Ecosystem for COllaborative Manufacturing PrOceSses – Intra- and Interfactory Integration and AutomaTION
(Grant Agreement No 723145)

D7.7 On-Site Readiness Assessment of Use Cases Based on Existing Sensor Infrastructure II

Date: 2019-02-06

Version 1.0

Published by the COMPOSITION Consortium

Dissemination Level: Public

Co-funded by the European Union's Horizon 2020 Framework Programme for Research and Innovation under Grant Agreement No 723145

Document control page

Document file: D7.7 On-Site Readiness Assessment of Use Cases Based on Existing Sensor Infrastructure II_V1.0.docx
Document version: 1.0
Document owner: TNI-UCC

Work package: WP7 – Integration of Internal and External Elements
Task: T7.3 – Analysis and Completion of Existing Sensor Infrastructure
Deliverable type: R

Document status: ☒ Approved by the document owner for internal review
☒ Approved for submission to the EC

Document history:

Version	Author(s)	Date	Summary of changes made
0.1	Peter Haigh (TNI-UCC),	2018-12-14	First version, for circulating to partners for inputs
0.2	Peter Haigh, Helene Udsen (IN-JET), Nikolaos Alexopoulos (CERTH), Alexandros Nizamis (CERTH), Christos Ntinas (CERTH), Theofilos Mastos (KLE), Angelos Papadopoulos (KLE)	2019-01-28	Inputs related to Asset Tracking and other edits to Conclusions, Exec Summary and Introduction Formatting and housekeeping inputs Inputs related to KLE and ELDIA deployment Inputs related to the benefits of the sensors' deployment at KLEEMANN's shop-floor
1.0	Helene Udsen, Peter Haigh	2019-02-06	Formatting, reviewer comments addressed. Final version submitted to the European Commission

Internal review history:

Reviewed by	Date	Summary of comments
Alexandros Nizamis (CERTH)	2019-01-29	The quality of the document is competent. Some improvements should be done in terms of text formatting as well as some minor changes in content and figures.
Dimitris Gkortzis (ELDIA)	2019-02-04	Clarification of ELDIA use case

Legal Notice

The information in this document is subject to change without notice.

The Members of the COMPOSITION Consortium make no warranty of any kind with regard to this document, including, but not limited to, the implied warranties of merchantability and fitness for a particular purpose. The Members of the COMPOSITION Consortium shall not be held liable for errors contained herein or direct, indirect, special, incidental or consequential damages in connection with the furnishing, performance, or use of this material.

Possible inaccuracies of information are under the responsibility of the project. This report reflects solely the views of its authors. The European Commission is not liable for any use that may be made of the information contained therein.

Index:

1 Executive Summary

This deliverable relates to WP7, Analysis and Completion of Existing Sensor Infrastructure Analysis and Completion of Existing Sensor Infrastructure, Task 7.3 Analysis and Completion of Existing Sensor Infrastructure. The first deliverable, *D7.6 On-Site Readiness Assessment of the Use Cases Based on Existing Sensor Infrastructure I* in Month 15 addressed the readiness of the 3 COMPOSITION industrial use case sites (BSL, KLE & ELDIA) for potential integration of wired and wireless sensors.

The key applications identified were asset tracking, conditional monitoring using acoustic and vibration sensors and container fill levels (scrap/metal).

Priority was given to identifying OTS (off the shelf) solutions and getting them to work and gather source data.

This deliverable documents the lab scale use cases that were deployed and identifies lessons learnt.

As the project evolved and technologies were evaluated greater understanding of capabilities and limitations were gained. In some cases, changes to the original development plan were adopted.

1.1 Asset Tracking

After evaluation of UWB it was found that whilst this potentially yields greatest accuracy, for this application the physical size and power consumption remains a significant challenge. To move these parameters to something that meets size and power requirements would require an iteration of the fundamental technology. Power consumption is a challenge that is being addressed in industry by companies such as Spark Microsystems (SM). BLE was the technology that was deployed as this met our requirements.

The BLE Asset tracking system was deployed in the white space area of BSL. Use of this area gave a realistic environment that is very close that experienced on the production floor. Being a prototype area, this gave the flexibility to experiment to a far greater degree.

At the time of writing this document testing is being undertaken in the White space at BSL

Whilst BLE is capable of tracking assets in the white space, future work may need to look at a fusion of technologies such as BLE, UWB and IMU to give location with high certainty and reliability.

1.2 Acoustic Sensing for Fan Condition Monitoring

A number of sensing techniques were evaluated including magnetics, current consumption and fan speed. Acoustic sensing was found to be the most sensitive and cost-effective technique to use. A prototype development platform was designed using commercial off the shelf components with wired Ethernet connection back to the COMPOSITION server where condition monitoring algorithms analysed the data for faulty fans.

One of the biggest challenges was reliability and stability of the system. This was due in part to the relatively harsh environment experienced by the sensors. An approach for future work here would be to implement with a bespoke industrial grade embedded micro wireless system.

1.2.1 UC-ELDIA-1 Fill Level Notification & UC-KLE-4 Scrap Metal Collection and Bidding Process

In both use-cases, two fill level sensors are installed in two containers on different sites. The fill level is monitored in both ELDIA's and KLEEMANN's bins and when it reaches a pre-defined threshold a bidding process starts for the sale or purchase of the scrap metal waste. This automated procedure is expected to reduce lead times and reduce scrap metal collection and bidding process costs. In ELDIA’s case the fill level notification serves as a warning that a container full of metal or paper or wood is reaching its maximum capacity, thus a substitution with a clean container should be arranged.

LoRa was chosen as the wireless protocol to use, as it has a very low power profile and one gateway can cover a radius of approximately 2km with outdoor devices. Furthermore, a LoRa sensor network was already deployed in CERTH at the time of development, so there was some experience about the behaviour of the protocol.

Because the bins were outdoor, IR sensor solutions were rejected as the direct sunlight exposure would create too much noise to the measurements. Interestingly, there aren’t many low power ultrasonic distance measurement solutions in the market and Maxbotix's ultrasonic sensors were proven the best choice as they also have a broad range of IP67 variants.

The biggest challenge was the deployment of the sensors in both ELDIA and KLEEMANN sites. The daily use of bins is very hostile to any third-party installation on them, as heavy waste is loaded and sometimes unloaded with cranes and forklifts. Consequently, there were many confines, due to safety reasons, that prohibited the free choice of a suitable placement that would contribute to a lower error rate and more precise measurements. Our input here is that in the future bins should be designed in a way that predicts and eases the installation of sensors on it.

1.3 UC-KLE-1 Predictive Maintenance

In this use case, two sensors are installed in two external motors of the BOSSI machine. A set of real-time vibration data is generated and analysed along with historical data, in order to provide early notification that the machine is close to a failure. Decision making regarding maintenance tasks is expected to be assisted and this will result in improvements in maintenance costs, down-times, process monitoring and quality of manufacturing.

A fill level sensor was designed and developed for the internal bins in order for the prediction engine to be able to estimate the fill level in the future and propose the optimal path to follow for the transportation of waste. A total of 14 sensors were installed on 5 special metal constructions manufactured by KLEEMANN.

A Time-of-Flight laser-ranging sensor was used, the communication was carried out via LoRa and the data were published on an MQTT topic on the cloud. Since the sensors are indoors, there may be a few package losses due to the limitations of LoRa in closed spaces.

The evaluation of the sensors took into account the limitations regarding the nature of the waste, environmental conditions and the limitations of the sensors themselves, such as the surface of the waste, the ambient light and the measurement distance respectively. Since the COMPOSITION scenarios are activated when the fill level percentage is 70-80%, those limitations don't cause any problems.

1.3.1 UC-KLE-3 Scrap Metal and Recyclable Waste Transportation

In this use case, twelve sensors are installed to the recycling and scrap metal bins inside KLEEMANN's factory to monitor fill levels. Real-time fill level notifications and suggestions for optimal collection routes will be provided. This will make the management of scrap metal and recyclable waste more efficient.

In this use case Wi-Fi was the chosen protocol for two main reasons. One reason was that there was already Wi-Fi coverage in KLEEMANN site. The other was the big volume of data required for training the prediction model.

The sensor used was a LIS3DH 3-axis MEMS accelerometer and was chosen primarily because of the high sampling rate at a low cost and small footprint, compared with other MEMS accelerometers.

The sensors create a lot of data per second, so scaling can prove a big challenge for the back-end infrastructure and network. Finally, due to the nature of the accelerometer orientation of the sensor matters during installation if we want to use multiple sensors with the same thresholds and compare their output.

2 Abbreviations and Acronyms

Table 1: Abbreviations and Acronyms

Acronym	Meaning
BLE	Bluetooth Low Energy
CMMS	Computerised Maintenance Management System
COTS	Commercial Off The Shelf
dBSPL	Column 2 text
GPS	Geostationary Positioning System
IIMS	Integrated Information Management system
IMU	Inertial Movement Unit
LoRa	Long Range (Proprietary low power wireless system from Semtech)
MCU	Microcontroller unit
MEMS	Micro Electro-Mechanical Systems
MQTT	Message Queuing Telemetry Transport
PC	Personal Computer
PCM	Pulse Code Modulation
RTLS	Real Time Location System
UWB	Ultra Wideband
WAV	Waveform Audio File

3 Introduction

3.1 Purpose, Context and Scope

The deliverable follows on from *D7.6 On-site readiness assessment of use cases based on existing sensor infrastructure I*. D7.6 was a description and assessment of planned use case deployments where as D7.7 describes the use cases that were deployed as part of the COMPOSITION project and goes on to discuss some of the lessons learnt.

The document gives a detailed description of the design and implementation of each deployment, showing at a system level the sensed data path from sensor to COMPOSITION server. In addition, for each use case lessons learnt from the deployment are presented.

3.2 Content and Structure

This document is structured as follows. Section 4 describes the BSL factory use cases that requires sensed data. The first sub section is a brief description of the work completed within D7.6 to remind the reader of some of the key pieces of work undertaken at the early stages of the project. Later sections go on to describe in some detail the design implementation and deployment of the two surviving BSL use cases. Section 5 repeats the same structure for the use cases at KLE and ELDIA. Section 6 completes the document with conclusions.

4 BSL Factory Use Cases

4.1 Review of Work Completed in D7.6

D7.6 describes the work undertaken to identify the use case, identify the requirements and the down selection of an appropriate solution. This section gives an overview of D7.6 to give the reader an understanding of the evolution of the project that got us to the final design and deployment described in sections 0 and 0 of this document.

4.1.1 BSL Use Cases

D7.6 originally identified three use cases as described below

1. *UC-BSL-2 Predictive Maintenance - Fan monitoring system* Use of sensors that can 'listen' and monitor performance (temperature, vibrations, power consumption) on and near fans (blowers) in reflow ovens. The 'signature data' from these can give early indication that a fan will fail in the near future.
2. *UC-BSL-3 Component Tracking - Asset tracking system.* Use of wireless sensors that can be attached to component reels, fixtures, jigs, sub-assembly trays, etc. so their location can be determined within a factory. The direct (material value) losses and indirect (time lost in production) losses in a factory can be considerable if such assets cannot be found quickly.
3. *UC-BSL-5 Equipment Monitoring and Line Visualization* - Extract data from existing factory systems and feed to partners to help develop systems that predict/detect production bottlenecks monitor and visualize general workflow, detect anomalies, etc. There is an ongoing requirement to show actual vs target production on the line. In this instance it is predominantly used to monitor a light tower in order to track the status of each piece of equipment. There was no foreseen need for additional sensors to be installed so it was not covered in this deliverable.

4.1.2 UC-BSL-2 Condition Monitoring D7.6 Review

There can be up to €20,000 worth of material running through a flow solder oven. If the oven fails material that was run through the machine has to be thrown away. Therefore, being able to detect faults before or as they occur yields a significant cost saving. Whilst there is some diagnostic sensing built into the fan such as temperature and fan speed this is inadequate for predictive maintenance where often a faulty condition is detected by the operator noticing additional noise coming from the fans. This noise often increases before there is a noticeable effect on the flow solder profile. Section 4.3.2 in D7.6 provides further details of the four oven types used in BSL today.

4.1.3 UC-BSL-3 Component Tracking D7.6 Review

Knowing where material is located within the factory in real time can yield significant cost savings in elimination of lost material and improved process flows.

D7.6 described the key requirements for asset tracking. In general material was to be tracked around the factory floor so that its location could be detected in real time. One of the key concerns was that the RF energy used by the tracking system should not interfere with those frequencies used for product, which is typically in the 433MHz band. Refer to section 4.3.1 in D7.6 for a full description of the requirements.

4.2 UC-BSL-2 Condition Monitoring

A number of different techniques were investigated to determine the best detection method. These included acoustic, electromagnetic, vibration and current consumption. The conclusion from this work was that acoustic sensing was the most appropriate method as one could see a clear distinction between a good and a faulty fan. Refer to D7.6 section 4.5.2 for more detailed information.

4.2.1 Design, Implementation and Deployment

The condition monitoring system is designed to detect fans that are about to go faulty before they affect the temperature profile of the oven. There are over 30 fans in the oven and utilisation of acoustic sensors means that we can detect zones for the condition of the fans to simplify the system and reduce costs.

4.2.1.1 Hardware Architecture

Figure 1 below gives an overview of the hardware architecture. The acoustic sensors are constructed from raspberry Pi interfacing with a SPH0645LM4H microphone via an I2C interface. The sensors record 20 seconds of sound every 5 minutes in a standard PCM format. This data is sent in an uncompressed WAV format to the Manufacturing floor PC via Ethernet. The WAV files are processed on the PC where each 20 second recording is converted into a single amplitude value. The WAV file is stored locally on a USB drive and the single amplitude value is stored on the COMPOSITION server remote from the factory.

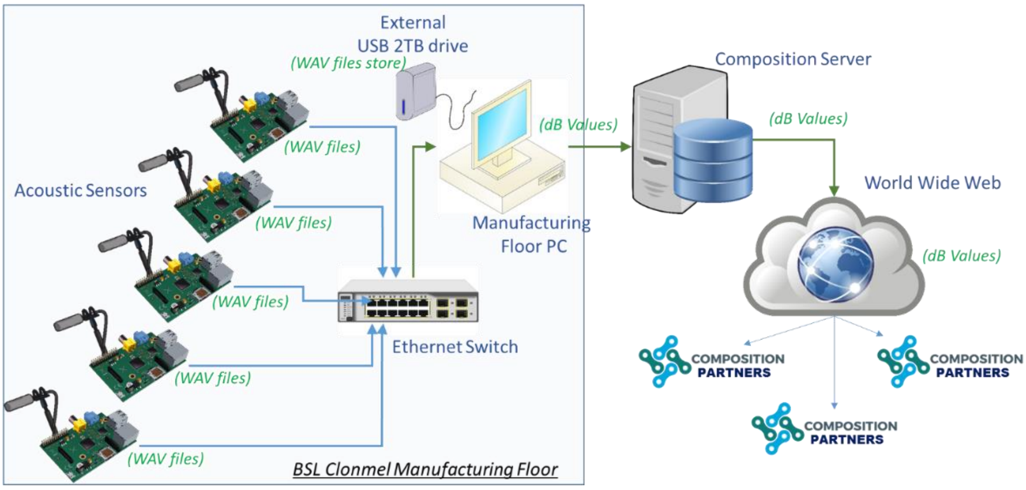

Figure 1: UC-BSL-1 Condition Monitoring Architecture

Each WAV file is 3.7MB in size the algorithms however only require to know the average amplitude in dBSPL this is calculated by taking the mean of peak envelope filtered data. The original WAV files are stored on a 2TB USB hard drive which would store 1 year of data from the 5 sensors.

4.2.1.2 Software Architecture

The software architecture consists of a number of elements that are distributed between the Raspberry PI acoustic sensor and COMPOSITION PC as described below in Figure 2. The Raspberry PI software records 20 seconds of sound every 5 minutes via a script that runs in a continuous loop. The data is saved in a WAV file format and stored on the local SD card. The 10th oldest file is deleted to prevent the SD card from overflowing.

The software on the COMPOSITION PC also runs in a continuous loop every 5 minutes that is controlled via the Windows task scheduler. The software takes the newest WAV file from each of the five Raspberry PIs, passes the data through a peak envelope filter and then takes the mean value to produce a single dB amplitude value. This is placed in a specific location on the PC where network level software then moves it to the COMPOSITION server.

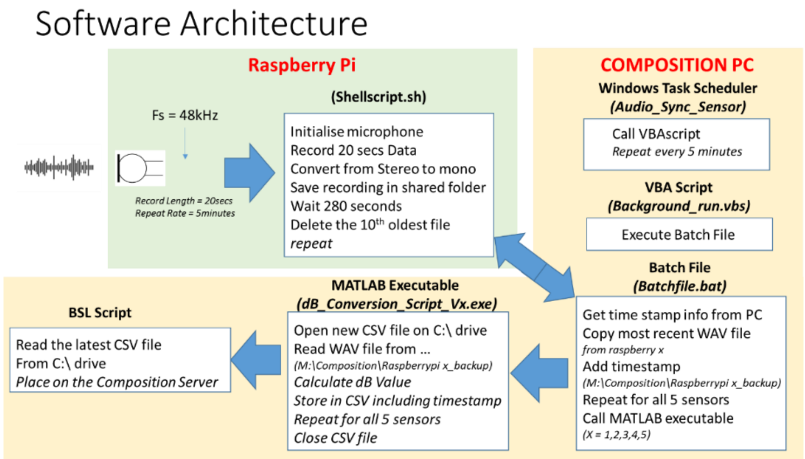

Figure 2: UC-BSL-1 Software Architecture

The shellscript.sh program is located in /home/pi directory on the raspberry PI and is the program that controls the recording of the audio, stereo to Mono conversion, recording duty cycle and storage and time-stamping of the 10 most recent recordings.

The batch file is located on the manufacturing PC. This file reads the latest WAV file from each of the Raspberry Pi sensors that are all mapped to a unique drive on the manufacturing PC. It then stores these files on the USB 2TB drive with a specific time-stamped file format that is in alignment with the Manufacturing PC time. The file then calls the MATLAB script that calculates the dB amplitude value.

The Matlab script calculates the average dB value of the 20 seconds of audio. This calculation includes peak envelope filtering as well as truncation of the data at either end to avoid any harsh transitions at switch on and switch off being included in the calculation.

For each set of 5 measurements a single CSV file is generated that stores the 5 values from each sensor. The timestamp in the filename is identical to that of the WAV files and so cross correlation is possible at a later date if required.

4.2.2 Additional Data

In addition to the data collected from the acoustic sensors, data is also collected from the inbuilt monitoring system of the flow solder oven and sent to the COMPOSITION server via the factory floor PC. The following information is contained within two sets of files.

The event file contains a list of events such as alarms, warnings, operator interventions and fault conditions.

The log file contains Set Temperature, Actual Temperature and Output Power.

4.3 UC-BSL-3 Asset Tracking

A number of technologies were reviewed. Three technologies BLE, UWB GPS underwent Lab evaluation. It was found that GPS did not lock properly to the radio signals within an indoor environment and so was unsuitable. UWB is able to give accurate positioning indoors with its wide bandwidth enabling multi path mitigation to enhance accuracy. However, this technology is still immature and its size and power requirements meant it was not suitable for battery operation on a small component tray. Section 4.5.2 and 4.6.1 in D7.6 discusses this selection further.

4.3.1 Design, Implementation and Deployment

4.3.1.1 Hardware Architecture

Tracking will be implemented using the Link-Labs AirFinder Bluetooth Low Energy tracking kit which employs a proximity sensing method. Tags listen out for beacons which continuously advertise their signature over BLE. By listening to and comparing the strength of these signals a tag decides which it is closest in proximity to. It then relays that information to an Internet Access Point which talks to the cloud based application, updating the known location of tags as they move from area to area. The following diagram shows an example of a tag entering the proximity of 'Location A' and sensing the signature of the beacon that defines this area causing a position update event.

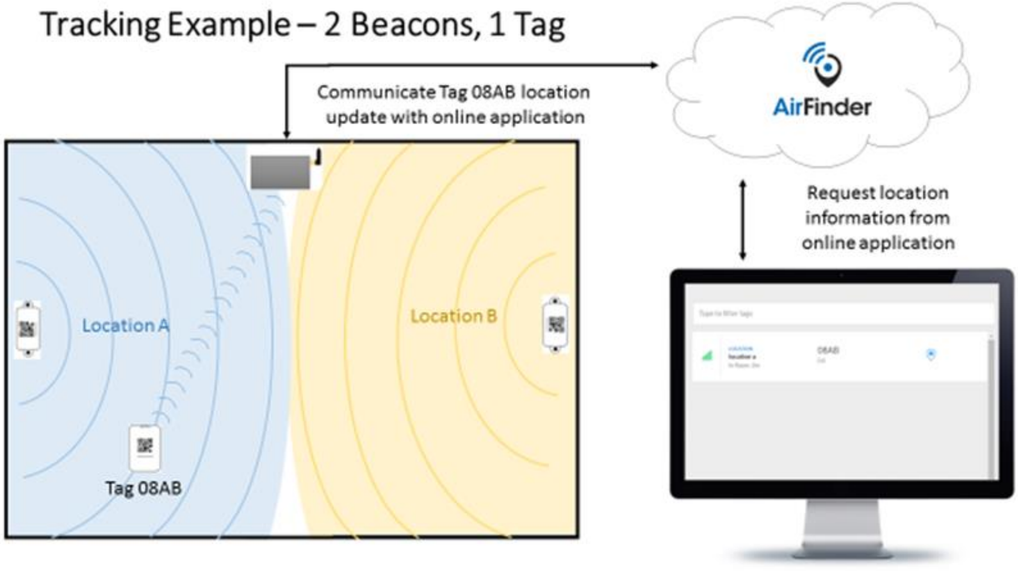

Figure 3: Air Finder Configuration Tracking Example

In the diagram the three components of the tracking system can be seen; a Tag, Beacons, and an Access Point. The following diagram shows the two different styles and dimensions of tags that can be used with this system.

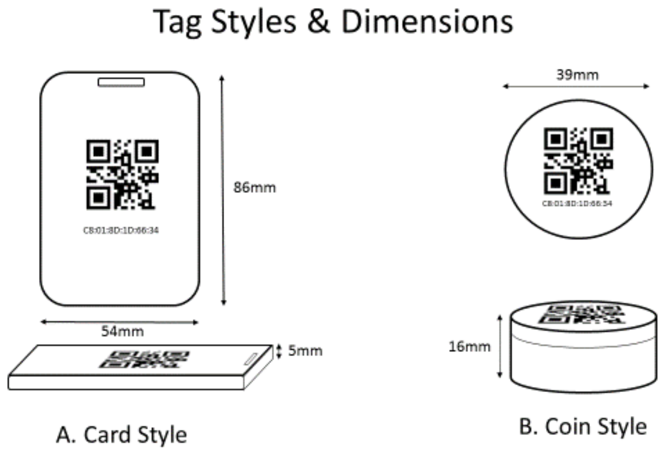

Figure 4: Tag Styles

The Air Finder tags are attached to assets that it would be desirable to track. Different forms may suit different assets depending on the form factor and conditions faced by the asset. They will be usually attached by strong adhesive tape for long term tracking, and either Velcro strips or zip-ties for shorter term tracking. Again, this depends on the needs of the specific use cases.

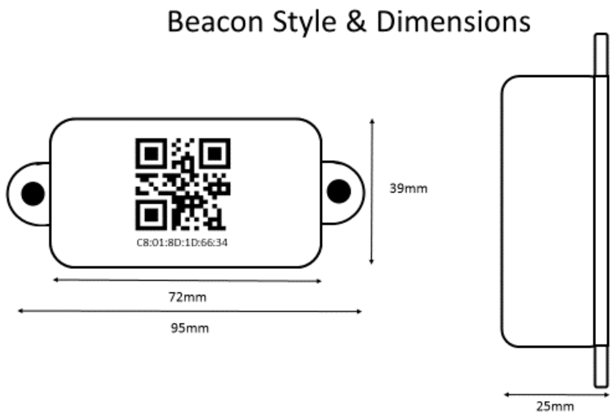

Figure 5: Beacon Dimensions

Beacons are attached at fixed points in the area you want to track in. Since the number of these defines the number of locations you have in the tracking area intuitively the higher the beacon density you have the better the potential accuracy of the tracking. Although, this only true up to a certain density where you introduce more areas of uncertainty which counteracts the effect of increasing accuracy. They can be attached at any fixed position through the 2 loops at either end of the packaging, through screws, zip-ties, etc. or with adhesive tape if either of those are not possible. Determining the optimal placement of the beacons is demonstrated later in this document.

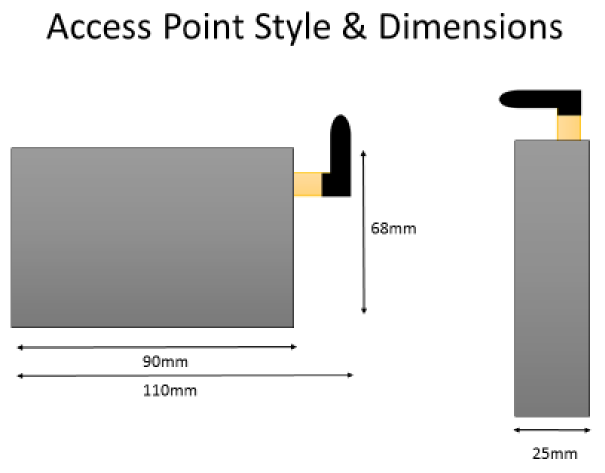

Figure 6: Access Point Dimensions

The Access Points are placed around the tracking area to communicate with the tags and also have a limited range (200 - 400m^2), so some sort of standard layout needed to be determined. These devices are powered via a micro-USB port, and also need to be connected to the internet using an Ethernet cable. This needs to be considered in the placement.

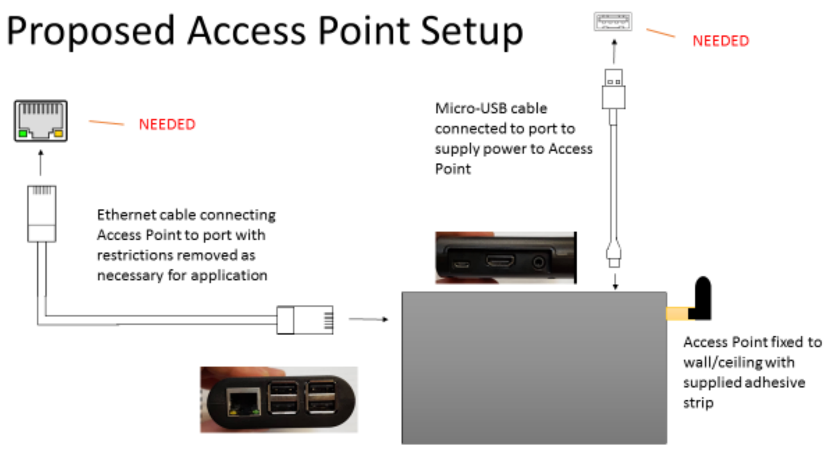

Figure 7: Access Point Requirements

4.3.1.2 Deployment

Since the technology we are using is proximity based we want a layout which minimizes the maximum distance a tag could be from any proximity beacon for a given floor space and number of tags. We also want to make sure that Access Point connectivity is available everywhere in the tracking area. We have been given the Rapid Prototyping room in BSL which has regular dimensions of 20 m x 10m.

We have a desired accuracy to track to in this space (roughly 2.5-3m) so we need to figure out the optimal layout which will minimize the number of beacons needed. If the room is pictured from above it can be imagined that the beacons' proximity are circles with radius 2.5m. So what we are trying to do is cover a 20m x 10m rectangle using the minimum amount of 2.5m circles. If a standard approach is taken where the beacons are set up in a square grid, we would get a packing that looks like this for a given area:

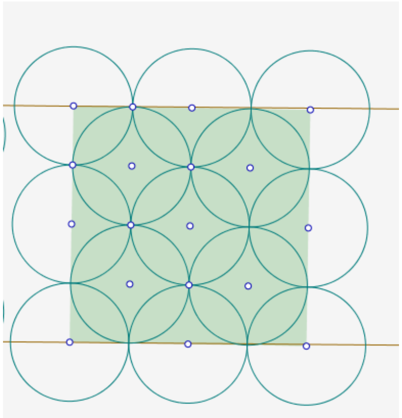

Figure 8. Tiling - Square Grid

What is obvious here is the large amount of overlap between the circles indicating inefficient packing. And this results in 13 circles being needed to completely cover an area this size. A clever approach would be to use hexagonal packing to cover the area in a more efficient way. The centres of these circles are placed so that if you were to draw the boundaries between each it would result in a grid of regular hexagons.

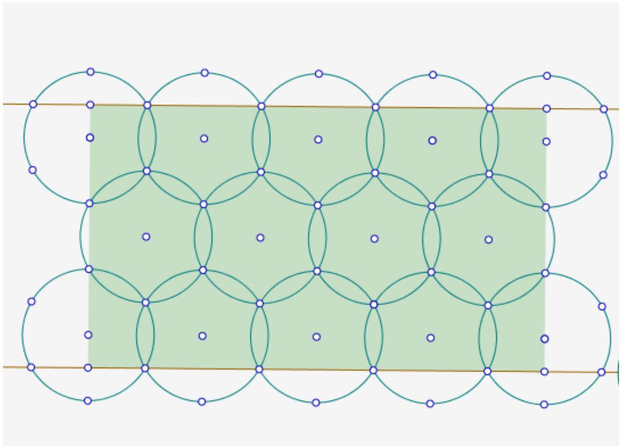

Figure 9: Tiling - Hexagonal Grid

Here we can see how 14 circles used in this fashion can cover twice the area shown in the first example where a square grid was used. This is the best regular pattern that can be used in this case and so a template was drawn up. We needed to determine the size of these hexagonal areas which would best suit the size room we are working with. Aiming for about 2.5m accuracy it made sense to tile 3 hexagons across so that they would cover the width of the room (10m). With this setup the grid makes sure that a tag is never more than 2.57m from a beacon anywhere in the room.

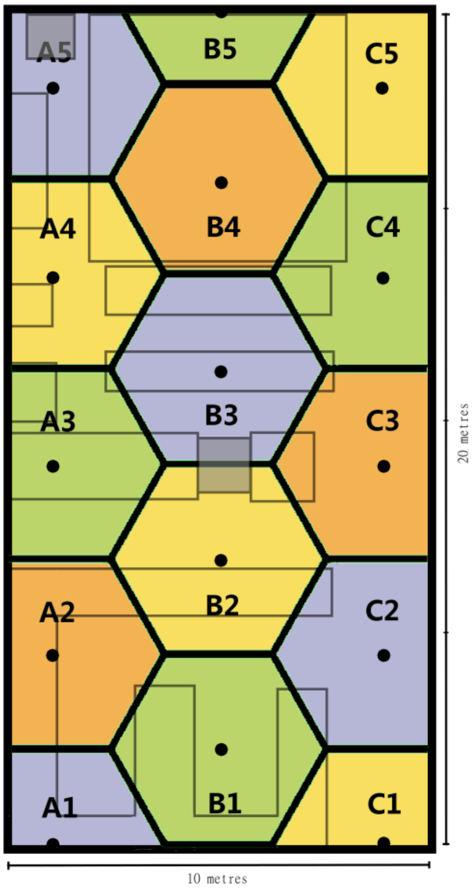

Figure 10: BSL Whitespace Map & Beacon Placement

This shows an overhead map of the BSL Whitespace where desks and other obstacles are outlined and a hexagonal grid of beacon proximities has been overlaid. It has been decided that locations for this prototype will be ordered in rows and columns, the first alphabetic character denotes the column from west to east and the number that follows denotes the row from south to north. This is a first pass design choice, this will be revisited following deployment as a better way of naming locations that is more intuitive may be discovered through long term use.

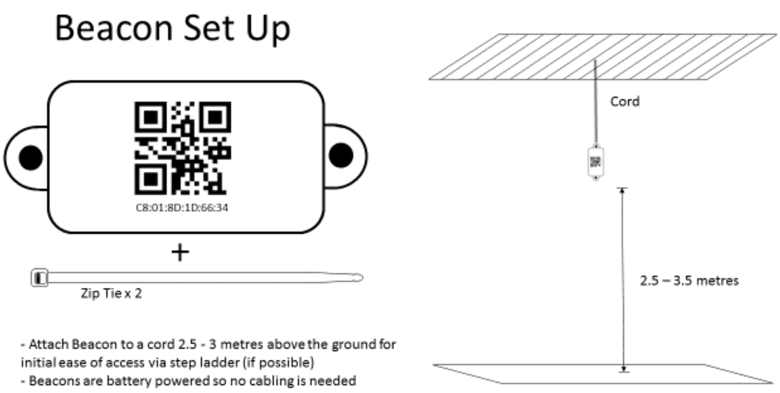

Figure 11: Beacon Set-Up

In the beacon placement grid, we are only looking at a two-dimensional picture of the tracking space where the height aspect is neglected. Taking into account the vertical dimension we notice that if the beacons are placed on the ceiling, then the closer you are to the ground the relative difference in signal strengths between 2 neighbouring beacons around the border becomes less and less. We need this difference to be great to allow less uncertainty when a tag senses which beacon it is proximate to. Using Friis (Friis) equation we estimate that the area of uncertainty between two beacons will be about 3 times as large at the floor compared to the ceiling. This is a problem but we need to attach the beacons to the ceiling in order to avoid obstacles, so the ideal solution is to hang the beacons down from the ceiling to a certain extent. To do this we will use a rubber cord to lower the beacons 3m from the ground. The tags should mostly be attached to objects on tables approximately 1m above the floor. This in theory should reduce the uncertainty areas by half as well as allowing easier access to the beacons during initial deployment testing.

Figure 12: Beacon Set-Up in Whitespace

4.3.1.3 Software and Alerts

The AirFinder application is cloud based and hosted by Amazon Web Services, which means while you can easily access the tracking information from anywhere, although the system needs to be constantly online otherwise you lose visibility of the movement of tags. Because of this we have decided to make use of the Ethernet port to provide internet access. The following data flow map best represents how the system communicates with the application and how the tracking data is accessed:

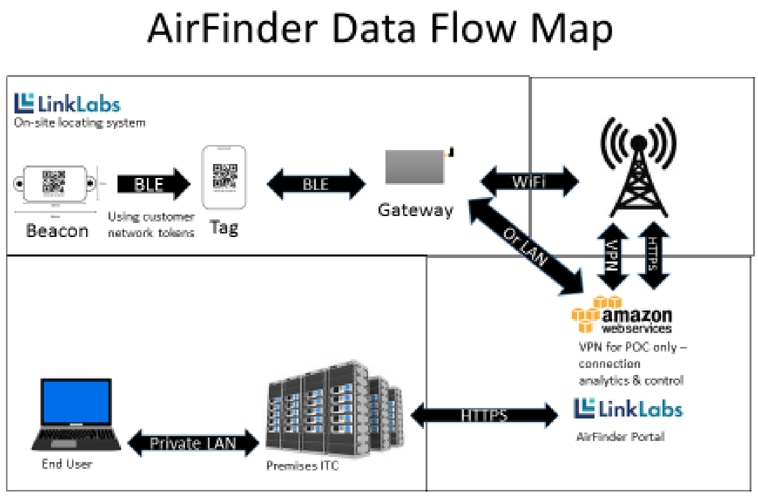

Figure 13: AirFinder Data Flow

The application itself is accessed via browser at app.airfinder.com. Each person who needs access to it may be added as either an admin or user with privileges dictated by the admin setting up the account. These privileges include the ability to add new locations, add tags, download reports, or set/receive activity alerts. The first step to set up the application is to add a new site which would be the facility in which you are setting up. Within this you divide the site into areas and then again divide the areas into locations or zones both of which are defined by the beacons. Locations are defined by single beacons whereas a zone may be defined by multiple beacons. This is useful in some cases where you only care about whether a tracked asset is in a general area rather than somewhere specifically.

Once the site you want to perform tracking in has been divided up into its respective areas, zones and locations and beacons have been assigned you may begin to add tags. The way this application allows classification of tracked assets is by group and category. Because these are loose ways of classifying assets, we will have to define how these tags are used. For our purposes we will say that groups will be general asset descriptors and categories will be more specific. So if an oscilloscope which is used for testing is tracked, we would call the group 'Test Equipment' and the category 'Oscilloscopes'. There is also an individual asset description with each tag which may be appended where more specific information can be appended if necessary. Adding a tag is as simple as entering the unique hex ID of the tag to be attached as well as all the relevant information about the asset into the application. Once entered the tag is tracked as long as it remains in the database.

The important aspect here is the physical security implications and there are many features here which are so in this case. The most important is the ability to send out alerts to relevant personnel when tags leave/enter zones, areas or locations they are not meant to. This is entirely configurable and the alerts can be sent via SMS or email. As well as this, full reports may be generated of the movement history of any tag from when it was initialised to the present.

4.3.2 Energy Harvesting Implementation

It has been desired since the start of the project to implement energy harvesting in as many aspects of this system as possible. Investigations were carried out to determine the power consumption of each piece of hardware in the tracking system. As well as this we needed to determine the available sources and availability of ambient energy which could be captured in each case, compare the two and determine the viability of implementing energy harvesting to increase the life of or replace the battery. The following is a table of available energy sources in BSL's factory floor and the power they could generate.

Table 2: Ambient Energy and Availability

Energy Source	Photovoltaic	Thermoelectric	Piezoelectric	RF
Expected Power	10-30 µW/cm^2 400 - 1000Lux	50-100 µW/cm^2 per °C	10-200 µW/cm^3	<1 µW/cm^2
Availability	600 – 1200 Lux 14 hours per weekday	Some machinery produces temperature differential when active	No consistent source of Hz – kHz vibration	Ambient RF is minimised to prevent interference with production

From the above table we can see that we can eliminate vibrational and RF energy harvesting as the sources are not present in the environment in a way that would produce consistent power. Between photovoltaic and thermal we have ways of generating decent amounts of power from the surroundings. There is quite high intensity indoor lighting present since this is a production environment where compared to office conditions brighter lighting is mostly preferred due to improved alertness, safety and productivity of workers. Of the many machines that cover the factory floor there are plenty that produce a temperature differential although they are located in specific areas. In the end it was decided to go with photovoltaic harvesting due to the greater coverage, potential energy and relative maturity of the technology.

The two pieces of the tracking kit that have a power consumption figure within range of what would be considered energy harvesting compatible would be the tags and the beacons. The main power draw in each component is caused by receiving and transmitting Bluetooth messages. The following graphs show the current draw from a 3 Volt source during operation.

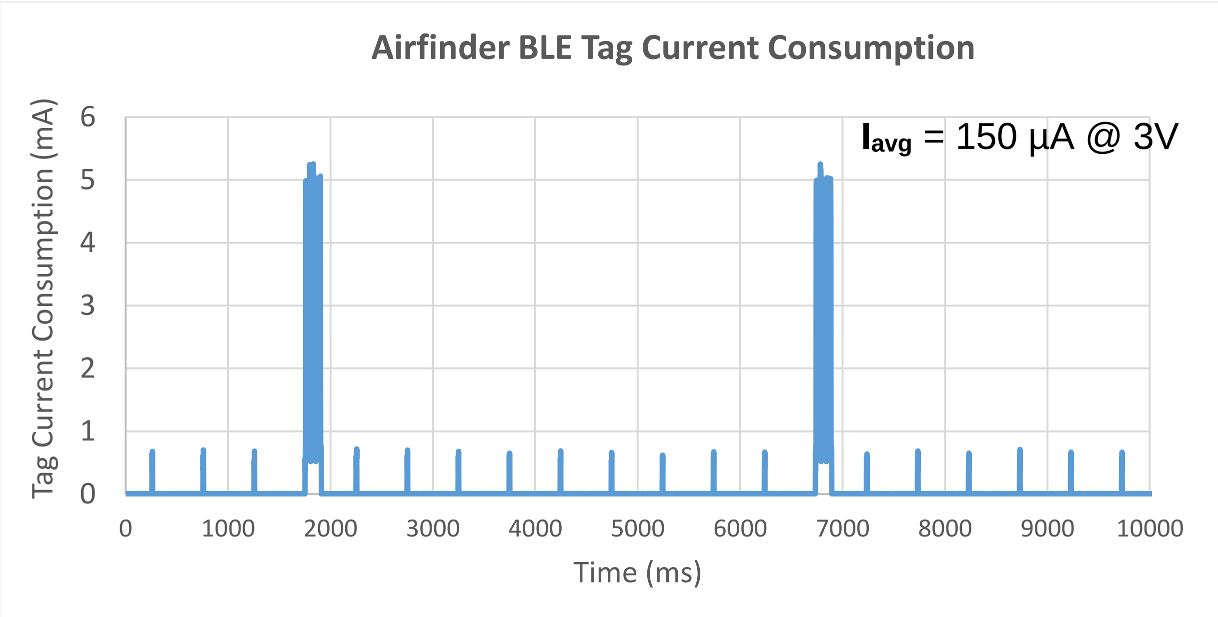

Figure 14: Tag Current Consumption

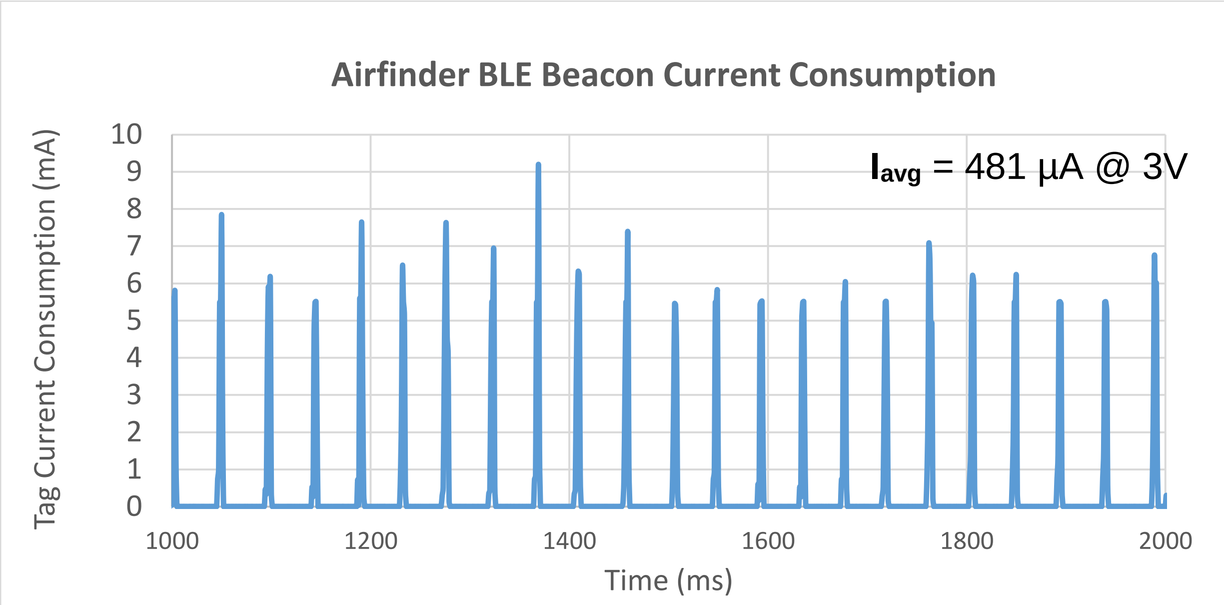

Figure 15: Beacon Current Consumption

The tag consumes on average 450µW which is well within the bounds for what may be harvested but the issue arises that since tags are intended to be mobile, we cannot guarantee that tags will be exposed to enough light or any other source of ambient energy to replace or compliment the battery as a power source. On the other hand, beacons are meant to be fixed in position and while the power consumption is 3 times as high (1.44mW) we can provide a consistent source of harvested power. The light conditions in the BSL factory floor and whitespace are very good for this, averaging about 1200 Lux at the beacon placement points. Using an AM1816 PV panel which measures 50cm^2 (53.3 x 94.7mm) we can generate up to 1.46mW in full lighting conditions. This figure is based on experimental data using the previously mentioned PV panel with our own energy harvesting power management board, and considers efficiency of the power management board, MPPT (maximum power point tracking) accuracy, storage and dc/dc conversion losses, etc. Although this can power the device completely during lighting hours, we still want the device to work while the lights have been switched off. In this we still need batteries to provide power during these times, and so the harvesting is to be used to complement and extend the life of the battery.

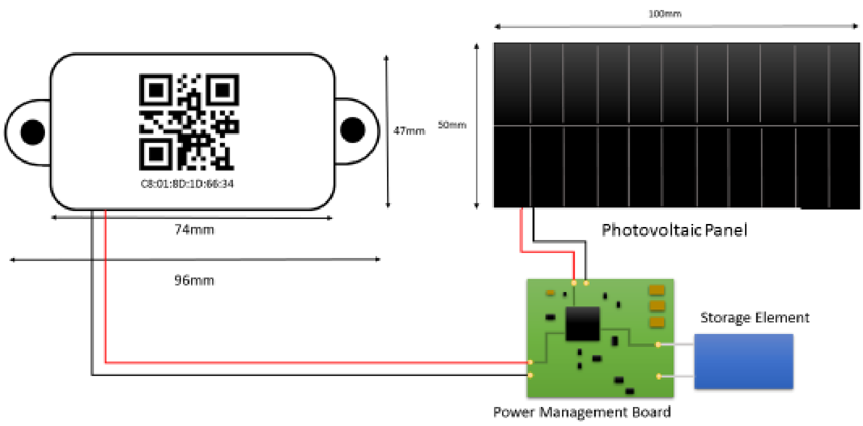

Figure 16: Proposed Beacon Setup

It was necessary to make predictions about the performance of this device over a long period of time. A script was written in MATLAB which modeled the system's behaviour over weeks at a time for different setups. Based on experimentally determined models of each component, when the script is run it outputs the utilized harvested power, supercap and battery state of charge over a date range, expected battery lifetime, and what contibution energy harvesting has made to this. The following plots are split into 3 sub-plots the first showing the power consumption profile of the end device. The second shows the prediced voltage level of the supercap (directly related to how much energy is stored) in blue and the red line indicates when energy harvested power is being used. The final plot shows the backup battery state of charge, in this case 2 AA batteries in series. The first plot shows the performance of an energy harvested beacon over a 3 week range assuming 1200 Lux, 14 hours of light a day only on weekdays.

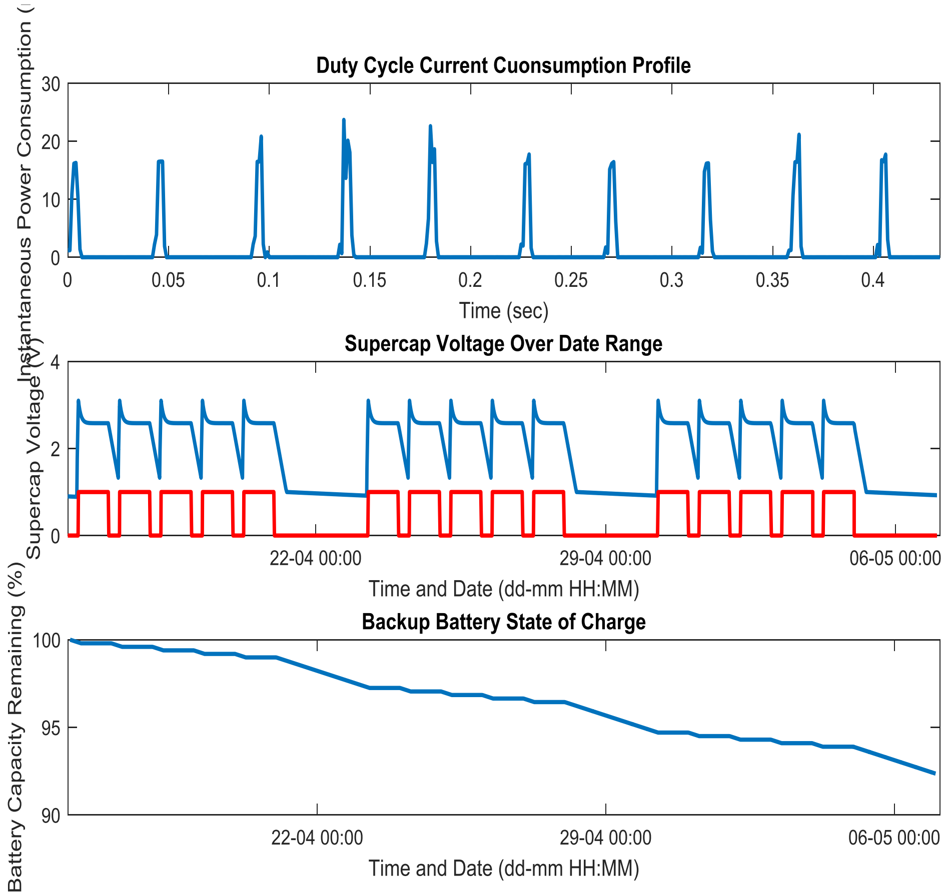

Figure 17: 3 Week Cycle at 1200 Lux

It can be seen during the weekdays the supercap quickly charges and starts providing power to the beacon. There is a small surplus of energy which starts being used when the lights go out but this is quickly drained and power is once again supplied by the battery. For this setup we can expect energy harvesting to provide power 52.5% of the time. This extends the battery lifetime by 110% from 260 days to 550 days.

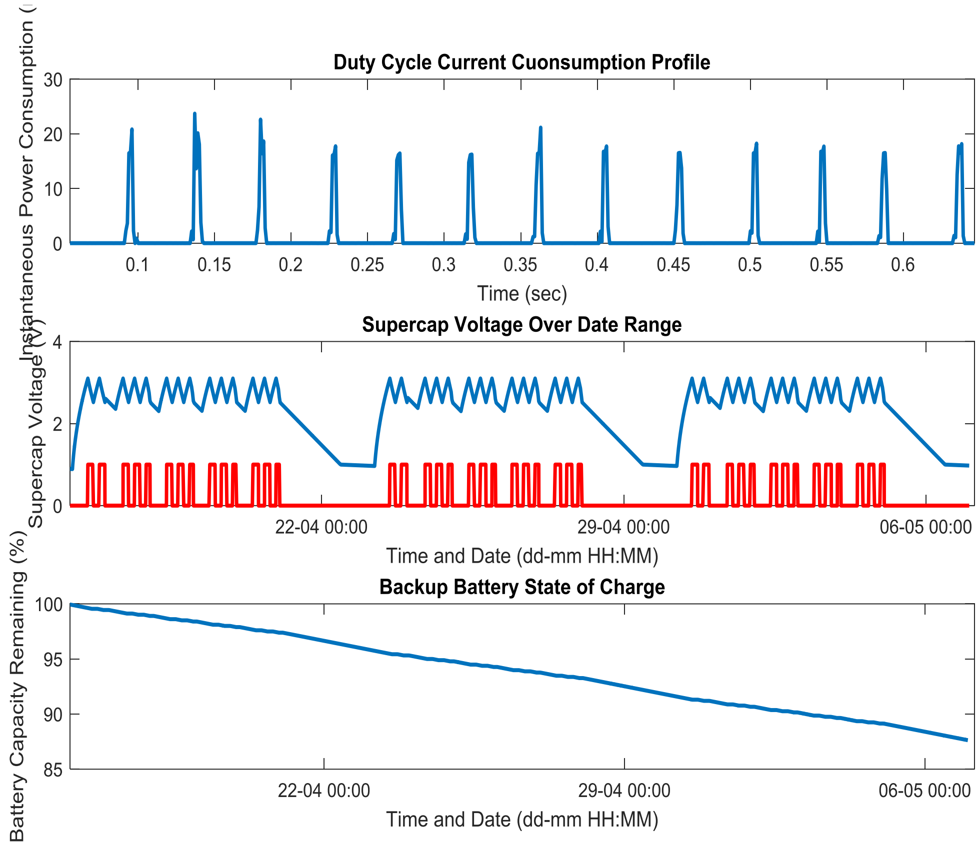

Figure 18: 3 Week Cycle at 700 Lux

The previous simulation was carried out again assuming a light intensity of 700 Lux which would be a worst-case scenario in terms of light level seen by the beacons. In this case the power which can be generated is 850µW which is only a fraction needed to run the beacon. Here instead the supercap charges and discharges several times over a single day providing on average 25.2% of the power consumed and extending the battery life by 33% from 260 to 348 days.

4.4 Lessons Learnt with UC-BSL-2 and UC-BSL-3

During the course of the project there are a number of technical challenges encountered which is part of the design process. Some high level lessons learnt are presented below which focus on events that may not necessarily be considered or generally known

4.4.1 Diagnostics and Housekeeping Systems

One of the key aspects of condition monitoring is the need for reliable data without interruption. The learning algorithms are greatly diminished if data gaps appear (as well as the risk of missing an event during down times). Over the course of the project it became clear that there needed to be an automated system that monitors the health of the sensors and gave automatic notification and alerts. In addition to this a process is required where by the alarm is acted upon and the issue fixed.

4.4.2 Condition Monitoring Algorithms Requires Reliable Data with no Gaps

To extend on the lesson learnt above, for these types of applications one may need to consider reliability and yield as part of the design to increase the level of robustness in the design. COTS components were used in the design which gives a certain level of robustness, which is usually appropriate for most sensing applications where occasional interrupts to data flow does not significantly affect the performance of the system. In this type of application un-interrupted data flow is very important.

4.4.3 The Use of Wireless in a Regulated Factory Environment

Although Wireless systems are pervasive today, this is something that is tightly controlled within the manufacturing floor at Boston Scientific. Some of the medical devices manufactured in Clonmel have very sensitive radio receivers and any electromagnetic interference can lead to production yield issues. These emission requirements are in excess of standard ETSI or FCC radio regulatory requirements.

4.4.4 Mechanical Deployment Issues Can be Challenging

The asset tracking system requires a number of beacons in fixed locations. In a real factory environment, a number of factors come in to play that may mean the beacons cannot be located where they are needed. One example is that we had to find a solution to the very high ceiling on deployment. In regulated factories a number of checks need to be undertaken before permissions can be sought. The lesson here is that this needs to be considered at the earliest stage of development.

5 KLE/ELDIA Factory Use Cases

The work in D7.6 identified three use cases UC-ELDIA-1 and 2 Fill Level notification, UC-KLE-1 Maintenance Decision Support and UC-KLE-4 Scrap metal collection and bidding process.

5.1 Review of Work Completed in D7.6

5.1.1 UC-ELDIA-1 Fill Level Notification & UC-KLE-4 Scrap Metal Collection and Bidding Process

Although it was stated in D7.6 that no sensor deployment will take place in ELDIA, one sensor was deployed on an ELDIA bin in ELDIA premises. KLEEMANN safety department advised us against on installing a sensor on one of the two bins, as the unloading process of that bin could easily destroy any installation we would attempt. In order to have a second deployment to compare results and sensor behaviour we chose to make one in ELDIA site. ELDIA also wanted to have a device in their premises in order to examine it and evaluate its performance first hand. Subsequently, one sensor was deployed on a bin at KLEEMANN and another on a bin at ELDIA.

After the initial deployment the sensor stated in D7.6 (HRLV - MaxSonar - EZ™ Series MB1013) was swapped with its IP67 variant (HRXL-MaxSonar MB 7380), when viability of an ultrasonic solution was ensured. The IP67 sensor proved to be more precise in outdoor conditions and it is certain that is going to have bigger life expectancy.

The LoRa gateways were integrated in KLEEMANN and ELDIA networks without problems and the connection with the LoRa Server is constant. This is the LoRa infrastructure required for the setup of a LoRa network as stated in D7.6.

5.1.2 UC-KLE-1 Predictive Maintenance

In D7.6 it was stated that two sensors would be installed on two Bossi motors, one on an internal motor and one on an external motor. During the development and installation process of the external sensor we were consulted by KLEEMANN safety department on safety issues concerning the installation of the internal sensor. It was pointed to us that the environment internally is very hostile, even for a sensor enclosed in a steel case the life expectancy of the device would be less than a month. In order to test the system behavior with more than one sensor and to compare behavior between sensors we added the second sensor on an external motor.

Finally, the KLEEMANN IT department provided us with an SSID and integrated the sensors in their network providing them connectivity with the Internet.

The installed sensors on the BOSSI machine are providing a set of vibration data, which together with the data from CMMS, are analysed in order to provide an early indication that the outside motors are close to a future breakdown. This notification is expected to assist the decision making of the maintenance planner and maintenance manager and help them better organise maintenance tasks. Overall it is expected to reduce the overall down-time from failures, to provide cost savings from improved process monitoring, to reduce cycle-times from process monitoring, to reduce scrap and repair costs and to improve the quality of manufacturing.

5.1.3 UC-KLE-3 Scrap Metal and Recyclable Waste Transportation

In D7.6 it was stated that the Time-of-Flight VL53L0X sensor would be used to measure the fill level of the recycling bins since its ranging capabilities were suitable for the task. For that reason, 12 sensors were used to monitor such bins, with the addition of 2 more to monitor 2 scrap metal bins. The LoRa infrastructure used for the UC-KLE-4 is used for this use case too.

The installed sensors in UC-KLE-3 will provide early (real-time) notification of the recyclable and scrap metal bins fill levels and suggest optimal routes for collecting bins within the factory. Overall, minimization of the total distance from bins to container and improvements in containers' fill level management are expected.

5.2 UC-ELIDIA-1 Fill Level Notification

5.2.1 Design, Implementation and Deployment

UC-ELDIA-1 and UC-KLE-4 share a sensor front-end of two fill level sensors deployed on two bins on separate sites that measure their fill percentage. UC-ELDIA-1 includes the distant monitoring of both industrial bins through the sensor network. UC-KLE-4 refers to KLEEMANN scrap metal bin where the sensor is monitoring the fill percentage and when it exceeds a pre-determined threshold a bidding process is triggered for the sale of the scrap metal.

The installed sensors in UC-KLE 4 are used for the automated notification of the scrap metal container fill level. This is expected to result in the optimization of the transportation of scrap metal and the selection of the best offer from the prospective bidders. Overall, the benefits from this use-case are expected to be the minimization of scrap metal collection and bidding process costs, the reduction in lead times, the improvements in receiving fast, efficient and high-quality services.

5.2.1.1 Hardware Architecture

Figure 19 provides the system architecture for UC-ELDIA-1, UC-KLE-4 and UC-KLE 3. Those use cases form a LoRa sensor network that is displayed below. On KLEEMANN and ELDIA sites various types of distance sensors have been deployed to different types of bins. For UC-ELDIA-1 and UC-KLE-4 two devices were developed that monitor two outdoor container-like scrap metal bins, one in KLEEMANN site and one in ELDIA site. The devices that were deployed consist of STM53L0 MCU controlling an IP-67 HRXL-MaxSonar MB 7380 ultrasonic distance sensor and an sx1272MB2xAS LoRa communication module. Special attention was given to the casing as the scrap metal can damage the device during loading or unloading. For this reason, a steel case was manufactured by KLEEMANN to enclose the devices and interface to the bins.

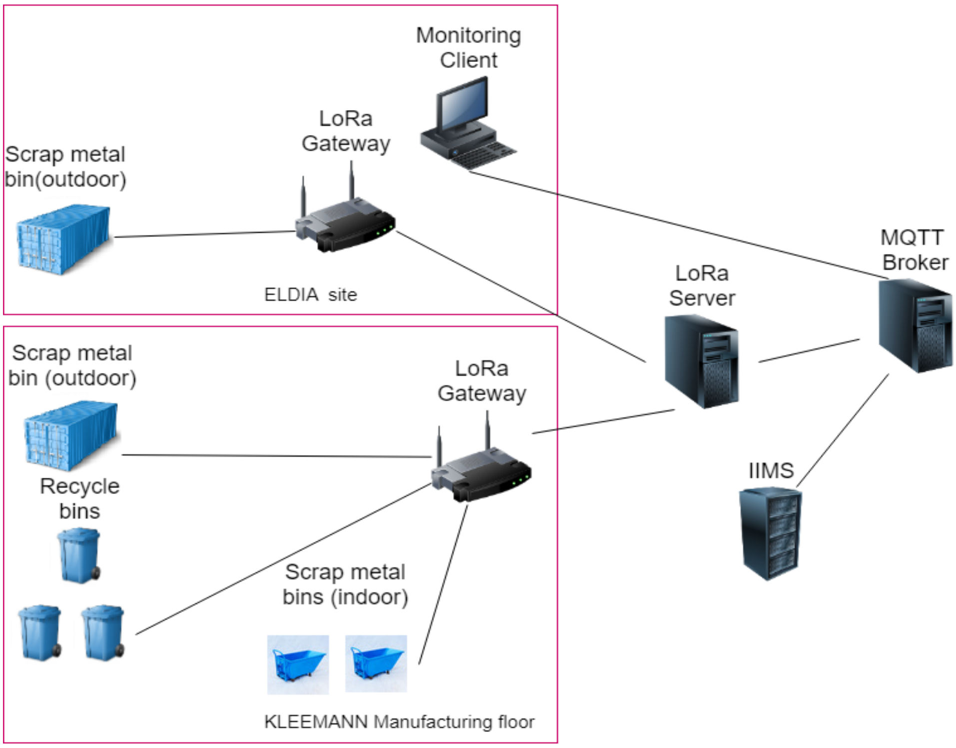

Figure 19: UC-ELDIA-1, UC-KLE-4 & UC-KLE 3 Architecture

In each site a LoRa gateway, connected to the internet via Ethernet, was deployed to transfer the wirelessly transmitted LoRa packets to the LoRa server. The LoRa Server, after decoding the packets, connects with the MQTT Broker and sends the data to the IIMS, where the data are being processed, stored and become available to consumers.

5.2.1.2 Software Architecture

The developed Fill Sensor is a low power distance measuring device connected on a LoRa network. On reset the MCU sets up the connection, the peripherals and the sleep cycle. Most of the time the device is in sleep mode and consumes minimal power. When it wakes up, it performs approximately 13 measurements. Those measurements are filtered first using a window of valid measurements, discarding the rest. The remaining measurements are sorted and a median filter is applied on them. The filtered output then is passed to the LoRa packet, along with battery life and bin-id, and is sent to the LoRa gateway.

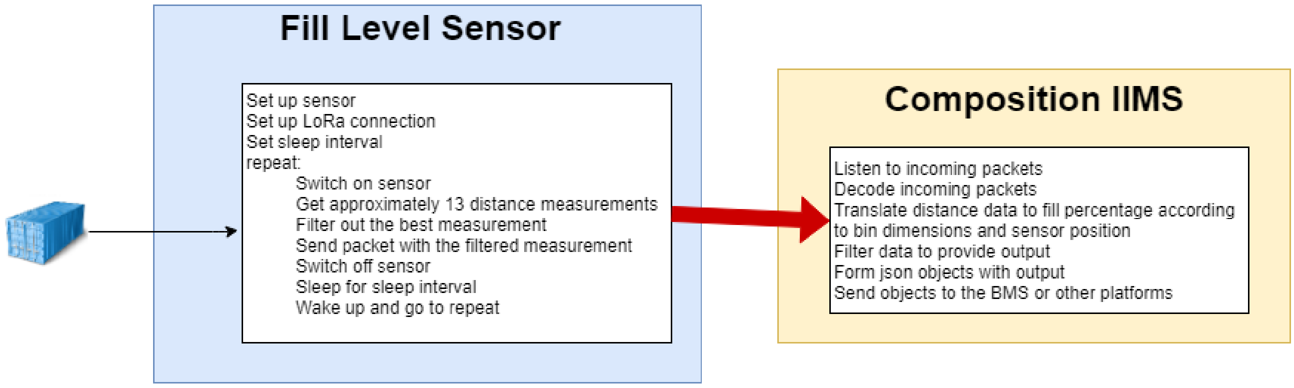

Figure 20: UC-ELDIA-1 & UC-KLE-4 Software Architecture

The packet then via the LoRa back-end is passed on the MQTT Broker of COMPOSITION and is decoded in IIMS. The distance then is translated into fill percentage according to bin dimensions and sensor placement. Then, further filtering is applied according to previous measurements and the filtered output is passed to BMS or other consumers.

5.3 UC-KLE-1 Predictive Maintenance

5.3.1 Design, Implementation and Deployment

The predictive maintenance use case features a sensor front-end that detects and transmits vibrations observed on two Bossi motors during operation. Also, a software back-end analyses the input data and predicts any motor malfunction in real-time. This framework can also be applied for historical data.

5.3.1.1 Hardware Architecture

Figure 21 below provides a system overview. The vibration sensors are LIS3DH accelerometer sensor break out boards each connected via SPI bus with an ESP32 SoC with integrated Wi-Fi communication. If a Bossi motor is operating the vibration sensor is detecting its vibrations and performs a sampling of raw measurements in 3-axis that is immediately sent over Wi-Fi and MQTT to the Broker. The procedure is repeated as long as the Bossi motor is operating.

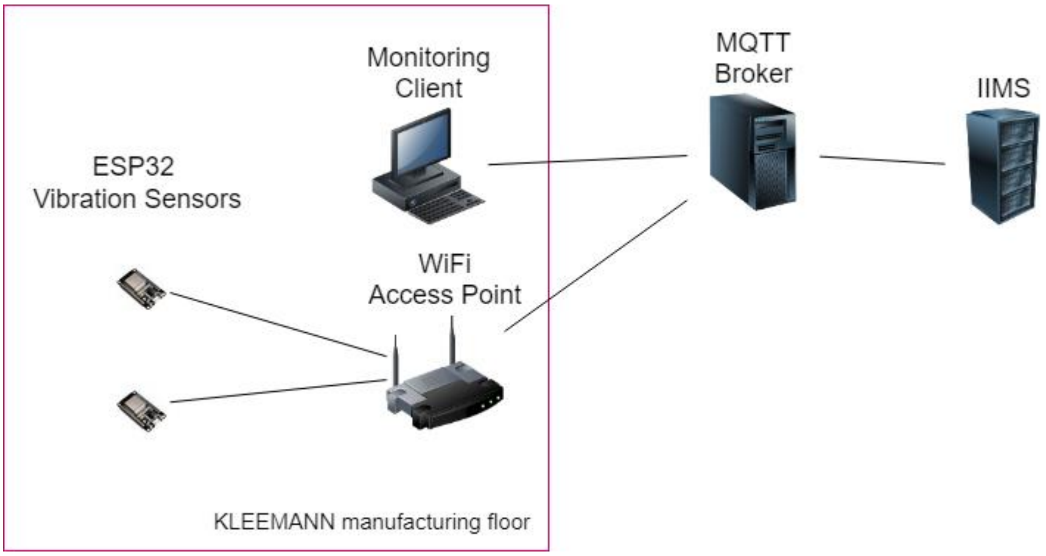

Figure 21: UC-KLE-1 Predictive Maintenance Architecture

IIMS, having subscribed to the topic where the sensors publish the raw data, is translating the data into accelerations measured in m/s^2, provides the data format BMS (or any test setup) demands and directs the data back to the broker publishing on the agreed topic. COMPOSITION Server provides the input to the BMS platform where the data are stored and relayed in real time for any client application that demands them.

5.3.1.2 Software Architecture

The Vibration sensor runs a bare metal program that detects vibrations above a threshold and when they are present, it performs subsequent samplings at 1.344 kHz. Each sampling contains 1344 samples of raw data per axis. Due to Wi-Fi packet restrictions data are sent per axis along with the sampling timestamp.

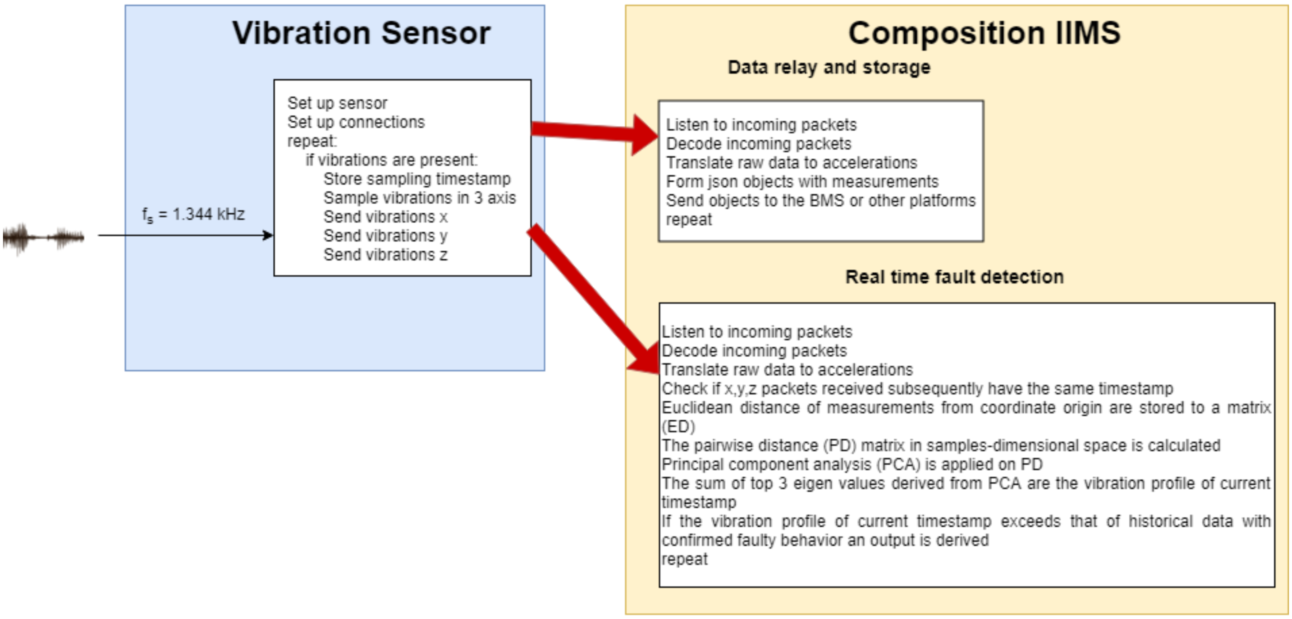

Figure 22: UC-KLE-6 Software Architecture

In the back-end there are two programs running. The Data relay and storage is the backend that is used for relaying and storing the data to BSM. The first thing that this script does is to translate the raw data values produced from the sensor into accelerations of m/s². Those translated values are then passed in the json objects that are formed, according to specifications from BMS or other data consumers, and then are transmitted to them via MQTT protocol. Various versions of this script were used to ease the development purposes. The Real time fault detection back end is a version of the first back end but with the data analysis and fault prediction mechanism integrated. Figure 22 gives a brief description of the data analysis and fault prediction algorithm.

5.4 UC-KLE-3 Scrap Metal Collection and Recyclable Waste Transportation

5.4.1 Design, Implementation and Deployment

The Scrap Metal and Recyclable Waste Transportation is triggered by a full bin in KLEEMANN's production. In order for the Prediction engine to be able to estimate and propose the optimal path to follow for the transportation of waste to a central bin outside the production line, a fill level sensor for the internal bins has been designed and developed. A total of 14 fill level sensors are deployed, 12 are installed on 3 sets of 4 bins containing recyclable materials (plastic, paper, aluminium and cardboard) and the other 2 are installed on bins containing scrap metal. The sensors are mounted on 5 special metal constructions manufactured by KLEEMANN, as shown in the figures below.

Figure 23: Indoor Fill Level Installations

5.4.1.1 Hardware Architecture

The VL53L0X Time-of-Flight laser-ranging sensor is used to capture raw measurements of distance between the sensor and the top of the waste. The sensor is controlled by an STM32L053 low power microcontroller. The microcontroller-sensor communication is carried through an I2C bus that can serve multiple sensors, thus giving us the option to monitor 4 bins simultaneously. The wireless communication is carried out via the

SX1272MBAS LoRa module for STM. The whole system is powered by 4 AA batteries and is placed inside a junction box and a custom made 3d-printed case as shown in the figure below.

Figure 24: Fill Level Sensor for (a) Single (b) Multiple Bins

The complete hardware architecture is shown in Figure 24 as the indoor bin sensors are part of the same LoRa network described in section 5.2.1.

5.4.1.2 Software Architecture

The microcontroller checks for faulty measurements, repeating the measurement process multiple times and returning a specific error value depending on the nature of the error. Data are transferred via the LoRa low power protocol to the LoRa Gateway. The gateway used is the LoRank 8. It is important to note that the LoRa protocol allows transmission of only very small packets of data, meaning that only raw measurements of distance are transferred along with the error value and battery level measurement. The gateway is connected to the internet via either Ethernet or Wi-Fi and publishes the data on an MQTT topic on the cloud. A listener connected to the same broker as the gateway, reprocesses the data, converting them from raw distance measurements to fill level percentages, assigns unique ids to each of the bins and creates the json object to send to the destination platform.

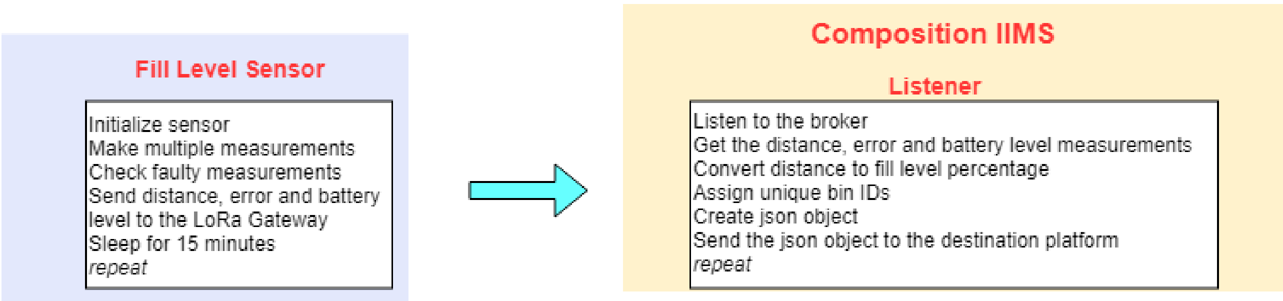

Figure 25: UC-KLE-3 Software Architecture

5.5 Lessons Learnt at KLE/ELDIA

5.5.1 UC-ELDIA-1 Fill Level Notification & UC-KLE-4 Scrap Metal Collection and Bidding Process

Although having the same device installed both in ELDIA and KLEEMANN, the ELDIA device performs better. The problem with KLEEMANN device is that the single type of scrap metal disposed appears to have reflecting or absorbing properties to ultrasonic. When bin in KLEEMANN is below 50% the sensor fails most of the time to take a measurement. Above that percentage measurements are frequent and precise. ELDIA bin contains various types of metal waste that have different properties. No such phenomenon was observed to ELDIA bin. Another reason why this is happening is the sensor placement. Due to limitations in the deployment the KLEEMANN sensor was deployed 1.2 meters above the bin, while in ELDIA it is 30 centimeters above. Generally the ultrasonic sensor appears to be more precise when the detected object is close and there are less chances for beam reflection or other type of failure to get a valid measurement. This means that as the bin progressively gets fuller the measurements are more precise.

Fill sensors were a custom solution deployed on bins that were not intended to have sensors on. As a future work it is recommended that the development of such devices is done in cooperation with bin manufacturers so that electronic devices can interface better with the bin and be protected against damage from waste.

5.5.2 UC-KLE-1 Predictive Maintenance

The Vibration sensor is providing a lot of data per second. The time between subsequent packets of data is heavily relying on the quality of the Wi-Fi network. A poor quality of network means packets delay, so less observability on the motor is provided. Scaling this use case (having sensors deployed on many motors) can also prove to be a challenge not only for the Wi-Fi network infrastructure but also for the Broker and the whole back-end.

Since the vibration sensor is based on an accelerometer, attention should be paid on the orientation of sensors if many are to be deployed. Orientation matters if we want to use the same metrics/thresholds or compare the behaviour of many motors. Bossi motors have the same orientation throughout the machine so this was not an issue

5.5.3 UC-KLE-3 Scrap Metal and Recyclable Waste Transportation

There are some limitations to the implemented solution due to the nature of the waste and the Time-of-Flight sensor limitations. The sensor can measure up to 1.2 meters, but the distance between the bottom of the recycle bin and the sensor is 1.3-1.4 meters so the minimum percentage we can measure is around 23%. The sensors are installed on such height so that they don't get in the way of the workers. However, since the COMPOSITION scenarios are activated when the fill level percentage is 70-80%, this limitation doesn't cause any problems.

Furthermore, the sensor measurement is based on the return of the laser light that was emitted. The surface of the scrap metal may reflect the light on various directions resulting in invalid ranging measurements below 50% in the scrap metal bins. The sensor raises a specific error informing us that the return signal is too low to give enough confidence on the distance measured. However, just like the previous case, this doesn't cause any problems. Moreover, depending on ambient light and distance we can get 3 to 12% ranging accuracy. Finally, if two objects are present on the field of view of the sensor, then the measured distance will be a weighted average.

6 Conclusions

Sensing in industrial environments is challenging and this is borne out in the work undertaken in the use cases described in this document. Providing reliable, accurate, low latency data, in industrial environments is not straightforward. It was found that the environmental, regulatory and internal factory regulations places significant requirements on sensing systems.

A number of prototype systems were designed to capture acoustic, vibration, ultrasonic, Time of Flight and RSSI data to determine, sound of failing fans, vibration of failing motors, fill levels of bins and location of material. The nature of prototypes is that there is inevitably some limitations to performance, however they need to be of sufficient quality to tease out the key challenges for future work. In this case the objective of these prototypes has been achieved as we now understand very clearly where future work needs to be focused.

The sections below describes in some more detail the conclusions on the use cases described in this document.

6.1.1 Condition Monitoring Use Cases

Condition monitoring of machines was implemented to determine their health. This in turn enables preventative maintenance, reduced down time and less waste. A number of such systems were implemented in this project.

One use case used acoustic sensors to listen to the health of fans in a flow solder oven to detect when they were about to go faulty, before their failure adversely effected the solder profile. The other use case detected vibration in a motor assembly to detect when a failure was about to occur.

In both of these it was found that reliable, low latency, accurate data was needed to enable the machine learning algorithms to function accurately and this proved to be challenging. In the case of the vibration monitoring system Wi-Fi was selected to wirelessly transmit data and it was found that the protocol does not provide for time deterministic data transfer. In the case of the acoustic sensors it was found that the use of commercial grade COTS components would occasionally be unreliable in the harsh environment in some parts of the oven. As mentioned above the machine learning algorithm places a very high bar on reliability of data flow, beyond that which is often required in other systems.

Other challenges found in this project was the amount of data collected at regular intervals can be very large. Moving large amounts of data with low latency can be a problem, especially if this is done wirelessly. It's unlikely that the systems as designed would be suitable for larger scale deployments and so any future developments would need to consider this aspect.

So in summary, these system gave a good level of performance but future developments need to focus on low latency, reliable, accurate scalable data transfer in tough industrial environments.

6.1.2 Asset Tracking Use Case

Accurate tracking of materials through large, complex, and dense factory environments is highly challenging. UWB uses Time of Flight with trilateration to detect position and is the most accurate technology but it is physically large, complex, expensive and power hungry. BLE uses proximity techniques that are not as accurate but they are cheap, low power and physically small. BLE is also more mature. For this project we opted to proof the concept of wireless location tracking in a smaller prototype area to see if it was possible to locate to below 2 meters this work is still on going. One of the issues with this technology going forward would be scalability in a larger floor area.

6.1.3 Fill Level Use Cases

For the Fill level use cases, many of the same challenges around robustness, scalability, reliability and accuracy as raised in the preceding sections also applies. Two techniques were used to determine the fill level of waste bins and both had issues with accurately detecting the fill level. It is proposed that some of these limitations could be overcome if the sensors were integrated as part of the bin itself and it is proposed that bin manufacturers should be approached to investigate the possibility as part of future work. These use cases used the LoRa wireless system which is a professional high grade high capacity wireless protocol and in this case no scalability issues were observed.

7 List of Figures and Tables

7.1 Figures

7.2 Tables

8 References

(SM) Spark Microsystems https://sparkmicro.com/our-technology/

(Friis) Friis equation for radiated RF pathloss http://www.antenna-theory.com/basics/friis.php