Towards an Agricultural IoT-Infrastructure for Micro-climate Measurements

Christian Hirsch, Radu Grosu
Technische Universität Wien
Computer Engineering / Cyber-Physical Systems
1040 Wien, Austria
Email: {christian.hirsch,radu.grosu}@tuwien.ac.at

Markus Redl
University of Natural Resources and Life Sciences, Vienna
Department of Crop Sciences / Division of Plant Protection
1180 Wien, Austria
Email: markus.redl@boku.ac.at

***Abstract*—Modeling, analysis, and control of crop diseases in a minimally invasive way continues to be a grand challenge in agriculture. While some diseases are well understood and possess accurate prediction models, there are many other diseases where not enough information is available yet. This precludes the construction of predictive models, and therefore the protection of such plants. In order to study these diseases it is necessary to acquire sufficient information from the field, that is, the plants' environment and the plants themselves. This in turn allows one to determine the environmental factors influencing or triggering a disease, and the way the disease spreads among the plants. We tackle this problem by building an agricultural Internet of Things (IoT) infrastructure which we call AIoT, for measuring the micro-climate in vineyards, which abound around Vienna. In this paper we present the overall architecture of AIoT, how it combines the swarm, the fog and the cloud, and describe in some detail the sensors we have developed for the swarm.**

I. Introduction

Crop diseases still represent a major risk for farmers to losing significant parts of their fields' yield. Unfortunately, for many important diseases, there is still not enough information available yet, for predicting and controlling these diseases. The farmers broadly treat their crops with pesticides, although the plants may not get infected at all. This is both expensive and unnecessarily harmful for the environment.

To overcome this problem, Agricultural Internets of Things (AIoTs) hold a great promise for collecting the missing data, and constructing accurate prediction models for the diseases of interest. In combination with weather stations and weather forecast, farmers can be alerted about impeding diseases, and thus reduce the type and amount of pesticides used.

A disease where the time of infection and outbreak, in connection with the weather, is still not known yet, is the powdery mildew in vineyards. In cooperation with the University of Natural Resources and Life Sciences, Vienna (BOKU Wien), we aim to build an AIoT infrastructure, for measuring the micro-climate in vineyards. BOKU Wien has 50 grapevines in their greenhouse and a vineyard of approx. 0.5ha at the outskirts of Vienna. In order to study the powdery mildew on grapevines the following environmental factors are of interest:

- Air temperature (at various locations).
- Relative humidity (at various locations).
- Wind speed (at various locations).
- Wind direction (at various locations).
- Leaf wetness (at various locations).
- Soil moisture (at various locations).
- Bark wetness (at various locations).
- UV radiation (UV-A, UV-B).
- Photon flux density (at various locations).
- Diameter of berries (at various locations).
- Temperature of berries (at various locations).

To study the primary infection of grapevines, the occurrence, and the spread of the powdery mildew has to be recorded. In order to accomplish that, it is necessary to first detect the primary infected grapevines in vineyards. Then, the location and the type of infection need to be evaluated/classified. To create a model that describes how the disease grows it is necessary to record the disease spread. Information that can help in this regard is as follows:

- Leaf mass and growth, stadium of phenolical growth.
- Splash effect, i.e., force of rain drops falling on the leaf.
- Spores concentration in the air (asco and conidia).

Finally, it is very important to investigate the proneness of grapevines to getting infected by the powdery mildew. The following information is necessary for this investigation:

- The density of berries on grapes.
- The color of leafs and berries.
- The thickness of the epicuticular waxes on the berries.
- The thickness of the cuticula of the berries.
- The lignin, cellulose and nitrogen content of berries.
- The size and the number of cells in the skin of the berries.
- The cell-wall thickness.
- The rate of transpiration of water.
- The rate of photosynthesis.
- The water content of berries.
- The ripeness of berries:
  - The rate of glucose and fructose.
  - The pH value.
  - The tartaric acid.
  - The malic acid.
  - The phenols.

The main contributions of this paper are as follows. First, we give an overview of the challenges faced in controlling the powdery mildew in vineyards. Second, we introduce a way to

overcoming them by designing an AIoT infrastructure. Third, we describe the already implemented parts of our AIoT. Our overarching goal is to creating a prediction model that allows us to control the powdery-mildew disease in vineyards.

This rest of the paper is structured as follows. Section II discusses related work. Section III describes our AIoT architecture and the ideas behind it. Section IV describes the already implemented parts of our AIoT. Finally, Section V concludes this work and discusses future work directions.

II. Related Work

Sensors measuring factors responsible for the health of a plant are available on the market. These are designed either for home use or for professional use. In the first case, they are usually cheap, but do not allow the access to their hardware and software. This is problematic, as one cannot reprogram them for scientific purposes. In the second case, although the products are more expensive, the hardware and software is still not accessible for scientific purposes.

One home-use example is the helloplant smart-plant sensor [1]. This measures the soil moisture, as well as the sun light in order to tell the user, via a smartphone app, when to water the plant, and where to place it. A professional-use example is the fieldmate and the soilmate from AppsforAgri B.V. [2], [3]. The fieldmate measures air temperature and humidity at two levels, as well as soil temperature at two levels. The measurements are transmitted to a cloud service, where the user can access the data. The user also gets forecast information, and the smartphone app tells the user when to employ pesticides to protect the plants.

Another trend in precision farming is to using Artificial Intelligence (AI), in order to improve the monitoring and the prediction techniques employed, and as a consequence to increase the health of animals and crops. Companies raised 800M USD in the last five years for improving robotics and machine-learning algorithms that are applied to the agricultural sector [4]. CB Insights splits the tech-companies working in agriculture, into the following categories: satellite-imagery analysis, in-field monitoring, crop and soil-heat assessment, agricultural robotics, and predictive analytics [4].

The work in [5] introduces a set of ontologies for describing sensors, actuators and samplers, including their observation, actuation and sampling activities. It was specially designed with the IoT in mind. The work in [6], is a good overview of the IoT. The authors show and compare various topologies that can be used for an IoT infrastructure. Furthermore, they describe in their work, all the necessary parts, means of communication and protocols currently used in this domain.

III. Architecture

In order to properly design the architecture of an AIoT for micro-climate measurements, it is first necessary to organize the sensors into classes and types. Furthermore, an appropriate topology for the AIoT network has to be designed.

A. Sensors

As described in Section I, there are many factors that can influence and trigger the powdery mildew. All these factors need to be measured with sensors. To get a better insight in the measurements to be effectuated, we first organize the sensors into the following three classes:

- Environmental sensors
- Imaging sensors
- Chemical sensors

In some cases multiple classes may apply to a single sensor.

Environmental Sensors: These are sensors measuring the environment in which they are deployed. This includes measurements that are typically done by weather stations, such as, temperature, humidity, atmospheric pressure, or CO2 equivalents. Because of the big availability and the low price of such sensors, they are suited to be placed all over a field.

Imaging Sensors: These are sensors that create a two (or three) dimensional matrix of values, representing an image. A camera is a typical example, measuring visible light, and creating a two dimensional image. More sophisticated cameras measuring invisible light, such as ultra-violet or infrared, or measuring some specific range of the light spectrum, are also possible. Cameras that measure multiple narrow light spectra are called multi- or hyper-spectral cameras, depending on the amount of spectra they are measuring. This kind of sensors are well suited to measure size, and to detect and classify objects. Infrared cameras can also be used to measure object temperature. Each pixel of the resulting image represents a particular temperature value.

Chemical Sensors: These sensors measure chemical substances, and are usually more expensive than environmental sensors. One reason for their higher cost is that they used less than the environmental sensors. As a consequence of this fact, it is in general too expensive to massively deploy them in the fields. However, it is still possible to mount these more expensive sensors on a robot or a drone, building this way a mobile sensing station. This station can then be driven through the field and take the required measurements.

B. Infrastructure and Communication

The basic AIoT infrastructure is shown in Fig. 1. It can be divided into three parts: The swarm shown in Fig. 1(a), the fog shown in Fig. 1(b) and the cloud shown in Fig. 1(c).

The swarm: This consists of set of sensing and actuation nodes, deployed at sufficiently many locations in space and time, such that the data they collect or act upon, enables the prediction and control of the considered disease. A swarm node may consist of a micro controller, one or multiple sensors and/or actuators, and a communication unit. This allows a swarm node to perform simple processing of the sensed data, and to specialize its control activities. It also allows it to protect the data through encryption and decryption.

The fog: This is the link between the swarm and the cloud. It directly communicates with the swarm nodes and transfers the data to the cloud, where it is stored. A fog node typically

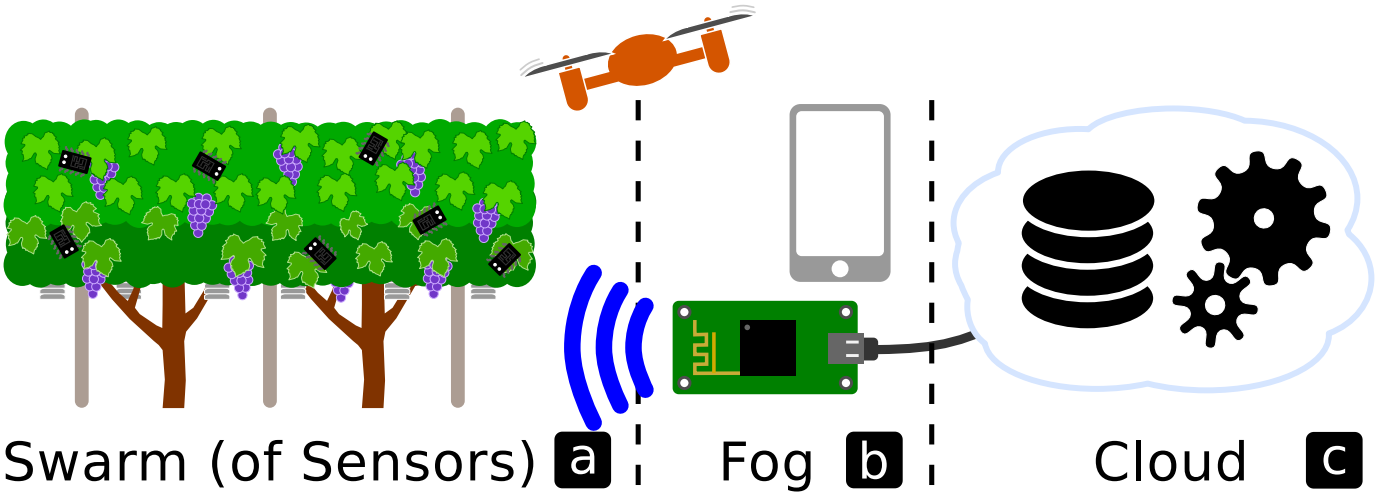

Fig. 1: Architecture, a) swarm, b) fog, c) cloud

consists of a more powerful computation unit, one or multiple units to communicate with the sensors, such as Bluetooth Low Energy (BLE), and with the cloud, such as a communication unit which is connected to the Internet.

With a more powerful computation unit, fog nodes are able to further process sensor data. This enables a fog node to filter data, set control actions (based on this data), etc. Some actions might be a trigger for an actuator activity. For example, starting and stopping an irrigation system, or sending notifications to a user. As fog nodes are in general directly connected to swarm nodes, they can react very fast. Fog nodes can either be simple computers, Single Board Computers (SBCs) or mobile robots, that communicate with the sensors and transmit the data to cloud. Depending on the swarm's means of communication, smart-phones and tablets can also be used to receive and forward data from the sensors to the cloud. The fog nodes can also deploy more sophisticated security mechanisms, and manage the identities of the swarm nodes.

The cloud: This provides services for users and machines. One major service is data storage. It is important to safely store the received sensor data in a database, and to offer services to for further processing of the data, typically through Application Programming Interfaces (APIs). Exploiting these APIs other programs can get data from the cloud, store it, and visualize it in a (web based) User Interface (UI) or smartphone app. The cloud can also provide machine-learning tools for analyzing the measured data, and build for example, prediction models based on the content of this data.

Time and Time-synchronization: Measurements recorded by a swarm node need a time-stamp in order to be valuable. Otherwise, they fail to provide useful information about system behavior: A measurement can be current, taken minutes ago, or even hours ago (cf. real-time systems [7]) To properly describe a system, it is necessary to know when a value has been measured. Time-stamping values can be done by a swarm node itself, or it can be done in the fog layer. However, this is only possible if the value the fog receives, was measured right away. If a swarm node has to store an entire sequence of measurements over time, and send the sequence to a fog node later on, then the swarm node might want to time-stamp its measurements. In general, the minimal amount of information a swarm node has to provide to a fog node is the time difference between successive measurements. The fog node can then compute the correct time stamps.

Consider now time synchronization in more detail [8]. Many electronic components have tolerances due to the nature of manufacturing. Consequently, they do not behave exactly the same way. This leads to micro-controllers having clocks that do not tick synchronously. The result is a swarm of nodes whose times drift away from the real time and from each other. If this drift sums up, measurements lose their value because they do not represent the sensor value at the time they claim. This can give wrong information to users and to machine learning tools and furthermore result in wrongly made decisions or models. It is therefore necessary to provide synchronization tools ensuring that the time at which measurements were taken is correct. Clock synchronization has to be done among both swarm nodes and fog nodes. Fog nodes can have a leading role to providing a master-clock to the swarm.

Computation: The large number of interconnected swarm and fog nodes, with quite powerful computation units, can be used to performing analysis, compression, filtering, and sensor fusion, before the data reaches the cloud. This computation is referred to as *Edge-computing* in [9]. At the swarm level, micro controllers are much more powerful than to just reading out a sensor's value. Their computation unit can be used to filter and pre-process measurements and it can also be used to do sensor fusion. This is useful in reducing the amount of data that needs to be stored on the swarm node and that needs to be transmitted. Computation, however, needs power which means that it comes to a cost of the battery life. Hence there is a trade-off between data pre-processing at a swarm node, and the node's lifetime or need for maintenance. Fog nodes are typically machines that have much more computational power

than the swarm nodes. Their lifetime is also not crucial as they are normally not powered by batteries. Thus fog nodes are perfectly suited for first analysis and a first response tasks, if necessary. The fog can act in a very fast manner, because it is located near the swarm, and directly communicates with its nodes. This makes the fog an ideal architectural component of the AIoT for time-critical tasks. These include simple control tasks, sending messages to users, or commands to actuators, based on the measurements of the swarm nodes.

At the cloud level, the IoT infrastructure has most of its computational power. While one major service of the cloud is data storage, another one is data analysis. Because almost no control tasks will be performed at this stage, real-time response in the cloud not crucial. The analysis of the data however is important, for learning various correlations and certain system and environment behaviors. The cloud supports planing, and the automatic synthesis of prediction models.

C. Further Challenges

The challenges below are not necessarily related to the successful creation of a disease-prediction model. However, their solution leads to a much more versatile infrastructure.

Security: This is one of these challenges. It is very important to protect the data that is exchanged between the different layers from being hijacked by unauthorized people. Modern micro-controllers, especially those used for IoT purposes, typically contain encryption units to securely encrypt data. But also physical security is important. Swarm or fog nodes, such as those that are placed all over a field, need to be securely placed to protect them from being stolen. Cloud nodes, are typically placed in data centers, where high security standards are common. But also at this place it is important to give the access rights only to those people who really need it.

Localization: This is another very important challenge. For swarm nodes that are physically spread all over a certain place, it is important for further processing to know where the swarm node measures its value. This can either be done manually, by assigning the location to the swarm node. But in case of several hundreds or thousands of such nodes, this might end up in a lot of manual work. Especially if it is supported to relocate the swarm nodes. Placing localization chips, such as GPS, on each swarm node is not an option, because of the increasing costs and the additional power consumption. For these cases it necessary to create a self-localizing swarm network.

IV. Implementation

Power consumption: This is a main focus in the design of a swarm node. As swarm nodes will be mounted all over the place in an outdoors field, where no power supply is to be typically expected, the nodes have to run on batteries. Moreover, recharging or changing batteries involves manpower, that should be reduced to a bare minimum. Hence, swarm nodes have to run as long as possible on a single battery.

One swarm node should preferably run at least one crop season without a need to recharge the battery. To achieve this goal, it is necessary to use low-power micro-controllers, communication units and sensors. Software also plays a major role in power consumption. The micro-controllers can be put to sleep, where power consumption is reduced to a bare minimum, but during this time, they can not do any calculations, any communication, or measurements. So there is a trade-off between doing actual work and battery lifetime. It is also important to measure the battery's status. This is crucial for taking steps to recharging nodes on time, and ensuring this way a continuous measurement of the system.

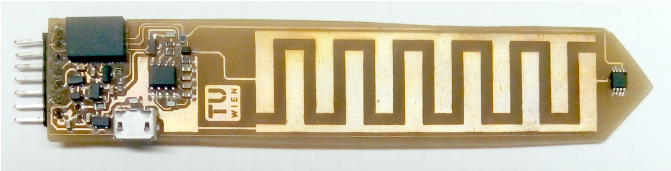

(a) Swarm node to measure soil temperature and moisture.

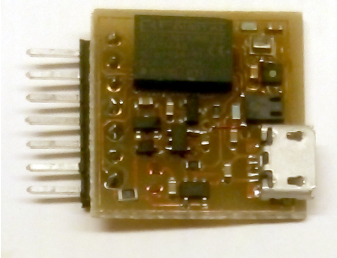

(b) Swarm node to measure environmental influences in the air.

Fig. 2: The developed swarm nodes based on the Simblee SoC.

Swarm-node design: As shown in Fig. 2, this is based on the Simblee System-on-Chip (SoC), with an ARM-Cortex M0 microprocessor, bluetooth low energy (BLE) communication stack, and integrated BLE antenna.[1] These nodes are equipped with a battery manager, for both loading the battery from a 5V power supply, and for monitoring the battery-charge level. We have currently designed two types of swarm nodes.

The first type of swarm nodes is shown in Fig. 2a. Its purpose is to measure the soil-moisture and the soil-temperature. The soil-moisture is measured using the capacitive method, where two isolated copper areas are exposed to a material that changes the capacitance of the copper areas. Between the copper areas and the material there is a small layer of protective varnish, so that the copper areas do not touch the material, to prevent the electrical connection between the two layers. With a clock circuit that alternates the voltage exposed to the copper capacitor, a voltage forms at the end, that can be measured by an Analog-Digital Converter (ADC). The lower the voltage is, the more wet material is exposed to the copper areas. This method is superior to methods where the resistance between two electrodes, that are placed in the material, is measured, because it does not oxidize and thus does not change its values over time. The soil-temperature is measured with an DS18B20 temperature-measuring unit.

[1] https://www.simblee.com/

The second type of swarm node is shown in Fig. 2b. It measures environmental temperature, relative-humidity, atmospheric-pressure, and air-quality. The sensor used to measure temperature and relative-humidity is the ENS210 by ams AG.[2] To measure atmospheric-pressure (and temperature), the BMP280 by Bosch sensortec[3] is used. For air quality measurements the CCS811 by ams AG sensor is used. Although the CCS811's application is indoor air quality measurements, we will use it also for outdoors for evaluation purposes.

Fog-node design: Fog nodes are based on the Raspberry Pi[4] SBC. They have WiFi and Ethernet-network capabilities, as well as, a BLE communication stack. Furthermore, they are cheap, and run Linux based OS which makes them a perfect match for a fog node in our simple AIoT infrastructure.

Cloud-node design: Finally, the cloud service runs on a permanently available server or personal computer. The software used for the cloud service is the open source project *ThingsBoard*[5]. This software already comes with the feature of receiving messages from sensors via various protocol types. It also stores the received data in a database. Furthermore, the software comes with various web-based tools for visualizing the data for a user. Thingsboard also offers an API, giving access to the data to other software and services.

V. Conclusion and Future Work

The work in this paper has introduced the main design principles for developing an AIoT infrastructure measuring micro-climate parameters, with the goal of creating disease-prediction models. The importance of having an infrastructure with many nodes and the opportunities it opens were discussed where appropriate. Although this work is based on the idea of building an AIoT infrastructure, it can also be used for other sectors, such as smart homes and smart buildings.

The swarm nodes we have built will have to be thoroughly evaluated and improved, before their outdoor placement. Additional swarm-node types will also have to be developed, for measuring other environmental influences, that are important to study the powdery mildew. After equipping the vineyard with a large number of swarm nodes, we will start to collect data of at least one season. The harvested data will be thereafter used to feeding machine learning algorithms that will learn various correlations among the data in a first step, and disease-prediction models in a second step.

References

[1] "helloplant smart plant sensor," http://helloplant.eu/, last accessed: 2018-02-14.

[2] AppsforAgri B.V., https://www.smartfarm.nl/en/fieldmate, last accessed: 2018-02-14.

[3] ——, https://www.smartfarm.nl/en/soilmate, last accessed: 2018-02-14.

[4] CP Insights Research, "Ai, robotics, and the future of precision agriculture," https://www.cbinsights.com/research/ai-robotics-agriculture-tech-startups-future/, 2017, last accessed: 2018-02-14.

[5] A. Haller, K. Janowicz, S. Cox, M. Lefrançois, K. Taylor, D. L. Phuoc, J. Lieberman, R. García-Castro, R. Atkinson, and C. Stadler, "The sosa/ssn ontology: A joint w3c and ogc standard specifying the semantics of sensors, observations, actuation, and sampling," preprint on webpage at http://www.semantic-web-journal.net/content/sosassn-ontology-joint-w3c-and-ogc-standard-specifying-semantics-sensors-observations.

[6] A. Al-Fuqaha, M. Guizani, M. Mohammadi, M. Aledhari, and M. Ayyash, "Internet of things: A survey on enabling technologies, protocols, and applications," *IEEE Communications Surveys Tutorials*, vol. 17, no. 4, pp. 2347–2376, Fourthquarter 2015.

[7] H. Kopetz, *Real-time systems: design principles for distributed embedded applications*. Springer Science & Business Media, 2011.

[8] A. Ceccarelli, F. Brancati, B. Frömel, and O. Höftberger, *Time and Resilient Master Clocks in Cyber-Physical Systems*. Springer, Cham, vol. 10099.

[9] W. Shi, J. Cao, Q. Zhang, Y. Li, and L. Xu, "Edge computing: Vision and challenges," *IEEE Internet of Things Journal*, vol. 3, no. 5, pp. 637–646, Oct 2016.

[2] http://ams.com

[3] https://www.bosch-sensortec.com

[4] https://www.raspberrypi.org/

[5] https://thingsboard.io/