An Agricultural Monitoring System for Smart Farming: Design and Implementation

Dr. K. P. Porkodi,

ASP/CSE, Vivekanandha College of Technology for Women

S. Narmatha, V. Kavitha, M. Sonipriya,

ME/CSE, Vivekanandha College of Technology for Women

Abstract

Industrial agriculture has already profited from the integration of current information technologies by increasing yields over the last few decades. In the context of "Smart Farming," new applications and potential for a more accurate, site-specific, and sustainable agriculture are now made possible by the burgeoning Internet of Things (IoT) and Wireless Sensor Networks (WSNs), which have low-cost sensors and actors. The design and architectural implementation of a comprehensive agricultural monitoring system are presented in this study. The main focus of the system is on in-situ evaluation of the leaf area index (LAI), a crucial agricultural metric. Moreover, we introduce real-world challenges and experiences gathered in diverse deployments. To quickly discuss the possibilities of our technology, first results are exemplarily displayed.

INTRODUCTION

Modern agriculture faces significant challenges due to the changing environment and rising food demand. According to projections, the world's population will grow dramatically by 2050, necessitating a 70% increase in food production [7]. The lack of arable land and water are both getting worse at the same time. As a result, there is a need for sustainability, an appropriate adaptation of farming techniques, and an adequate selection of crop kinds. Conditions need to be adjusted for when it comes to irrigation management, crop treatment, pest control, and fertilization. Farming operations must be carried out in a smart manner to conserve limited resources.

Modern agriculture faces significant challenges as a result of both the changing environment and the rising food demand. The global addition to agriculture and agronomy, research skills from many other fields must be effectively combined in order to accomplish this. Utilizing scientific advances and cutting-edge

technologies from other fields has already received a lot of attention. Precision agriculture is a concept that has emerged as a result of the increasing integration of digital advancements into the agricultural industry [12]. Internet of Things (IoT) principles have recently expanded Precision Agriculture with intelligent, distributed, and cooperating sensors and technologies that are now well established in other industrial sectors as well as in home automation [6]. This expansion, known as "Smart Farming," includes all phases from sensor-based data collection and communication to in order to help farmers to make better and more sustainable decisions.

IoT technologies have the ability to offer fresh perspectives and current situational awareness with a significantly higher level of spatiotemporal granularity of monitoring in the context of arable farming [6]. They encourage knowledge of the variables affecting crop development and yields, which is essential for a sustainable agriculture. IoT solutions enable farms dramatically save resources through sitespecific management, which boosts farm output. Additionally, IoT-based crop monitoring enhances yield modelling and the accuracy of production forecasts. In general, the burgeoning digital revolution, especially the incorporation of IoT into modern agriculture, is a significant enabler that makes it possible to automate many processes and provide them with useful extra information.

Since sensor-based data collection forms the backbone of the IoT chain from crops to farmers, the success of smart farming depends heavily on the deployment of sensor devices in the field. These products range from simple, low-cost sensors with limited resources to sophisticated, high-accuracy sensor platforms that may be very expensive. A wide-scale agricultural monitoring often requires a huge number of sensors. Therefore, from an economic standpoint, the cost of each sensor is essential for the return on investment. (RoI). Such inexpensive sensors are frequently employed in Wireless Sensor Networks (WSNs) for ambient sensing of physically quantifiable factors [1]. These are extensive and intended for long-term deployments and are wirelessly connected. They are also very resource-constrained because they are primarily battery-powered.

The huge number of cooperating devices continuously gather environmental data to make up for the poor sensing precision of individual sensors. Furthermore, airborne imagery-based remote sensing can be used to supplement ground-based WSNs. This is often obtained by satellites or, more recently, by drones, and is used, among other things, to determine agricultural growth. This helps to approximate WSN-based in-situ information over even broader areas, which is very advantageous. During the last decade, WSNs were already deploye in the agricultural domain, improving remote monitoring of agricultural resources and products, e.g., [8], [13]. Their potential of enhancing productivity and waste reduction has been shown to be quite promising [2]. Recent advances in cloud technology have enabled successful integration of similar deployments onto IoT systems [15]. In this work, we present an agricultural crop growth monitoring system and report on our experience of real-world deployments.

The focus of these deployments is on a specific crop metric, notably the leaf area index (LAI). The LAI is a frequently employed crucial metric that offers details regarding the photosynthetic activity and important circumstances of plants, see [10]. The parameter, which is connected to vegetative biomass, is simply the ratio of leaf area to ground surface area. It is extremely interesting in the context of agriculture and utilised for yield modelling because it also functions as an indication for yield-reducing processes brought on by diseases or poor management. Long-term continuous crop monitoring, which permits LAI profiles with a fine-grained spatio-temporal resolution, is the main objective of our system.

As a result, we are using the sensor prototype we previously created [4] (see Fig. 1). In the range of photosynthetically active radiation (PAR), it detects the ambient light. The transmittance of solar irradiation through the canopy, which enables an assessment of the LAI, can be determined from two simultaneous PAR measurements, one below and the other above the canopy.

II. ACTUAL DIFFICULTIES

Although a lot of progress has been made in the last ten years, real-world WSN deployment and maintenance remain highly difficult tasks. For long-term outdoor deployments, there are additional hurdles besides the typical WSN challenges, such as the hardware limitations of small sensor devices, their power consumption, and low-power transmission (cf. [1]). The major reasons of such challenges can be divided into two categories: obstacles brought on by the environment and, more specifically for agricultural deployments, challenges brought on by wildlife. Figure 1 gives a brief idea of the technical issues we encountered during our deployments through a chosen collection of images.

A. Environmental Challenges

The operability of a WSN is significantly influenced by the surrounding environment. Sensor motes are exposed to extreme weather, likely involving significant temperature fluctuations, rainfalls. Also humidity and soil moisture tends to be comparatively high, particularly if devices are directly deployed in the field. For instance, it was discovered during our first deployment that the cases we designed were not entirely sealed and durably water-resistant. As a result, some sensors developed condensation water under the diffuser cap (see Fig. 1(a)).Motes are additionally vulnerable to corrosion and short circuits, which may cause operational instability. Even worse, catastrophic hardware failures can happen and have in fact done so (cf. Fig. 1(b)).

In addition, sensors located at ground level may be affected by mud during rainy seasons as opposed to dust during dry seasons (cf. Fig. 1(c)).However, as radio link quality are subject to changing environmental extremely circumstances, adverse weather affects not only impair sensor hardware but also the connectivity of the entire network. Moreover, agricultural lands normally situated far away from the energy gridWhile this has not impact on low-power and already batterydriven WSN devices, the lack in reliable power sources does affect more complex IoT components such as weather stations and Internet gateways. Hence, it has to fall back on solar energy solutions. Unfortunately, it appeared that the installed solar energy equipment in some of our deployments was not completely dependable and led to partial disruptions.



(a) Condensation water in the diffuser cap of a PAR sensor.



(c) Mud on a ground-level PAR sensor.



(b) Burned-out hardware box caused by cable fire.



(d) Bite marks from wildlife.

B. Wildlife-related Challenges

Conflicts with wildlife and animals are inevitable whenever sensors are placed in remote locations for an unattended operation. For ground-level equipment, moved sensors, nibbled cables, or even bit marks on sensors (cf. Fig. 1(d)) are not unusual occurrences. However, sensors positioned on taller stands above crop level may also be momentarily obscured by birds. Particularly in our use case that depends on PAR measurements, this covering is totally interfering. Finally, likewise insects forming nests in sensor housings can have harmful impacts. Overall, similar problems in both categories were also highlighted by [11] more than ten years ago, but subsequent works on the subject demonstrate that they are still relevant and exist in the community. Additionally, agricultural WSNs are reportedly plagued by farming activities and, sadly, vandalism [8], both of which thankfully did not happen during our deployments.

Industrial outdoor cases, improved solar energy gear, expert uninterruptible power supply (UPS) systems, stronger and shielded cables, or electric fences could all help to lessen some of these problems. However, it is costly and a tradeoffs between such additional costs and operational safety arise. Nevertheless, It is hardly conceivable to tackle all challenges that potentially could emerge in real-world deployments and unforeseen issues should still be expected. A non-disruptive WSN operation could not be guaranteed in practice, and false sensor readings could never be completely avoided.

A. Concept & Architecture

We employ tried-and-true concepts in order to handle the two different challenges outlined in the preceding section. Our main strategies are redundant hardware, straightforward software, and system-wide remote control. The main component of the architecture we have created is a WSN-based monitoring system that is specifically designed for in-situ LAI assessment. In order to quantify incoming solar irradiance in the PAR range, two crucial measurement sites are needed: a groundlevel sensor below the canopy (G) and a corresponding above reference sensor (R), see [4]. To achieve equivalent irradiation conditions, the devices must not be placed in the same location and must not be too far apart spatially.

. So, for the purpose of LAI monitoring WSNs, we think that clustering a number of geographically dispersed ground sensors within communication range under the leadership of a single above reference is feasible. As a result, within each cluster, sensor motes were arranged in clusters using a straightforward star-topology. On the assumption that cluster heads are constantly active and do not have power constraints. This seems fair given that tiny solar panels may actually be used to power this type of sensor technology. In conclusion, there is no need for routing protocols in our current system. Due to the fact that cluster heads are continually accessible to ground sensors. time synchronisation methods are also not necessary. As a result, they might also modify their reference sampling to account for ground packet receive events. Various group leaders a central base station is linked to each cluster head in turn. Multi-hop routing may be required for these connections, depending on the size of the WSN, but is not yet taken into account in our deployments. Instead, WLAN is used to connect cluster heads.

As shown in the architectural overview of our system in Figure 2, the entire architectural concept includes not only the WSN itself but also an IoT-based infrastructure. The central base station serves as a gateway to a traditional IPbased IoT network for that reason. Through communication through public land mobile networks (PLMNs), which is made possible by an A farm management information system (FMIS) in general (or a customised server in our instance) receives the data generated by the base station and the sensors and uses it for data analytics and visualisation. MQTT [3], a publish-subscribe messaging protocol with hierarchically structured topics for individual message streams, implements data transmission. MQTT clients allow subscribers and publishers to exchange messages based on predefined topics.



Connections to these clients and their registrations are handled by a central broker. According to [5], [14], MQTT is appropriate for exchanging periodically acquired sensor data from agricultural WSNs. MQTT is now only employed at the IoT layer of our design; however, in accordance with [14], the integration of MQTT within the WSN is planned as future work. Because we use industry-standard open-source IoT software and commercial off-the-shelf (COTS) hardware, our architecture is generally modular in design and flexible. It is possible to effortlessly connect other agricultural sensor types, such as soil moisture and temperature sensors. Additionally, a smartphone-based LAI assessment [9] has been connected in a prototype manner. For our upcoming work, connecting to complementary technologies like remote sensing would be both possible and fascinating.

Weather-related issues might interfere with wireless connectivity in challenging outdoor settings. It is well known that link quality can vary greatly in practise. Humid air and moist plants weaken 2.4 GHz radio band signals, which are utilised in WLANs and WSNs. In actuality, we saw substantial mistake rates during our deployments. Forward error correction (FEC) methods, such as network coding, have already been shown to reduce them.be practical in our circumstance [14]. However, wired links can potentially break as a result of farming or wildlife activity. We cautiously chose to use both wired and wireless connections for redundant sensor data transmissions in our initial deployment generation due to intrinsically poor data delivery and to boost redundancy.

We also reused USB connections for additional powering all sensors because changing batteries during the deployment would be quite disruptive. Additionally, using USB, reprogramming sensors event of software failures in the or reconfigurations was simple. This made it possible to test and assess various software programmes. We made use of clusters of four ground mosses in accordance with the cluster concept. Each mote was connected via USB to a Raspberry Pi 3, a little, feature-rich, and reasonably priced single-board computer that is popular in the IoT space and runs Linux. A central ALIX.6F2 router with UMTS

connectivity enabled a WLAN for these laptops to be wirelessly connected to. We have already seen some software stalling and crashes in lab settings. As a result, we made the decision to add mechanical timers to consistently turn off all components during the nighttime hours when sensor readings are not necessary for operational safety. By doing this, we guarantee that the system will function at least the next day.



Waterproof Housing

In the second wave of deployment, the WSN was reconfigured to operate entirely wirelessly, including battery-powered sensors and wireless communication. In addition, TelosBs were employed as WSN gateways and Raspberry Pis as cluster heads. To improve radio performance and communication range, these motes also have external antennae. A central Pi with an LTE modem was used in place of the ALIX router that served as the IoT gateway for consistency and software unification. Real time clocks (RTCs) were additionally installed in each Pi, greatly simplifying the operation. The hardware prices of the two in-situ components of the monitoring system are listed in Table I for completeness.

1) Hardware Components for In-Situ WSN: The TelosB1 platform (8 MHz TI MSP430 MCU, 10 kB RAM), which is based on the low-cost IEEE 802.15.4 compliant sensor prototype disclosed in our early study [4], is employed as the primary sensor in our monitoring WSN. Three integrated environmental sensors for temperature, humidity, and light are a part of this open-source COTS device. It has been demonstrated that the latter sensor is ideal for PAR measurements that enable the generation of accurate LAI estimates when used in conjunction with an appropriate optical filter and diffuser attachment. Additionally, additional external sensors can be linked utilising the platform's GPIOs and SPI or I2C interface.

2) Software Components: The sensor platform is equipped with our system's fundamental acquisition software. We modified the sensing application from [4] and utilise TinyOS2, an open-source operating system widely used on devices with little resources. This programme is still maintained as basic and lightweight as it can be, with a focus on the most important functionality. Additionally, we use extensive logging of all data that can be collected via the radio interface and the universal asynchronous receiver transmitter (UART), together with time stamps and sequence numbers. (SNs). This also includes additional reports concerning other occurrences, including NTP synchronisation events or hardware component resets or reboots involving all associated components. Once more, the first and second generations of our deployments must be distinguished. A constant sampling rate of 30 samples per hour was employed in the first, which relies on a wired backbone for energy supply, however the rate is precautionarily decreased to 6 samples per hour for the second generation's battery mode.

Here, TinyOS's duty cycling feature known as low power listening (LPL), which switches devices to low-power modes when inactive, is engaged. This indicates that a 10-minute LPL period is configured in order to further lower the energy requirements of ground-level sensor devices. Each sensor sample in both generations consists of a collection of measurements from every environmental sensor on the platform, including temperature, humidity, and PAR, in that order. The PAR sensor is sampled many times in a brief burst of 25 readings with a temporal gap of 50 ms in order to obtain more accurate averages and also to deliberately examine small-scale changes. As soon as all sensors have been sampled, the collected data is combined and added as payload to an 802.15.4 frame, along with continuously incremented SNs.

The broadcast transmission of this data frame follows. All active motes within communication range can therefore receive it, especially cluster heads or base stations that forward payloads to the backbone system. Ground sensors can receive transmissions from one another by using broadcast transmission. Since information on the received signal strength indicator (RSSI) and link quality indicator (LQI) can be acquired from the 802.15.4 radio chip, they are used to keep track of link characteristics from neighbouring devices. In the first deployment generation with alwaysactive motes, that is very pertinent. Here, similar data is continuously gathered and transmitted along with each sample for the purpose of link quality monitoring and could be utilised for network protocol study in the future.

In fact, the chosen sample rate is a large oversampling in both deployment generations and is especially unnecessary for a pure LAI assessment. However, we want to build a

versatile, always-on testbed for WSNs as well as a large data set. This set is meant to be sufficient for in-depth research of a variety of parameters, including investigations into connection quality or pertinent effects on WSN-based LAI assessment. Since it is unlikely that hardware and software components will operate reliably during a deployment, we developed a number of safety measures. Passive SN synchronization, which takes the role of a stringent time synchronisation protocol, is one technique. This means that a sensor inspects the contained SN whenever it gets a transmission from a neighbouring mote. The current SN is modified if there is a gap between it and its own SN state, which was probably produced by a software reset. The second deployment generation's LPL makes it necessary to modify theSN synchronisation. Here, motes are configured to continue functioning for a set amount of time following reboots.

Additionally, we included safety features to the fully functional gadgets. For an unattended operation, for instance, Raspberry Pis are in charge of watching the incoming data from each attached cluster sensor and rebooting the afflicted sensors automatically. Additionally, the central gateway creates a persistent SSH reverse tunnel to the Internet server, enabling remote access to Pis and sensors that might be manually reprogrammed and customised to overcome unforeseen difficulties.

3) Energy Considerations

An empirical energy evaluation was undertaken to determine the energy requirements of motes and to confirm whether the battery capacity is adequate to ensure operation throughout the growth season of typical crop varieties. Using a Fluke 289 True-RMS industrial multimeter, we measured the electric current in Ampere over the course of 100 hours and discovered that the average consumption was close to 0.2 mA. This is in line with the TelosB specification, and even a very conservative energy prediction yields a sensor lifespan that is more than sufficient for the majority of agricultural applications and typical crop growing cycles.

C. Remote Observation

The Mosquitto3 MQTT broker, which is running on an Internet server, is the heart of the Internet of Things infrastructure. Periodically, the WSN-IoT gateway sends the broker particular messages with sensor data. These messages are effectively subscribed to, serialized, and saved in a database using Google protocol buffers4 (protobuf). A web-based graphical user interface (GUI) is offered by an Apache server, which also makes database queries. It is in charge of performing data analytics and ensuring that user-generated material, like graphs of the temperature and humidity or LAI trajectories, is visualised appropriately.

Information about the current sensor status is also retrievable in the GUI. As a result, the server keeps track of the network's functionality and alerts users when sensors malfunction or connectivity is lost. Failures are identified and messages are provided in the GUI using periodic sensor data as keep-alives. Additionally, instant messenger-based smartphone notifications are also possible. The server simultaneously allows access to certain sensors for the purposes of remote reconfiguration and reprogramming. The Internet server also uses meteorological data from outside sources, such as the Deutscher Wetterdienst (DWD)5, which it can integrate with WSN data.

IV. PRELIMINARY EXPERIMENTAL RESULTS & REAL-WORLD DEPLOYMENTS

In Lower Saxony, Germany, at two locations with experimental crop fields during the 2016 and 2017 growing seasons—(1) at the Institute for Crop and Soil Science, Julius Kuhn-Institute (JKI), Braunschweig, and (2) at the Faculty of Agricultural Sciences—we have experimentally deployed our monitoring system. as well as Landscape Architecture (AuL), both at the Osnabruck University of Applied Science. We noticed during these missions. Three economically significant crop types **winter wheat, rape, and maize** have been developed by LAI. Table II provides a crucial statistical summary of the installations and sample characteristics.

An examination of the gathered data appears to have a lot of promise. It might offer fresh perspectives, such as information on how the environment affects WSN-based LAI assessments. An investigation of this scope is now under way and will be included in our upcoming work, though. As a result, it is outside the purview of this article, which focuses on system design and lessons learned from actual deployments. However, as an overview, Figure 3 exemplarily illustrates some preliminary outcomes. A ground-level sensor that was placed in a wheat field has collected the temperature and humidity curves over the course of two days, as shown in Figure 3(a). It is possible to see that both curves have a midday peak and that they are similar but inverse to one another. The humidity curve also supports the comparatively high humidity that ground-level sensors on both days sensed.

The daily time series of PAR values acquired by a single cluster, consisting of one above reference sensor (R in blue) and four ground-level sensors (Gi, shaded in green), are shown in the second subfigure (Fig. 3(b)). LAI estimations can be obtained from the ratio Gi to R. (cf. [4]). However, the time series in this subfigure also show that PAR curves are only found to have some stability around dawn or dusk, whereas during the rest of the day, they are highly variable. Therefore, a proper processing will be necessary for an accurate LAI assessment. The averaged link characteristics within cluster C1 are lastly displayed in Figure 3(c) for a month at the start of the wheat growing season. LQI average transmissions from ground sensors G1 through G4 to their cluster head serve as a representation of these. (shaded in green). Moreover, the cluster head of C1 occasionally picked up transmissions of sensors from the nearby cluster C2 (blue-shaded). Our method makes it simple to divide frequency among many clusters in large-scale deployments. Our method makes it simple to divide frequencies across

several clusters in large-scale deployments, but it is not applicable here due to the sparse channel use. The LQI curves demonstrate that throughout the particular month depicted in the subfigure, connection quality were normally fluctuating and, furthermore, also declining. The crop increase is probably the cause of the later observation.

Our discussion of this intriguing relationship will include an analysis upcoming projects.



Conclusion

We described a comprehensive IoT-based agricultural monitoring system in this research. An in-situ WSN that is specifically designed for the collecting of sensor data that is of relevance to smart farming is the core part of this system. The sensor network's primary goal is to continuously assess the LAI, which is important for an accurate monitoring of crop growth processes. This sensor network is linked to a central server via a PLMN connection and a MOTT-based IoT architecture. The server is in charge of data persistence, analytics, and visualisation that farmers can utilise to aid in decision-making. We intend to expand the monitoring range of our system by including new types of environmental sensors into it in further work. This might allow for a more thorough examination, improved decision making, and the development of new agricultural uses and insights.

ACKNOWLEDGMENTS

This work was partially funded by the German Federal Ministry of Education and Research (BMBF) within the Program "Innovations for Tomorrow's Production, Services, and Work" (02K14A19K). The authors are responsible for the contents of this publication. The authors would like to thank Lars Huning, Thomas Jarmer, and Bastian Siegmann for substantial discussions and support as well as Siegfried Schittenhelm (JKI) and Dieter Trautz (AuL) for the great opportunity to realize the deployment at their experimental crop fields.

REFERENCES

[1] I. F. Akyildiz, W. Su, Y. Sankarasubramaniam, and E. Cayirci, "Wireless Sensor Networks: A Survey," Computer Networks, vol. 38, no. 4, pp. 393–422, Mar. 2002.

[2] M. H. Almarshadi and S. M. Ismail, "A Review of Wireless Sensors and Networks' Applications in Agriculture," Journal of Applied Sciences Research, vol. 7, no. 3, pp. 299–308, 2011.

[3] A. Banks and R. G. (Editors), "MQTT Version 3.1.1," OASIS Standard, October 2014. [Online]. Available: http://docs.oasisopen.org/mqtt/mqtt/ v3.1.1/os/mqtt-v3.1.1os.html

[4] J. Bauer, B. Siegmann, T. Jarmer, and N. Aschenbruck, "On the Potential of Wireless Sensor Networks for the In-Situ Assessment of Crop Leaf Area Index," Computers and Electronics in Agriculture, vol. 128, pp. 149–159, 2016.

[5] J. Bauer and N. Aschenbruck, "Measuring and Adapting MQTT in Cellular Networks for Collaborative Smart Farming," in Proc. of the 42nd IEEE Conference on Local Computer Networks (LCN), Singapore, 2017, pp. 294–302.

[6] Beecham Research Ltd., "Enabling The Smart Agriculture Revolution – The Future of Farming through the IoT Perspective," Tech. Rep., 2016.

[7] Food and Agriculture Organization (FAO) of the United Nations, "How to Feed the World in 2050," 2009. [Online]. Available: http://www.fao.org/fileadmin/templates/wsfs/do cs/ expert paper/How to Feed the World in 2050.pdf [8] R. Hartung, U. Kulau, B. Gernert, S. Rottmann, and L. Wolf, "On the experiences with testbeds and applications in precision farming," in Proc. of the 1st ACM International Workshop on the Engineering of Reliable, Robust, and Secure Embedded Wireless Sensing Systems (FAILSAFE), Delft, Netherlands, 2017, pp. 54–61.

[9] L. Huning, J. Bauer, and N. Aschenbruck, "A Privacy Preserving Mobile Crowdsensing Architecture for a Smart Farming Application," in Proc. of the 1st ACM Workshop on Mobile Crowdsensing Systems and Applications (CrowdSenSys), Delft, Netherlands, 2017, pp. 62–67.

[10] I. Jonckheere, S. Fleck, K. Nackaerts, B. Muys, P. Coppin, M. Weiss, and F. Baret, "Review of methods for in situ leaf area index determination: Part I. Theories, sensors and hemispherical photography," Agricultural and Forest Meteorology, vol. 121, no. 1–2, pp. 19–35, 2004.

[11] K. Langendoen, A. Baggio, and O. Visser, "Murphy loves potatoes: experiences from a pilot sensor network deployment in precision agriculture," in Proc. of 20th Int. Parallel and Distributed Processing Symposium (IPDPS), Rhodes Island, Greece, 2006.

[12] J. M. Lowenberg-DeBoer, "The precision agriculture revolution: Making the modern farmer," Foreign affairs, vol. 94, pp. 105–112, 2015.

[13] A.-u. Rehman, A. Z. Abbasi, N. Islam, and Z. A. Shaikh, "A Review of Wireless Sensors and Networks' Applications in Agriculture," Computer Standards & Interfaces, vol. 36, no. 2, pp. 263–270, 2014.

[14] B. Schutz, J. Bauer, and N. Aschenbruck, "Improving Energy Efficiency" of MQTT-SN in Lossy Environments Using Seed-based Network Coding," in Proc. of the 42nd IEEE Conference on Local Computer Networks (LCN), Singapore, 2017, pp. 0–0.

[15] D. Vasisht, Z. Kapetanovic, J.-h. Won, X. Jin, R. Chandra, A. Kapoor, S. N. Sinha, M. Sudarshan, and S. Stratman, "Farmbeats: An IoT Platform for Data-driven Agriculture," in Proc. of the 14th USENIX Conference on Networked Systems Design and Implementation (NSDI'17), Boston, MA, USA, 2017, pp. 515–528.