Climate-KIC

D2.4 Technical Architecture

CTT2.0 Carbon Track and Trace

Dirk Ahlers, Patrick Driscoll (NTNU)

Corresponding author: dirk.ahlers@ntnu.no

Trondheim, Norway | 15 December 2016

climate-kic.org

Climate-KIC is supported by the
EIT, a body of the European Union

Contents

1. Preface

1.1. About LoCaL

This report was written through support from Low Carbon City Lab (LoCaL). LoCaL aims to reduce 1Gt of CO2 and mobilize €25 billion of climate finance for cities annually by 2050. It is an innovation platform aiming to provide cities with better tools for assessing greenhouse gas emissions, planning, investing and evaluating progress. Started in 2015, LoCaL is a growing community of more than 20 organisations dedicated to unlocking climate finance for cities. This report was realized as part of the project Closing the Gap through Transformative LoCaL Action (CGTLA) under LoCaL.. LoCaL is a Climate-KIC flagship programme.

http://local.climate-kic.org. Contact: victor.gancel@climate-kic.org

1.2 About Climate KIC

Climate-KIC is the EU's largest public private partnership addressing climate change through innovation to build a zero carbon economy. We address climate change across four priority themes: urban areas, land use, production systems, climate metrics and finance. Education is at the heart of these themes to inspire and empower the next generation of climate leaders. We run programmes for students, start-ups and innovators across Europe via centres in major cities, convening a community of the best people and organisations. Our approach starts with improving the way people live in cities. Our focus on industry creates the products required for a better living environment, and we look to optimise land use to produce the food people need. Climate-KIC is supported by the European Institute of Innovation and Technology (EIT), a body of the European Union.

1.3 About Carbon Track and Trace

The Carbon Track and Trace (CTT) project is intended to provide cities with real-time greenhouse gas (GHG) measurement capability. Traditional methods of building and maintaining municipal GHG emission inventories are expensive, time-consuming, and are of questionable utility for mitigation decision and planning support processes. CTT couples low-cost, open source sensors to a Big Data analytics platform that provides cities and regions with a unique capacity to directly measure the impacts of their policy and planning decisions and to develop a semi-autonomous system for building, maintaining, and reporting their annual GHG emissions.

http://carbontrackandtrace.com/

2. Executive Summary

This document describes the technical architecture of CTT and the coordination and cooperation of its components. It describes the overall CTT concept, the technical architecture, deployment and installation in the field, data flows, components, and analysis and visualization systems that are being built. It details components that discuss and implement the technical architecture, components, and analysis approaches. This report lays the basis for the overall system and reiterates important aspects together with design decisions and stronger focus on overall system deployment and components and the different technical contributions and tools built throughout CTT. One objective of the project that can thus be met is to test the viability of a distributed GHG measurement system approach especially in terms of data quality and trade-offs between high-accuracy and low-cost approaches.

3. System Architecture and Components

The technical architecture of CTT derives from the overall conceptual approach of using real-time measured data and structuring the project along the data flow from deployed sensors up to tools for stakeholders. This section describes the basis of the overall system and the coordination and cooperation of its components. It links to other deliverables and a scientific publication [NIK2016] (see references) that discuss in detail and implement the technical architecture, components, and analysis approaches. The CTT concept serves as the original structure, based on which a detailed technical architecture is developed, which is then detailed with individual components.

Concept

The idea of CTT is to enable city officials, decisions makers, citizens, and other stakeholders to access emission measurements throughout a city. To achieve this goal, we define a general concept and overall architecture of the system as outlined in the concept figure. It shows the system components and the simplified dataflow, starting from individual sensors through gateway antennas to a cloud data storage into an analytics backend that provides insight and visualizations to a range of stakeholders in various degrees of abstraction. The architecture is kept as flexible as possible to be able to exchange components easily with clear interfaces between them. This facilitates collaboration and development within the project for separation of concerns. For example, the wireless IoT backbone used in the project can be scaled out to drive a city's IoT projects, but on the other hand, in a city with existing IoT backbone, the project should be able to integrate easily without the need for its dedicated network. The same holds for the data storage, which may instead use a city's existing open data portal. With this in mind, detailed hardware and software/protocol components can be exchanged with limited effort.

On the larger picture, CTT aims to better understand the gaps in current approaches and the tradeoffs inherent in the technical solutions that are being piloted. These are for example the conflicting priorities between high-quality measurements and the low-cost approach, the gap between concentration measurements and emission estimates, as well as the overall project aim of understanding relations between sensor-based measurements and yearly statistical incentories.

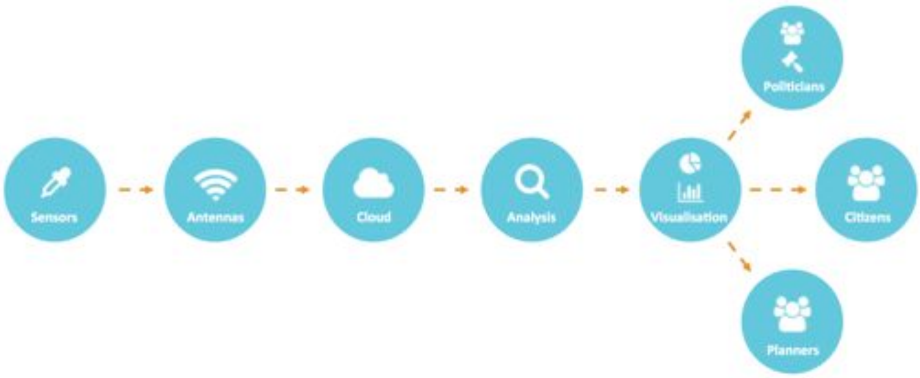

Figure 1 Concept Architecture

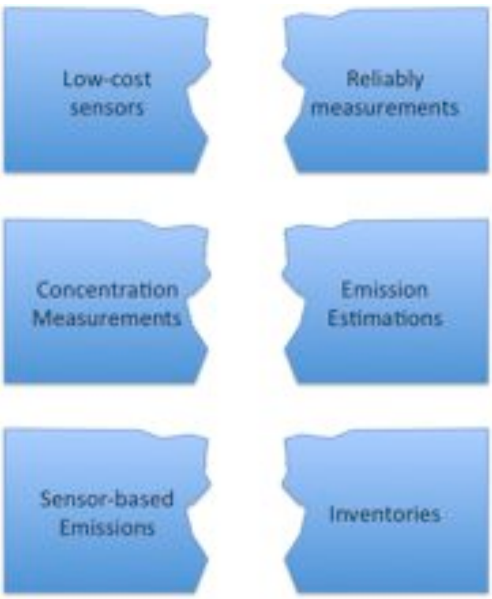

Figure 2 Gaps and Tradeoffs

Deployed Architecture

The actual implemented prototype of CTT (as an instance of the concept architecture) for the cities of Trondheim and Vejle sets up the whole chain from sensor over gateway antennas and storage to analysis and visualization systems through a range of components.

Figure 3 Component Deployment Architecture

Sensor communication and protocols

Details of used protocols and dataflow are shown in the following figure.

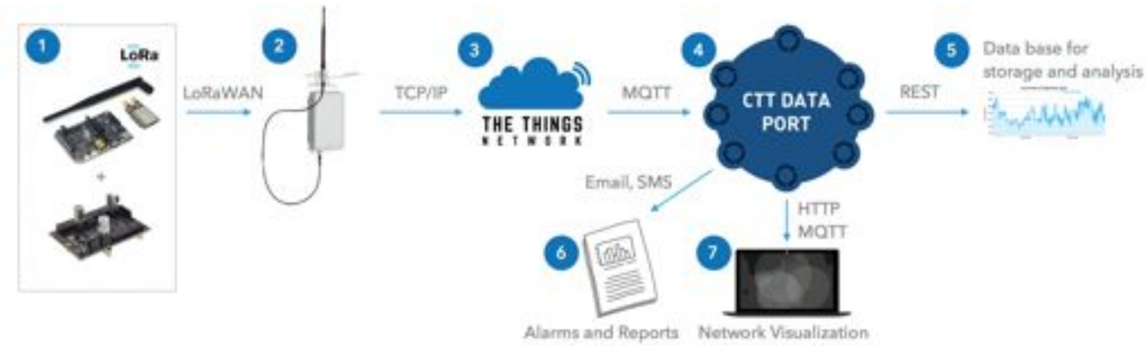

Figure 4 Network architecture and protocols [NIK2016]

Data pipeline architecture

Sensor nodes:

A sensor node is a Libelium Waspmote sensor box with installed sensors. Currently all communication from the nodes is via LoRaWAN and the local gateways.

The current data frame payloads are dependent on deployed sensors per node. There is limited variation in that we have 2 main deployment sets:

A. CO2, NO2, temp, hum, pres, PMx, bat
B. CO2, NO2, temp, hum, pres, bat
C. CO2, bat (only the first sensor at Elgesetergate, Trondheim; will be replaced with a full deployment)

Each data frame combines metadata from the nodes and measurements from the sensors into a binary string that is sent over the LoRaWAN. The encoding needs to be done carefully, as the

Libelium libraries encode values differently depending on the type, for example float in little endian format and int in Big endian. Similarly, care has to be taken decoding binary into hex or ASCII. Details are discussed in D2.3.

The code installed on the nodes is available on GitHub, details are described in D2.3.

Gateways:

The gateways listen to the LoRaWAN nodes within range and forward any received packets to TTN over an ethernet connection. We currently use Kerlink and Multitech outdoor gateways to act as concentrators for a number of sensors in their coverage area. They run a LoRa packet forwarder that reads and transforms RF packets received over the wireless and forwards it over a TCP/IP based network. In Trondheim they send data through the Wireless Trondheim network, in Vejle the single gateway is attached to the city network within the city hall building and separated by a firewall.

The things network (TTN):

Currently TTN is used as a concentrator backend. The reason is that it works well with the Libelium system and is a useful sink for data. TTN collects and holds values received through a LoRaWAN network for a short while. It allows managing of nodes and data can be retrieved through the messaging protocol MQTT or node.js in a JSON frame.

Nodes have a device ID and they are also configured as part of a pool of sensors, currently set to one pool per city identified by an application ID within TTN. It is predefined in the configuration for each node. This is used to identify and filter messages downstream in the message receivers, storage and analysis backends. The application ID is linked to a closed messaging concentrator within TTN that allows to access all devices grouped within that application.

Application side:

Applications can directly access measurements by listening to a defined application ID in TTN and retrieving live data streams. An alternative is the data stored within a MonetDB database through the system described in D3.1 and using standard database queries. Applications have to take care to correctly parse values and identify error codes.

A standard parser is discussed in D3.1, details are given in D2.3.

A node has one or more sensors attached. If a sensor is in a node there should always be a value transmitted. Some error codes have been defined to identify errors for monitoring or for exclusion from analysis.

-99: 'sensorError' error on the sensor in the waspmote: 'something went wrong', possibly unconnected or something else

-98: 'nonmeasured' Intentionally not measured due to low battery

Deployment of Sensors and Gateways

The main requirements of the deployment in the city are a good coverage of the urban area with sensors (as mapped out by the municipalities), a stable deployment package of the sensor units, which have to be weatherproof (thus an enclosed unit, but with exposed sensors), and a connection of the units to the central data storage, realized through a coverage of the city area with a wireless transmission system. To further ease installation, a base plate was developed (D2.3) that contains all sensor unit components.

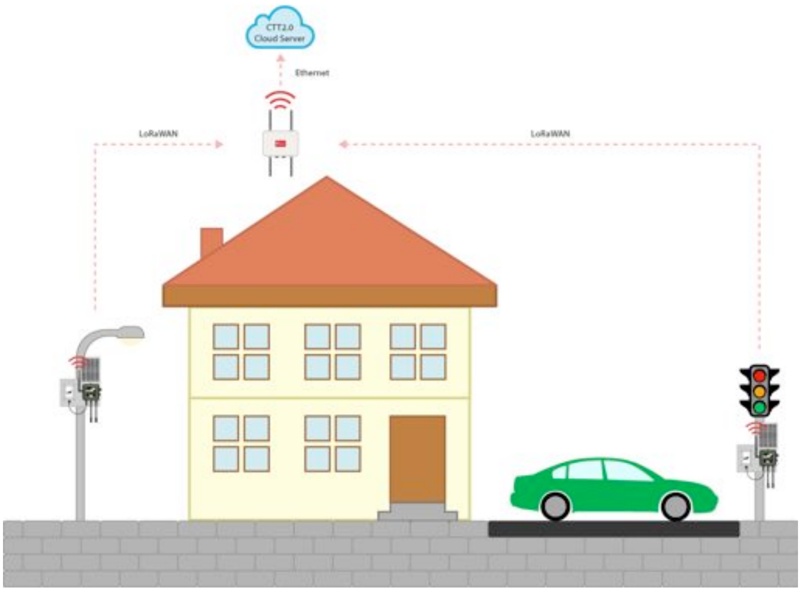

Figure 5 Urban deployment sketch, figure courtesy of Libelium, from the Libelium/CTT case study: Enhancing environmental control and reducing emissions in Nordic Smart Cities (http://www.libelium.com/enhancing-environmental-control-and-reducing-emissions-in-nordic-smart-cities/)

Sensor deployment

Figure 6 PSSEP deployed in Trondheim (left) and Vejle (right) using deployment plate

Figure 7 Deployment locations in Vejle

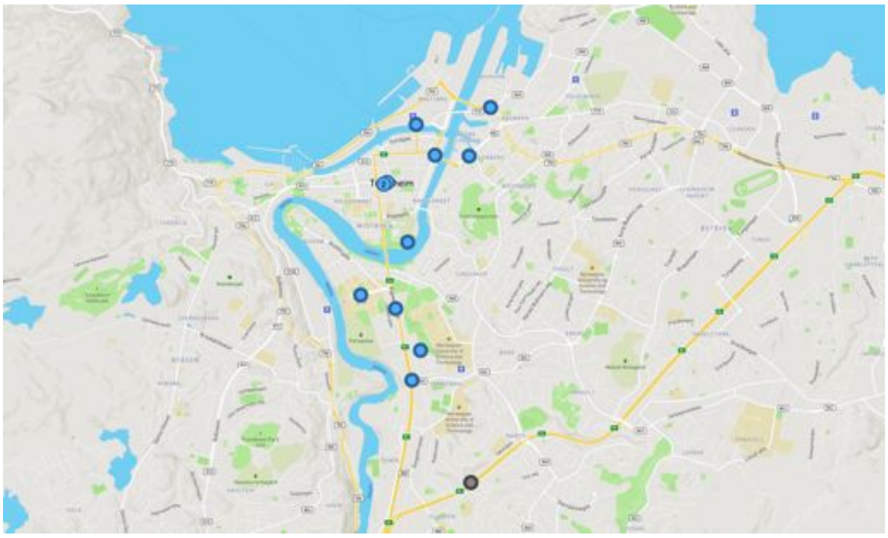

Figure 8 Deployment locations in Trondheim[1]

[1] Picture from: http://dataport.item.ntnu.no/

Gateway deployment

Figure 9 Deployment of LoRaWAN antenna gateway outdoors on the roof of the Student Society building in Trondheim (left) and elevated indoors in a clock tower in Vejle (right)

4. Analysis and Visualization

The platform, architecture, and the data provided by the sensors have been the basis for a number of analysis and visualization components and applications. These are presented in the following.

DataPort

The Dataport (http://dataport.item.ntnu.no/) is a tool designed to monitor the network health and state of the connected IoT devices such as gateways and nodes.

It provides a graphical interface showing the network by connected gateways and sensors, signals received, associations between gateways and sensors, and message contents. For faster reaction, the Dataport connects through MQTT to the network and uses its own InfluxDB database for storage of the network metadata along with the payload of the measurements.

The future goal is to use the tool also as an alerting service that can generate detailed reports and alerts when part of the infrastructure fail to be able to respond to sensor, node, or gateway outages. Details are found in [NIK2016].

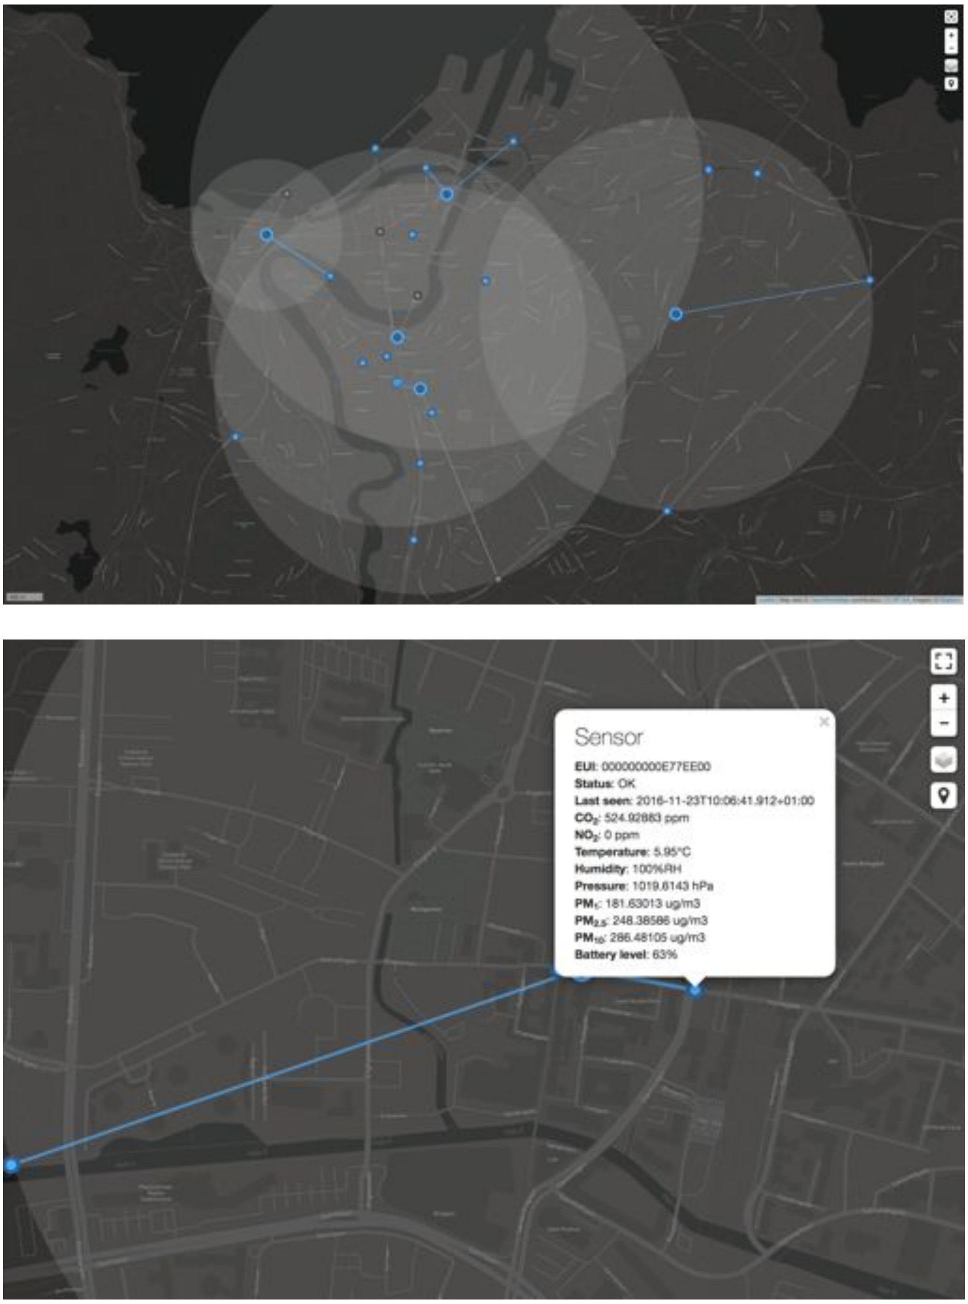

Figure 10 DataPort Full overview of a city area, and Detail view or sensor data for a node

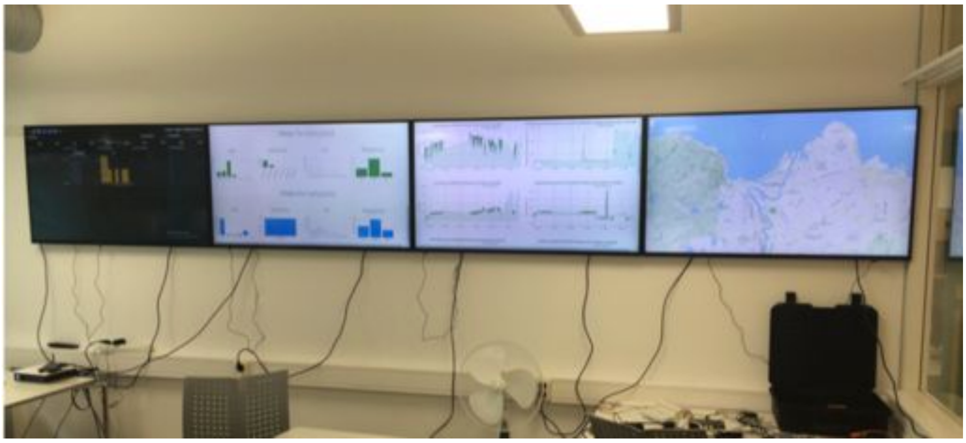

Figure 11 For a better overview, the possibility of a larger integrated dashboard on network monitoring and data visualization has been explored in the NTNU Telematics Lab.

Data Storage and Analysis

As noted above, the flow of the sensor data uses MQTT connect the nodes through TTN with the CTT backend. A Numascale R appliance has been set up as described in D3.1 that uses a MonetDB to store measurements for the sensor system and can make it available.

The storage system uses a combination of R and Python code. The software and data flow for this subsystem is shown below.

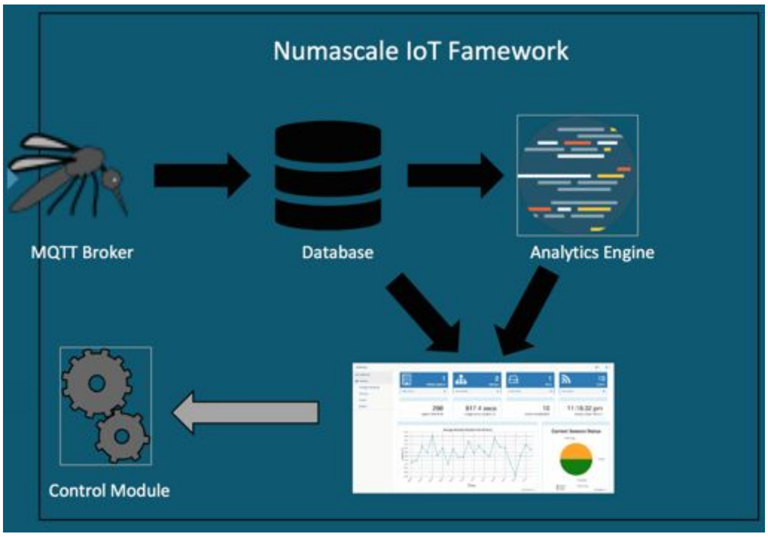

Figure 12 Data Storage system

Link with NILU Air Quality Data for Trondheim

On the basis of the MonetDB and the Numascale R Appliance, a set of proof-of-concept analytics studies was done, including the use of external data. The main example is the combination of NILU (Norwegian Air Quality Research Institute) air quality data from official measurement stations in Trondheim. NILU NO_X data was combined with weather data to show correlations between temperatures, wind, air pressure, and emissions.

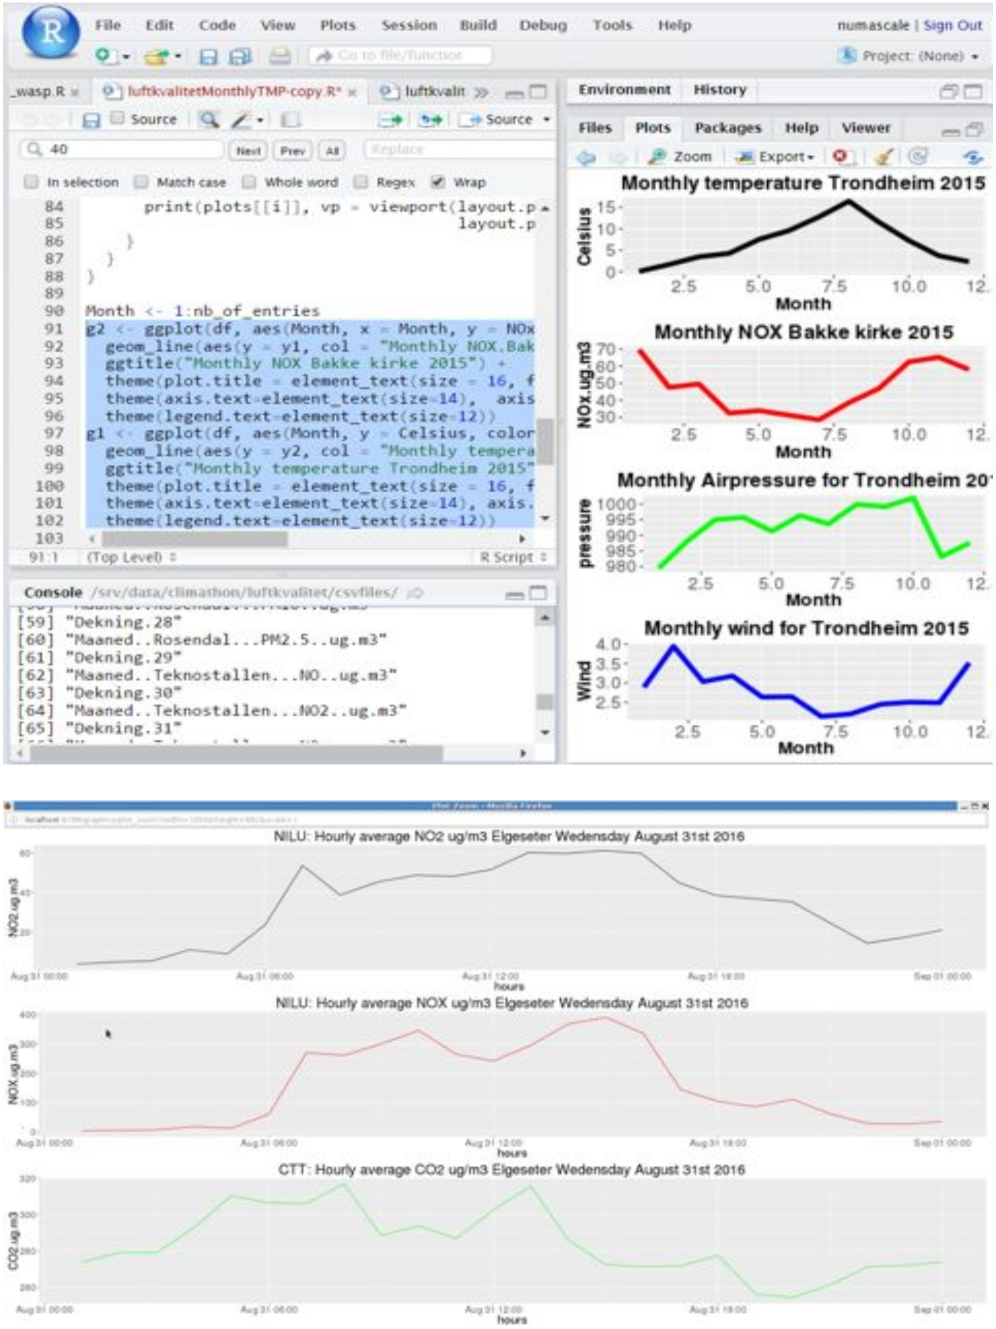

Figure 13 NILU air quality data analysis

Traffic and Air Quality correlation

An air quality dashboard has been developed as part of D3.2 by DTU that can show live and historical CO2 values, calculate air quality metrics on top of it, and initially explore correlations of traffic flow to emissions. This has been built as a demo for Vejle.

It can represent traffic in Vejle and the level of CO_2 measured by CTT in equipped streets of the city center and allows a user to interactively explore the data. Details are found in D3.2.

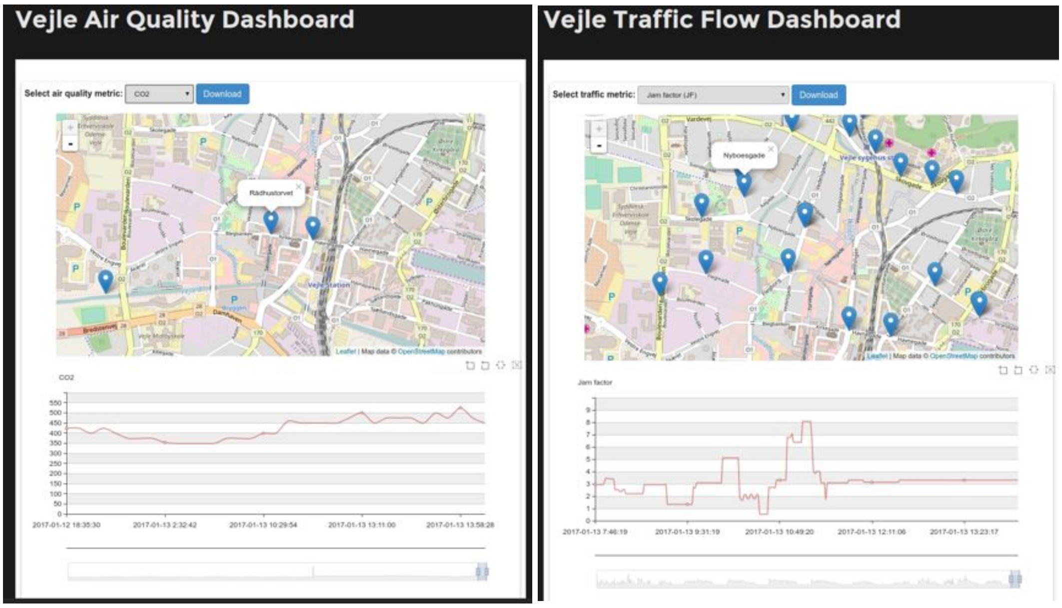

Figure 14 Traffic Flow and Air Quality Dashboard showing historical traffic data and live and historical CO_2 measurements in Vejle

3D City Model Sensor Integration

A Prototype of a 3D GIS city model has been built upon the developed CTT system with inputs from deployed sensors. The pilot city was Vejle due to easier data conversion of the 3D model. In partnership with Virtual City Systems, Vejle Kommune, DTU, and NTNU, CTT has extended the existing CityGML 3D model of the city of Vejle to include real-time sensor data from the installed Waspmote Smart Environment Pros. An integration of real-time traffic data was not possible for technical reasons. However, this was achieved in the above Vejle Traffic dashboard as shown in D3.2. The details of the integration and setup are described in D4.4.

Figure 15 virtualcityMAP visualizing the 3D city model and the sensor locations.

Figure 16 Visualizing the measurements of a sensor in a graph linked to a situated symbol in the 3D scene.

Additional Visualizations, Applications, Analyses

During the development and debugging of the system, a number of other visualizations have been built. The following shows a small overview of data collected. A much wider range of visualizations for collected data and discussion can be found in D2.3.

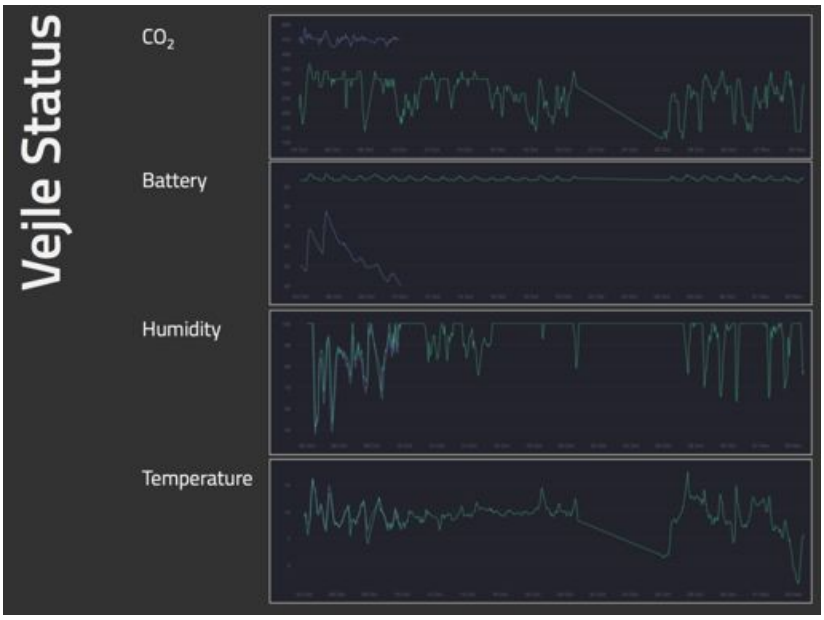

Figure 17 Vejle Dashboard mockup with a selection of real-time data.

Another question was how to possible base measured data from low-cost sensors with uncertainty of measurements to some other source of data that would allow a comparison. Here the OCO-2 satellite data was used. The overflight paths were plotted for Trondheim and the variance of the measured concentrations was analysed, showing a certain range of values around 400ppm.

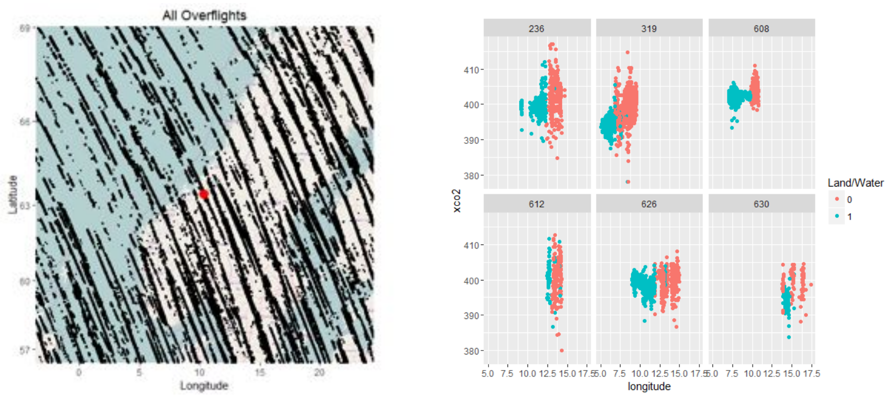

Figure 18 OCO-2 overflights (Black dots) around Trondheim (Red dot) (left), Analysis of variance of measured CO2 (right). (Made by Kasper Einarson)

More details about the inclusion of different forms of external data, such as traffic, air quality, or remote sensing satellite data are discussed in D3.1 and D3.2.

5. Conclusion

The CTT architecture and main components have been presented. The resulting system has shown flexibility in linking up to a number of other systems and frontends within the CTT project and demonstrates that the system is sufficiently flexible to easily swap out components based on different requirements or when better solutions become available. This shows that a working system has been successfully implemented in the two pilot cities of Trondheim and Vejle. Improvements will still have to be made to all components of this first prototype developed within only 9 months. Integrating any changes should be rather straightforward given the gathered experience up to this point. A major open issue is the data quality of the sensors themselves as well as the need to capture more data to develop better models based on a sufficient background of historical measurements. This will then enable better models to be built on top of the approaches demonstrated so far. That would then also be the way forward to apply more complex models on better quality data to move from measurements of concentrations to fluxes and emissions.

References

CTT2.0 Deliverable D2.3: Report and evaluation of low cost sensor system (Sensor Field Reports). Fredrik Valde Anthonisen, Dirk Ahlers, Patrick Driscoll (NTNU)

CTT2.0 Deliverable D3.1: Prototype and Deployment of a Big Data Analytics Platform. Atle Vesterkjær, Patrick Merlot (Numascale)

CTT2.0 Deliverable D3.2: Big Data Analytics. Lasse Engbo Christiansen, Tobias Straarup Andreasen, Per Sieverts Nielsen (DTU)

CTT2.0 Deliverable D4.4: Prototype 3-D GIS city model of Vejle with GHG emissions from installed sensors. Lutz Ross (virtualcitySYSTEMS)

CTT2.0 Report: Sensor Network Overview. Thomas Ulleberg (Wireless Trondheim)

A Measurement-Driven Approach to Understand Urban Greenhouse Gas Emissions in Nordic Cities. Dirk Ahlers, Patrick Driscoll, Frank Alexander Kraemer, Fredrik Anthonisen, John Krogstie. NIK — Norwegian Informatics Conference 2016.