scalability, the clients will be notified of new available data rather than polling themselves. It

will also increase the responsiveness of the application. The message hub is implemented using SignalR⁵.

5.5.2 COP Dashboard

During an event monitored with MONICA, the CC will interact with event staff and build situational awareness using the main COP dashboard. The main interface is based on a map of the event venue with symbols marking different objects of interest. The main interface is an HTML-based app that can be used with any type of device that supports HTML 5. It is possible to run it on computers and tablets, but it is primarily designed to be used on devices that have a larger screen. It is possible to run the COP UI on a mobile phone, but it will have limitations due to the screen size.

Currently there are three main user interfaces that can be selected:

- Crowd/Security Monitoring that displays all the information related to Crowd/Security monitoring.
- Sound monitoring that displays all the information concerning sound levels etc.
- Staff view.

Depending on what is monitored at different pilot sites the COP dashboard is tailored to only show relevant information for the event. In the COP it is also possible to filter out information to unclutter the dashboard. Figure 5.20 shows an example of the COP view for sound monitoring while Figure 5.23 shows an example of the COP view for Crowd/Security Monitoring.

5.5.3 Professional Sound APPs

One example of an app built with the Professional API is the Professional Sound Application that is intended to be used by sound managers, sound engineers and event managers. With the application the users can follow the sound measured by the SLMs in real time. The application visualises 1/3 Octave spectra, LAeq and LCEq. It is also possible to combine two SLMs and compare them easily, see Figure 5.24. Furthermore, there is functionality to report sound feedback in different locations which will be shown as “sound incidents” on the COP dashboard.

⁵<https://en.wikipedia.org/wiki/SignalR>



Figure 5.23 COP dashboard showing incidents and crowd counting.

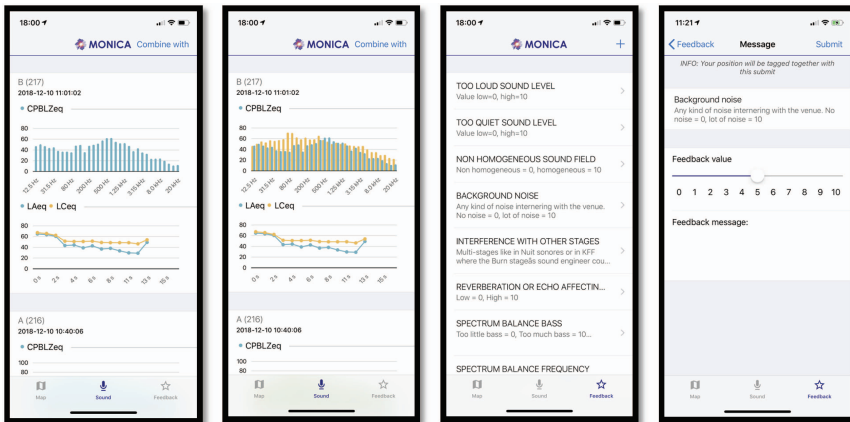


Figure 5.24 Views of the MONICA professional sound app.

5.6 Conclusion

This chapter presented a set of IoT solutions, developed by the MONICA project, able to improve safety and security as well as to reduce the noise level for neighbours in open-air cultural events inside cities. These solutions are based on a large-scale deployment of innovative wearables and devices interconnected with closed-loop back-end services integrated into an interoperable cloud-based platform.

To support crowd management and monitoring use cases, MONICA has developed three types of wearables: a low-cost crowd wristband providing coarse localisation suitable for crowd monitoring, a more expensive wristband for security staff providing a more accurate localisation based on UWB and a third type of wearable with tracking capabilities based on GNSS and LoRa communication able to cover larger areas with a reduced installation cost. Finally, a smart glasses application has been developed to support the security staff to quickly share information about critical situations during an event. In addition to wearables, MONICA employed a set of crowd monitoring solutions relying on existing camera infrastructure. In particular, crowd analysis algorithms have been developed including crowd counting, crowd density estimation, crowd flow analysis, gate counting, and crowd anomaly detection.

Finally, concerning the sound monitoring use cases, MONICA has developed a sound monitoring system comprising IoT SLMs, the Sound Heat Map module and the DSS. For each pilot, these sound components have been used to cover two main areas of interest: the event venue and the neighbourhood. The sound data were processed by the DSS, which adopts a “fuzzy logic” approach to provide more elaborated results (e.g. 15 minutes average and sound contribution analysis) and noise-related alerts as output to be visualised in the COP dashboard. Furthermore, the implementation of the Sound Heat Map module provided a wider and graphic estimation of the sound levels by using interpolation algorithms and geometric modelling of the event area.

Acknowledgements

This work has received funding from the European Union’s Horizon 2020 research and innovation programme as part of the “Management Of Networked IoT Wearables – Very Large-Scale Demonstration of Cultural Societal Applications” (MONICA) project under grant agreement No. 732350.

This work reflects only the authors' views and the commission is not responsible for any use that may be made of the information it contains.

References

- [1] High Level Architecture (HLA), Release 4.0, AIOTI WG03 – IoT Standardisation, online at: <https://aioti.eu/wp-content/uploads/2018/06/AIOTI-HLA-R4.0.7.1-Final.pdf>
- [2] OGC SensorThings Standard, online at: <https://www.opengeospatial.org/standards/sensorthings>
- [3] Grafana, online at: <https://grafana.com/grafana> (retrieved: 2019-09-05).
- [4] Github Repository: Grafana OGC Sensorthings Plugin, online at: <https://github.com/linksmart/grafana-sensorthings-datasource>
- [5] M. C. Bor, J. Vidler and U. Roedig, “LoRa for the Internet of Things”. In: EWSN, Vol. 16 (EWSN, 2016) pp. 361–366.
- [6] K. Haddad, P. Munoz, E. Gallo, B. Vincent, M-H Song, “Application of Internet of Things technology for sound monitoring during large-scale outdoor events”, Inter-Noise, Madrid, 2019.
- [7] R. E. Haber, et al., “Application of knowledge-based systems for supervision and control of machining processes”. In: Handbook of Software Engineering and Knowledge Engineering: Volume II: Emerging Technologies. 2002. pp. 673–709.
- [8] L. A. Zadeh, “Fuzzy sets”, Information and control, 1965, 8.3, pp. 338–353.
- [9] J. M. Mendel, “Fuzzy logic systems for engineering: a tutorial”, Proceedings of the IEEE, 1995, 83.3, pp. 345–377.
- [10] M. J. Skibniewski, et al., “Quantitative constructability analysis with a neuro-fuzzy knowledge-based multi-criterion decision support system”, Automation in Construction, 1999, 8.5, pp. 553–565.
- [11] IoT European Large-Scale Pilots Programme (IoT-LSPs), online at: <http://european-iot-pilots.eu/>

