

FarmBeats: An IoT Platform for Data-Driven Agriculture

Deepak Vasisht², Zerina Kapetanovic³, Jong-ho Won⁴, Xinxin Jin⁵

Ranveer Chandra¹, Ashish Kapoor¹, Sudipta Sinha¹, Madhusudhan Sudarshan¹, Sean Stratman⁶

¹Microsoft, ²MIT, ³University of Washington, ⁴Purdue University, ⁵UCSD, ⁶Dancing Crow Farm

Abstract – Data-driven techniques help boost agricultural productivity by increasing yields, reducing losses and cutting down input costs. However, these techniques have seen sparse adoption owing to high costs of manual data collection and limited connectivity solutions. In this paper, we present FarmBeats, an end-to-end IoT platform for agriculture that enables seamless data collection from various sensors, cameras and drones. FarmBeats’s system design that explicitly accounts for weather-related power and Internet outages has enabled six month long deployments in two US farms.

1 INTRODUCTION

The demand for food is expected to double by 2050, primarily fueled by an increase in population and upward social mobility [58]. Achieving this increase in food production is even more challenging because of receding water levels, climate change and shrinking amount of arable land. According to International Food Policy Research Institute, data-driven techniques can help us achieve this goal by increasing farm productivity by as much as 67% by 2050 and cutting down agricultural losses [20].

In fact, field trials have shown that techniques that use sensor measurements to vary water input across the farm at a fine granularity (precision irrigation) can increase farm productivity by as much as 45% while reducing the water intake by 35% [3]. Similar techniques to vary other farm inputs like seeds, soil nutrients, etc. have proven to be beneficial [25, 37]. More recently, the advent of aerial imagery systems, such as drones, has enabled farmers to get richer sensor data from the farms. Drones can help farmers map their fields, monitor crop canopy remotely and check for anomalies. Over time, all this data can indicate useful practices in farms and make suggestions based on previous crop cycles; resulting in higher yields, lower inputs and less environmental impact.

While these techniques for agriculture have shown promising results, their adoption is limited to less than 20 percent farmers owing to the high cost of manual sensor data collection (according to US Department of Agriculture [30]). Automating sensor data collection requires establishing network connection to these sensors. However, existing connectivity solutions [11, 18] require a cellular data logger to be attached to each sensor (see Table 1 for a detailed comparison). These loggers cost around \$1000 each in equipment cost with over \$100 in subscription fee. Further, they are limited in the amount of data that

they can send to few kilobytes per day. Clearly, these solutions do not scale up for large farms and cannot support high bandwidth sensors like cameras and drones, which rely on sending all their data to the cloud for processing [10, 49]. This situation is further worsened by the fact that farms typically have limited cellular coverage [24] and are prone to weather-based Internet outages.

In this paper, we present FarmBeats, an end-to-end IoT platform for data-driven agriculture, that enables seamless data collection from various sensor types, i.e., cameras, drones and soil sensors, with very different bandwidth constraints. FarmBeats can ensure system availability even in the face of power and Internet outages caused by bad weather; scenarios that are fairly common for a farm. Further, FarmBeats enables cloud connectivity for the sensor data to enable persistent storage as well as long-term or cross-farm analytics. We have deployed FarmBeats in two farms in the US over a period of six months and used FarmBeats to enable three applications for the farmer: precision agriculture, monitoring temperature and humidity in food storage, and monitoring animal shelters. In designing FarmBeats, we solve three key challenges.

First, to enable connectivity within the farm, FarmBeats leverages recent work in unlicensed TV White Spaces (TVWS) [6, 16, 44] to setup a high bandwidth link from the farmer’s home Internet connection to an IoT base station on the farm. Sensors, cameras and drones can connect to this base station over a Wi-Fi frontend. This ensures high bandwidth connectivity within the farm. However, due to the lack of power on the farm, the base station is powered by battery-backed solar power which suffers from power unreliability depending on weather conditions. As shown in past work [22, 51], cloudy weather can reduce solar power output significantly and drain the batteries of the base station to shut it down. To solve this problem, FarmBeats uses a novel weather-aware IoT base station design. Specifically, it uses weather forecasts to appropriately duty cycle different components of the base station. To the best of our knowledge, this is the first weather-aware IoT base station design.

Second, Internet connection to the farm is typically weak making it challenging to ship high bandwidth drone videos (multiple GBs) to the cloud. Furthermore, farms are prone to weather-related network outages that last weeks. Such system unavailability impedes a farmer’s

ability to take adequate preventive actions, do UAV inspections and leads to loss of valuable sensor data. Thus, FarmBeats uses a Gateway based design, wherein a PC at the farmer’s home serves as a gateway for the farm data. The FarmBeats Gateway serves two purposes: a) it performs significant computation locally on the farm data to consolidate it into summaries that can be shipped to the cloud for long-term and cross-farm analytics, and b) the gateway is capable of independent operation to handle periods of network outage, thus leading to continuous availability for the farmer.

Finally, while drones are one of the most exciting farm sensors today, they suffer from poor battery life. Getting aerial imagery for a farm requires multiple drone flights and a long wait time in between when the batteries are being charged. We use the fact that farms are typically very windy, since they are open spaces. Thus, we incorporate a novel path planning algorithm in the FarmBeats gateway, that leverages wind to help the drone accelerate and decelerate, thereby conserving battery. This algorithm is motivated by how sailors use winds to navigate sailboats.

We use the FarmBeats system to enable precision agriculture applications on two farms: one in Washington state and the other in upstate New York. While traditional farming treats the farm as a homogeneous piece of land, precision agriculture adapts the farm inputs over different parts of the farm depending on the requirement. Precision agriculture techniques require a precision map with information about each location in the farm, for example, the soil temperature, soil moisture, nutrient levels, etc. To construct this precision map, existing solutions for precision agriculture require a dense deployment of in-ground sensors [30]. A dense deployment of sensors becomes expensive (as well as cumbersome to manage) as the size of the farm grows. Unless these sensors are deployed densely within a farm, the estimated precision map can be very inaccurate, as we show in Section 7. Since FarmBeats’s gateway has access to both the drone videos and sensor data, it enables a novel low-cost mechanism that uses drone videos in combination with sparse ground sensors to generate precision maps for the farm. To the best of our knowledge, this is the first system that can combine the temporal data from sensors, with the spatial data from drones to construct an instantaneous precision map of the farm, such as the one in Figure 5.

Beyond FarmBeats’s application in precision agriculture, farmers have so far used FarmBeats for two other applications. First, the farmers have been using FarmBeats to monitor temperature and humidity in storage spaces to ensure that the produce does not go bad. Second, the farmers have plugged in cameras at different locations, to monitor cow sheds, selling stations etc¹.

¹Supplementary Material includes detailed description of FarmBeats applications and usage.

Contributions: To summarize, FarmBeats makes the following key contributions:

- **Long-term large scale deployment:** Our deployments have run over 6 months in each of the farms and collected over 10 million sensor measurements, 1 million camera images and 100 drone videos
- **Novel Weather-Aware IoT Base Station Design:** Adding weather awareness into the IoT base station reduced the base station down time to zero as opposed to greater than 30% downtime during the same month in the previous year in an earlier version of our deployment
- **Novel Inference Techniques for Compression of Aerial Imagery Data:** FarmBeats’s gateway achieved a median compression of 1000 times from an aerial drone video to the sensor summaries sent to the cloud. Further, the gateway remained available even when the Internet connectivity to the farm faced a week-long outage
- **Wind-Assisted Drone Flight Planning Algorithm:** FarmBeats’s flight planning algorithm improves the area covered by a single drone flight by 30%

2 IOT PLATFORM: OBJECTIVES

In building FarmBeats, we target the following goals:

- **Availability:** The platform should have negligible downtime. When there is an outage (for example, due to power or network failure), data collection from the sensors should not stop and the platform should continue to deliver services to the farmers.
- **Capacity:** It should support sensors with widely varying requirements: pH sensors reporting few bytes of data to drones sending gigabytes of video. Similarly, the system should be capable of supporting end-user applications with varying needs: from a precision irrigation application that needs the latest sensor data for the entire farm to a crop suggestion application that needs just high level productivity data but across several growing seasons
- **Cloud Connectivity:** Several farming applications, such as crop cycle prediction, seeding suggestions, farming practice advisory, etc. rely on long term data analytics. Besides, a farmer may want to access some applications even when he is not on the farm. Thus, the IoT platform must enable pushing data to the cloud.
- **Data Freshness:** Stale sensor data from the farm can make applications suggest incorrect courses of action to the farmer. Gaps in historical data can also cause applications to misbehave. Moreover, stale data leads to bad user experience. Thus, the platform must strive to maintain maximum data freshness.

3 THE FARMBEATS IOT PLATFORM

While these objectives have been fairly successfully achieved by home IoT platforms like Amazon Echo, achieving these objectives in an agricultural setting introduces several challenges for two main reasons: access and environmental variability. As discussed before (and

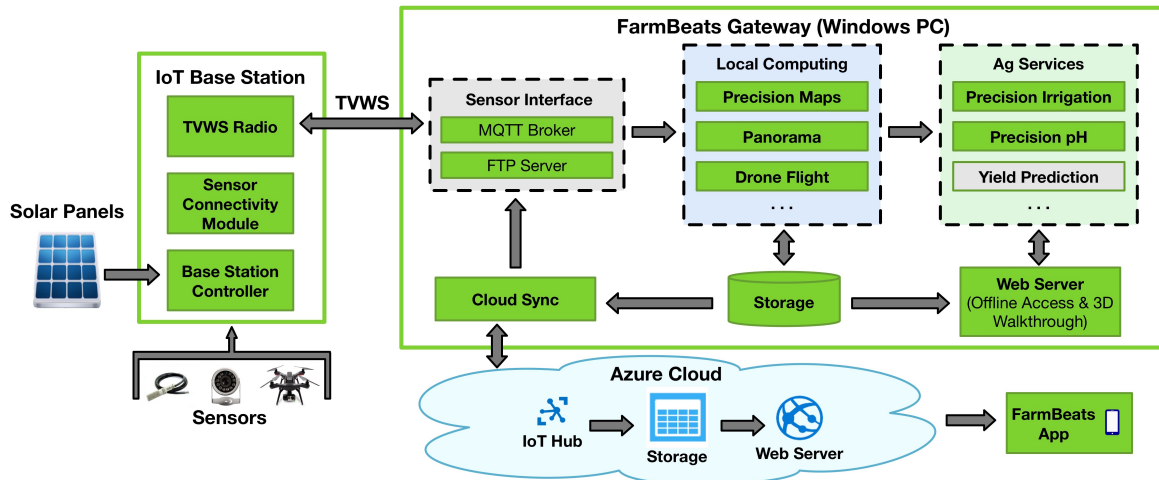


Figure 1: FarmBeats System Overview

as shown in Table 1), farms do not have access to power and high-bandwidth Internet connectivity unlike indoor IoT systems. Furthermore, energy harvested from the environment and weak network connectivity to the farm is susceptible to failures due to weather variability. So, the key question for the design of FarmBeats is: how does one design an IoT platform to meet the objectives in a highly variable, resource constrained environment?

3.1 Design Decisions

An overview of the system is given in Figure 1. Here, we discuss the main design decisions.

To achieve farm connectivity over long range, we leverage recent work in the TV White Spaces [6, 16, 44] to setup a high-bandwidth connection from the farmer’s home to the farm. However, sensors, drones and cameras typically do not support TVWS. Thus, in order to maintain compatibility with sensors along with long-range high bandwidth connectivity, we deploy a two-layer hybrid network. We use a TVWS link to connect the farmer’s home Internet connection to a few IoT base stations on the farm. Since it is a high bandwidth backhaul link, each base station can accommodate sensors, as well as cameras and drones. At the second layer, the IoT base station provides a Wi-Fi interface for connections from sensors and other devices. The Wi-Fi interface ensures that the farmer can not only connect most off-the-shelf farming sensors, cameras and drones; but they can also use their phone to access farming productivity apps.²

Variability in harvested solar energy leads to IoT base station downtime in overcast conditions. In fact, in our early deployments, power failures due to environmental factors were the major cause of unavailability. While past work has dealt with this problem in the context of single sensors [22, 51, 59] by duty cycling the sensors, the same

²Future iterations of the systems would add multiple interfaces to the base station to enable compatibility with more sensor types.

approach does not work for a base station. Specifically, the base station has multiple components with different power requirements and duty cycling costs. For example, a farmer is typically inactive at night and is unlikely to check the farm data. So, turning the TVWS device off (which consumes 5x more power than the rest of the base station) can enable the base station to collect data (in a cache) from the sensors more frequently. Further, FarmBeats enables the farmer to turn the base station on to access Wi-Fi for productivity applications, while they are on the farm. This adds another layer of uncertainty in the duty cycling plan. Thus, we propose a novel duty cycle policy (in Section 4) wherein the different components of the base station are duty cycled at different rates; while explicitly accommodating these constraints.

Finally, given the weak internet connectivity to the farm, a naive approach of pushing all the data to the cloud does not work. We make the key observation that the data requirements of the farming applications can be broadly classified into two main categories: immediate detailed data and long-term summarized data. Table 2 summarizes how the industrial and research applications of farm data can be classified into these two categories. This categorization enables a gateway based IoT design for FarmBeats. The local gateway sits at the farmer’s home at the other end of the White Space link and performs two functions: a) creates summaries for future use and ships them to the cloud and b) delivers applications that can be provided locally. The summaries are several orders of magnitude lower in size than the raw farm data (3-4 orders of magnitude smaller in case of the precision agriculture application discussed later) and hence, respect the harsh bandwidth constraints.

3.2 Architecture

The FarmBeats system has the following components: **Sensors & Drones:** FarmBeats uses off-the-shelf sensors for its applications. Each sensor measures specific

Technology	Cost	Data Restriction
Cellular Connection (Decagon Devices)	Per sensor fee: 1000\$ + 100\$ annual fee	Restricted to sensor data; Uploads every 15 mins at best
Mesh Networks (Ranch Systems)	Base station: 3500\$ + 750\$ annual fee; Per sensor fee: 1100\$ + 60\$ annual fee	Maximum 25 mesh nodes per base station
Satellite (Iridium)	Per sensor fee: 800\$ + 100\$ monthly fee	Restricted to 2.5 Kbps

Table 1: Cost Comparison of Farm Sensor Networking Solutions

Data Requirement	Applications
Immediate Descriptive Data	Precision irrigation, virtual walkthroughs, productivity apps, farm monitoring, ...
Long-term Summarized Data	Crop suggestions, seed distribution, yield monitoring, financial management, animal health statistics, ...

Table 2: Application classification based on requirements characteristics of the farm, such as soil moisture and soil pH, and reports this data to the IoT base station over a Wi-Fi connection. In addition to soil sensors, FarmBeats supports cameras for farm monitoring and drones. The cameras are either connected to the IoT base station over Ethernet or report data over Wi-Fi. They take periodic snapshots and transmit this data to the IoT base station. UAV flights are either periodically scheduled or manually initiated using the FarmBeats app on the farmer’s phone. **IoT Base Station:** The IoT base station on the farm is powered by solar panels, backed by batteries and has three components:

- The TVWS device ensures that the base station on the farm can send the data to the gateway, which then, sends it up to the cloud.
- The sensor connectivity module establishes a connection between the base station and the sensors deployed on the farm. In FarmBeats’s current implementation, this module is just a Wi-Fi router.
- Finally, the Base Station Controller is responsible for two functions. First, it serves as a cache for the sensor data collected by the sensor module and syncs this data with the IoT gateway when the TVWS device is switched on. Second, it plans and enforces the duty cycle rates depending on the current battery status and weather conditions.

IoT Gateway: As mentioned before, the goal of the IoT gateway is to enable local services and create summaries from existing data to be sent to the cloud. We use a PC form factor device as the FarmBeats gateway, which is typically placed in the farmer’s house or office, whichever has Internet access. The gateway provides an interface for applications to run and create summaries to be sent to the cloud as well as to post data to the local web server. Furthermore, it includes a web service for the farmer to access detailed data when they are on the farm network. This also ensures that FarmBeats remains available even when the cloud connection is not present. Fi-

nally, it includes built-in algorithms for drone path planning and for compressing drone data before being sent to the cloud (described in Section 5). We illustrate in Section 5.3 how applications function on the gateway with the example of precision agriculture applications.

Three aspects of the FarmBeats gateway differentiate it from prior IoT gateways. First, the FarmBeats gateway implements a web service, providing unique services that are different from the FarmBeats web service in the cloud. Second, the gateway can operate offline, and still offer the most important services. Finally, as shown later in the context of precision agriculture, having access to data from multiple types of sensors enables unique feature-based summarization technologies for the drone videos and sensor data.

Services & the Cloud: The Gateway ships data summaries to the cloud, which provides a storage system for long-term data and a web interface for the farmer. The cloud enables three functions: data access outside the farm network (e.g. when traveling), long term applications like crop suggestions, and cross-farm analytics.

4 DUTY CYCLING THE BASE STATION

As discussed before, FarmBeats’s solar-powered IoT base station on the farm is duty cycled to explicitly account for weather forecasts and current charge state of the batteries. Two aspects of the base station make this problem challenging: a) The sensor connectivity module has significantly lower power requirements than the TVWS device. Thus, we need to intelligently proportion power between these components to achieve optimum performance. b) FarmBeats allows farmers to manually turn the base station on to connect to the Internet to use productivity apps on their phone. This adds a variable component to power consumption.

4.1 Duty Cycling Goals

The key goals for the duty cycling algorithm are:

- **Energy Neutrality:** Like past work in the context of duty-cycling sensors backed by energy harvesting sensor systems [22,59], FarmBeats aims to achieve the objective of energy neutrality. For a given planning period, the goal is to consume at max as much power as can be harvested from the solar panels.
- **Variable Access:** FarmBeats allows farmers to access Wi-Fi connectivity on-demand. This power consumption

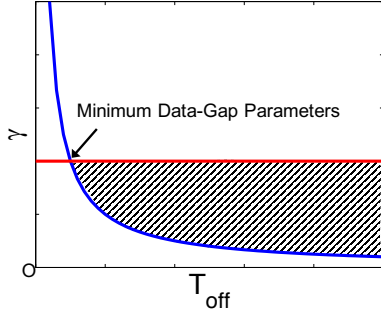


Figure 2: **Duty Cycling Approach:** The shaded region shows the feasibility region of Equations 2 and 4. The latency is minimized when both the equations are satisfied on the boundaries.

is usage-driven and varies across days. FarmBeats must plan ahead for this variable delay.

- **Minimize Data Gaps:** We use the term ‘Data Gaps’ to denote continuous time-intervals with no sensor measurements available. Such gaps need to be minimized to avoid missing out on interesting data trends. So, FarmBeats’s duty cycling algorithm aims to minimize the length of the largest data gaps, under the constraints of energy neutrality and variable access.

4.2 Power Budget

The sole power source for the base station is a set of solar panels (backed by a battery). The solar power output varies with the time of day and the weather conditions. We use standard methods [51] to estimate the output of the solar panels, given the weather conditions. Let us say that the energy output from the solar panels over the next planning period is S_I . Because the estimation is not perfect and there is usage variability, there maybe some credit or debit from the previous planning period. Let us denote this credit by C_I . So, the total power budget for the base station over the next planning period is $S_I + C_I$.

4.3 Duty Cycling Approach

The duty cycle decisions are made on the order of a planning period, T_p . Since our deployments use solar powered base stations, we set T_p to be one day. We define the average energy loss due to battery leakage and the very low power base station controller during one T_p to be E_D . For the farmer to have on-demand Wi-Fi access, we allocate a fixed time budget of T_v . If we denote the power consumption of the TVWS device by P_T and the power consumption of the sensor connectivity module by P_S , then, we need to allocate $T_v(P_T + P_S)$ for variable Wi-Fi access. Now the key question is, how do we proportion the remaining power budget?

Duty Cycling the TVWS device: The TVWS module is needed to sync the data in the base station cache with the gateway. Let us assume that we have a schedule, S , the set of sync times advised for the base-station to sync with the FarmBeats gateway. This could depend on the farmer’s

usage patterns, sensor types and can be either manually programmed or automatically inferred. The sync times in the set S have a corresponding set of weights given by set W . An example of a high-weighted sync time could be sunrise, as that is when the farmer begins their day. Thus, they would like to access the latest sensor data when the activities of the day are planned.

To ascertain the subset of syncs that need to be performed, we make a simple observation. If the sensors haven’t sent any data to the base station, the base station need not turn on the TVWS device. Specifically, it uses the following greedy algorithm to identify the syncs to be executed. Let us denote by, $S_1 \subset S$, the subset of syncs that are to be executed. This subset is initialized as an empty set. FarmBeats starts by adding the highest priority sync to S_1 . After it has done that, it subtracts $|S_1|P_T T_S$ from the power budget, where $|\cdot|$ denotes set cardinality and T_S denotes the time to perform a sync operation. Then, FarmBeats computes the corresponding duty-cycle rate for the sensor connectivity module. If this rate ensures that the second highest weighted sync in S will have additional data from the sensors to sync with the gateway, it adds this sync operation to the set S_1 . It repeats this process in decreasing order of weights until it reaches a state where one of the syncs in S_1 has no new data to share. As we add more sync operations to S_1 , the power budget for the sensor connectivity modules decreases. With a lower power budget, the sensor connectivity module can collect data from the sensors less often and hence it becomes less likely for frequent syncs to see new data. Thus, the algorithm implicitly regulates the sync operations between the gateway and the base station.

Duty Cycling the Sensor Connectivity Module: We denote the duty cycling rate for the sensor connectivity module by γ . In particular, it is turned off for a time period, T_{off} , followed by an on period of, T_{on} and $\gamma = T_{on}/T_{off}$. Using the notation we have established so far, the energy expenditure of the system is $E_D + (P_S + P_T)T_v + P_T T_S |S_1| + P_S T_p \gamma$. Since the goal of the planning algorithm is to estimate T_{on} and T_{off} such that the energy expenditure does not exceed the energy budget during the planning period, this imposes the following constraint:

$$S_I + C_I \geq E_D + (P_S + P_T)T_v + P_T T_S |S_1| + P_S T_p \gamma \quad (1)$$

$$\implies \gamma \leq \frac{S_I + C_I - E_D - (P_S + P_T)T_v - P_T T_S |S_1|}{P_S T_p} \quad (2)$$

Let us denote $T_{connect}$ as the time taken for the sensor connectivity module to turn on and establish a connection to the sensors. Further, let T_{sensor} be the time that it takes for all the sensors to wake up and transmit to the base station. Since the ON time of the module has to be long enough for the sensors to be able to communicate their data to the base station, this imposes a further constraint:

$$T_{ON} \geq T_{connect} + T_{transfer} \quad (3)$$

$$\implies \gamma T_{off} \geq T_{connect} + T_{transfer} \quad (4)$$

Since our goal is to minimize the data gap under the power constraints imposed by Equations 2 and 4, we aim to minimize T_{off} . The inequalities from Equations 2 and 4 define a convex region in the 2-dimensional space of (γ, T_{off}) , shown as the shaded region in Figure 2. Since the cost function T_{off} is linear, the minimum occurs on a corner of the intersection region defined by the two inequalities. Specifically, the minimum latency is achieved when the two inequalities are exactly met. The solution is shown graphically in Figure 2.

4.4 Discussion

At this point, it is worth noting that:

- By explicitly accounting for the credit term, C_I , the formulation absorbs the variability in on-demand Wi-Fi usage patterns. If the on-demand Wi-Fi usage patterns are stable, the term C_I goes down to zero.
- By incorporating flexibility in sync times between the gateway and the base station, FarmBeats can easily adapt to farm applications with different requirements.
- We have not yet discussed the duty cycling of sensor nodes. In our implementation, we set the duty cycle off time for sensors to be less than $T_{transfer}$ to ensure that the sensor can transfer data when the sensor connectivity module is on. An alternative implementation would allow the base station to send wake-up times to sensors. Our design choice was motivated by the availability of very low-power sensors that consume 3-4 orders of magnitude less power than the base station on average.

5 THE FARMBEATS GATEWAY

In this section, we discuss two key components of the FarmBeats gateway: UAV path planning and stitched imagery (orthomosaic) generation from UAV videos. We also illustrate how the FarmBeats gateway implements a precision agriculture application.

5.1 UAV Path Planning

Most UAVs operate in line sweep patterns. Specifically, given a sequence of waypoints defined by their GPS coordinates, they move from one waypoint to the next, in order. However, in the context of agriculture, our objective is to optimize for the area covered in a single flight. Thus, we aim to minimize the time taken to cover a given area. To that end, we make the observation that increasing the number of waypoints to cover the same area increases the time taken to cover it, even though the total path length may be the same. This is because the quadrotor has to decelerate at each waypoint and come to a halt before it can turn around and accelerate again. We present a novel flight planning algorithm that minimizes

the number of waypoints required to cover a given part of the farm.

Existing commercial systems like Pix4D [42], DroneDeploy [15], etc. offer area coverage services, these systems cover a given area using an east-to-west flight path, without any regards to the number of waypoints required. Recent research proposals like [17] do not guarantee the minimum number of waypoints either. Our area coverage algorithm *Min-waypoint* described below guarantees that the UAV covers an area with the minimum number of waypoints:

- Given an area, construct its convex hull.
- Determine the direction of sweeping lines. For each edge and its antipodal (diametrically opposite) vertex [50], draw two parallel lines and measures the distance between them. The slope of the edge corresponding to the minimum distance between the edge and the antipodal vertex becomes the direction of the sweeping lines.
- Determine the waypoints depending on the flight altitude, the camera’s field of view, and desired image quality.
- Given a start-point and end-point of the flight path, order the waypoints to minimize the total travel distance.

Adaptive Wind-assisted Yaw Control: Since farms are large open spaces and typically very windy, we observed that quadrotors that have an asymmetric physical profile can exploit the wind either for more efficient propulsion or deceleration. Figure 3(b) shows an example of a quadrotor (DJI Inspire 1) that has an asymmetrical profile, where its front and the side are considerably different; thus, it can exploit the wind similar to sailboats. Intuitively, when the quadrotor is flying downwind (i.e. wind is helping the quadrotor), the side profile of the quadrotor should face the wind since the side profile has a larger area and hence, will be able to extract the maximum assistance from the wind. In our experiments on the farm, the quadrotor requires significantly more energy (80% higher at 4m/s) to maintain its speed upwind in comparison to the downwind flight.

To leverage this observation, we designed a novel yaw control algorithm to exploit the wind energy on the farm. Specifically, yaw is the angle of the quadrotor with respect to the vertical axis. While we don’t describe the algorithm in detail, on a high level, Figure 3 describes how the yaw control algorithm would operate for a quadrotor that has a larger area on the sideways profile. For the downwind segment from the start point to the first waypoint, the adaptive control starts by making the yaw perpendicular to the flight path, thereby maximally utilizing the favorable wind as the quadrotor accelerates. However, as the velocity increases, the air drag generated by the quadrotors profile also increases. Consequently, once the quadrotor accelerates the yaw is reduced so as to maximally exploit the wind, while minimizing the parasitic drag due to the side profile. Similarly, the deceleration




Figure 3: FarmBeats’s path planning algorithm uses the asymmetry in front and side profiles of a drone like DJI Inspire 1 (in (b)) to leverage wind to its advantage

phase can very effectively exploit the air drag by making its yaw perpendicular to the flight path. This action is analogous to the action that a skier takes to stop.

5.2 Generating Orthomosaics from UAV Videos

UAVs generate a prohibitive amount of video that is difficult to transfer to the cloud due to poor network connectivity on farms. For example, a 4 minute flight with a UAV capturing 1080p video at 30 frames per second generates almost a Gigabyte of video data. We make the observation that the unit of interest for the farmer is not the drone video itself, but an overview of the farm that can be provided by a geo-referenced panoramic overview, which is one-two order of magnitude more compact than the full resolution video (see Figure 4). The stitched orthomosaic generated from the drone video provides a high resolution visual summary of the farm from a low altitude vantage point, revealing minute details. In fact, existing agricultural drone solutions ([10, 49]) ship the videos to the cloud and convert them into orthomosaics to show to the farmer. Thus, we incorporate the orthomosaics processing pipeline into the FarmBeats Gateway, to process the drone videos locally.

Broadly speaking, the panoramic views can be constructed from the UAV video using two approaches, based on either (i) aerial 3D mapping [42, 48] or (ii) image stitching and mosaicking [4, 7, 36, 53, 56]. While the aerial 3D mapping is a general-purpose method to reconstruct high resolution 3D surface maps of the environment from aerial videos, the image stitching methods treat the world as planar and simply stitch the different images together by finding their relative positions.

Computing high-resolution surface maps is both compute and memory intensive and is not suitable for the resource-constrained farm gateway. On the other hand, while image stitching methods can be incorporated into the gateway, the planar terrain assumption becomes invalid on the farm. Uneven ground geometry, trees, animals or man-made structures observed in the video generates parallax which cannot be handled by the image registration algorithms that assume a planar scene. As we show later in Section 7 and as observed in prior work [27], existing image stitchers – Microsoft ICE [36], AutoPano [4] tend to produce distorted orthomosaics in such scenarios. This presents us with an uncomfortable trade-

off: either fly high such that the farm appears planar and sacrifice fine details of the farm, or ship the large aerial videos to the cloud for processing.

Our approach: In order to break this tradeoff, we have developed a hybrid technique which combines key components from both 3D mapping and image stitching methods. On a high level, we use techniques from the aerial 3D mapping systems, just to estimate the relative position of different video frames; without computing the expensive high resolution digital surface maps. Since this process can be performed at a much lower resolution, this allows us to get rid of the harsh compute and memory requirements, while removing the inaccuracies due to non-planar nature of the farm. Once these relative positions have been computed, we can then use standard stitching software (like Microsoft ICE) to stitch together these images. The performance achievements of this hybrid approach are evaluated further in Section 7.

5.3 Generating Precision Maps

As discussed before, precision agriculture relies on accurate precision maps of the farm that indicate the distribution of a specific characteristic throughout the farm. The FarmBeats gateway naturally enables a novel approach to precision map generation that can use the aerial imagery from drones to perform spatial inference of sensor values from sparsely deployed sensors.

Specifically, FarmBeats uses the orthomosaic generated from the drone videos together with the sensor values observed by the sensors planted in the soil, and generates predictions for the entire farm. For example, sensors that observed soil temperature at the discrete locations can inform the machine learning pipeline to make predictions about every location in the farm by considering spatial proximity as well visual similarity of the locations to the sites with the sensors.

FarmBeats’s gateway embeds a machine learning pipeline that draws on probabilistic graphical models that embed Gaussian processes [43]. The key intuition in the proposed model is *spatial and visual smoothness*: areas that are similar should observe similar sensor readings. Specifically, the model relies on two kinds of similarities:

- **Visual Smoothness:** Areas that look similar have similar sensor values. For example, a recently irrigated area

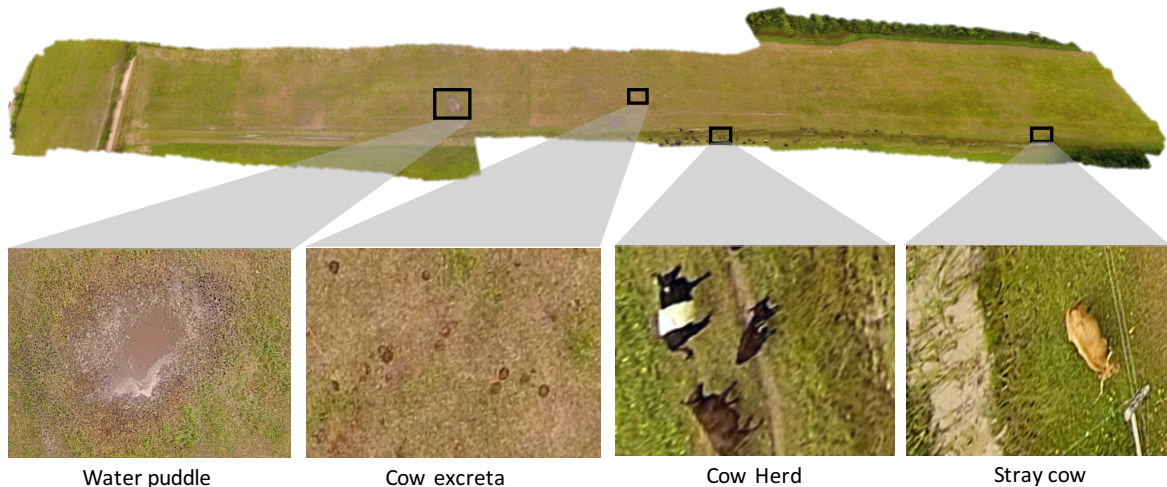


Figure 4: **Orthomosaic Generation:** The high resolution orthomosaic generated by FarmBeats for a 5 acre patch in the large farm reveals important visual details to the farmer, such as those shown in the insets – puddles that can make part of the land unavailable for agriculture, cow excreta that becomes manure and enriches the soil, location of individual cows grazing on the farm and their distance from the nearby electric fence.

would look darker and hence, has more moisture.

- **Spatial Smoothness:** Since we are measuring physical properties of the soil and the environment, the sensor readings for locations that are nearby should be similar.

We encode these two intuitions into a graphical model using standard techniques and formulate it as a Gaussian process regression model [43].

In our current design, FarmBeats uses the precision maps as units of summarization for the UAV data and ships them to the cloud. This has two advantages over the using orthomosaics as the unit of summary. First, they incorporate sensor data from the farm into drone videos. Second, they can be compressed to two to three orders of magnitude smaller size than a orthomosaic. So, while the orthomosaic is good for giving the farmer a detailed overview of the farm, precision maps are better for long term storage and shipping. We envision that for other machine learning applications as well, feature maps like the precision maps of the field would be the summaries that get shipped to the cloud, while the descriptive data delivers short-term applications on the gateway.

6 DEPLOYMENT

We deployed FarmBeats in two farms located in Washington (WA) state and in upstate New York (NY), with an area of 5 acres and 100 acres, respectively. The farmer in WA grows vegetables that he sells in the local farmers market. The farm in upstate NY follows the community supported agriculture (CSA) model, and grows vegetables, fruits, grains, as well as dairy, poultry, and meat. Our deployments consist of: sensors, cameras, UAV, the IoT base station, a gateway PC, the cloud service and a dashboard (mobile app and a web page).

Sensors: Each farm was equipped with sensors that measure soil temperature, pH, and moisture. In case of

sensors without Wi-Fi support, we interfaced them with Arduinos, Particle Photons or NodeMCUs to add Wi-Fi capability. While the exact number of sensors varied over the deployments and the application of interest, we have deployed over 100 different sensors. Additionally, We deployed Microseven IP [33] cameras in different parts of the field to monitor the farm, as well as to capture IR images of crops. To avoid potential damage from environmental impacts, each sensing platform was encased in a weatherproof box. An example of a sensor deployment can be seen in Figure 6(a).

Drones: We used the DJI Phantom 2, Phantom 3 and Inspire 1 for our drone flights.³ We created an auto-pilot application using the DJI Mobile SDK [12] to interface with FarmBeats. The user can use the app to first select the flight altitude and determine the area to be covered on an interactive map. FarmBeats’s app then plans a flight path using the algorithm proposed in Section 5.1. After the drone completes its mission, it automatically returns to its home position and transfers the video recording during the flight to the gateway, through the IoT base station.

IoT Base Station: At each IoT base station deployment, we set up a TVWS network using the FCC certified Adaptrum ACRS 2 radios [2] operating at 20 dBm, and 11 dBi directional antennas with 90 degree sectors. The internet connectivity was provided by the home internet connection of the farmers. To power the base station we setup a solar charging system, which comprised of two 60 Watt solar panels connected to a solar charge controller. The powering system is backed by four 12V-44Ah batteries connected in parallel. The power output goes through an 8-port Digital Logger PoE switch [28]. This provides us the capability to turn on or off individual components of the base station. A Raspberry Pi 3 with 64 GB SD

³We received an exemption from the FAA to fly the UAV.

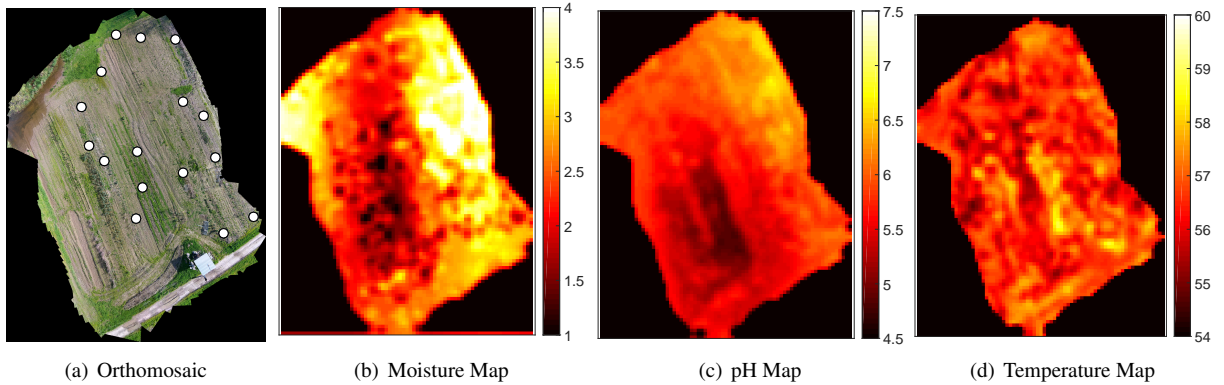


Figure 5: **Precision Maps:** (a) A 40 MPixel orthomosaic created from a 3 minute flight over 2 acre area of a farm. Our system infers dense sensor measurements from very few sensors deployed on the farm (indicated by white circles). (b) The predicted soil moisture map (our sensors measures moisture on a scale of 1 to 5). Note that the top left region in the image where the ground appears wet was correctly predicted to have high moisture even though no moisture sensors were present in that part of the farm. (c) The predicted pH map (pH is measured from 0-14, 7 is neutral and 0 is the most acidic). Our system identified that the whole field is slightly acidic, but the bottom left/center is more acidic than the rest. (d) The predicted soil temperature map (in Fahrenheit scale).

card serves as the base station controller. The sensors interfaced with the base station through a 802.11b router, with a range of over 100 m.

Gateway: The gateway is a Lenovo Thinkpad in the WA farm and a Dell Inspiron laptop in the upstate NY farm.

Cloud: We use the Azure IoT Suite ([34]) for FarmBeats. The sensor readings, camera images, and drone video summaries are populated through the Azure IoT Hub ([35]), to storage. We use blobs for images, and tables for the sensor readings. Although in our current implementation, the different farms share the Azure account, with table-level access control, we plan to have different cloud service accounts for the different farms, as FarmBeats scales up.

7 RESULTS

We evaluate the components of FarmBeats below:

7.1 Weather Aware Base Station

The FarmBeats base station leverages the algorithm in Section 4 to duty cycle different components. It uses the OpenWeather API [40] to get the weather forecasts and plans the duty cycling scheme for the next day. The weather information gives us the cloudiness percentage for each period of three hours. The cloudiness percentage over three days is plotted in Figure 7(a).

Over this set of three days, we compare three power-awareness schemes. We define the start of the day as 6AM local time. We periodically record the state-of-charge of our solar power backed batteries. First, we let the base station be always on. As shown in Figure 7(b), the battery charge goes up during a sunny day and down during the night. While the base station remains energy neutral during the first day, during subsequent days its battery drains because of cloudy weather, leading to unavailability on the third day. Then, we evaluate the alternate approach. We set the base station to a conservative

duty cycling period. While this ensures that the base station is available on cloudy days, the base station battery charges up to 100% during the sunny days thus wasting solar power that could have been utilized. Moreover, its duty cycling interval collects *15 times* less data than the optimal FarmBeats solution, plotted in 7(d).

FarmBeats collects data on the first two days more frequently owing to high availability of solar power. However, on the third day, it switches to a conservative duty cycling schedule to save power. Of the 15x gain in data collection frequency achieved over a fixed duty cycle, a factor of 2 is because of the TVWS client being duty cycled at a different rate than the Wi-Fi router. An earlier version of our deployment which did not duty cycle the base station faced a downtime of 30% in a cloudy month as opposed to zero downtime for our power-aware design in the same month. Thus, FarmBeats’s power-aware design achieves its goal of maximizing data-freshness while maintaining energy neutrality.

7.2 UAV Flight Planning

As mentioned in Section 5.1, we use an efficient area coverage algorithm in addition to leveraging wind assistance to extend drone battery life. To understand the impact of area coverage algorithms on drone flight time, we compare performance of FarmBeats in covering a given area as compared to the state-of-the-art *East-to-west* algorithm (used by Pix4D, DroneMapper, etc.). As shown in Figure 6(b), the east-to-west algorithm generates sweeping patterns from the east to the west or vice-versa regardless of the area shape. However, FarmBeats generates a path that minimizes the number of waypoints.

Next, we compare the time taken to complete flights planned by the two algorithms to cover a given area. The maximum speed was set to 10m/s and the altitude was set to 20m. Figure 6(c) plots the time taken to complete a flight with the two algorithms in different area geome-

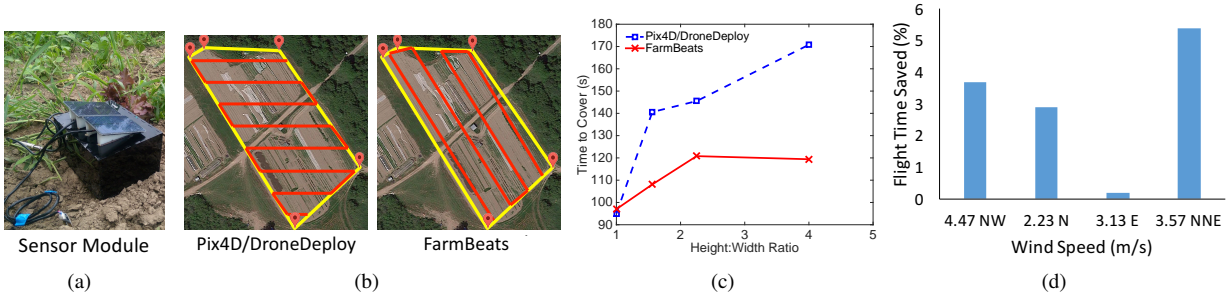


Figure 6: (a) A weather-resistant, solar-powered FarmBeats sensor module. (b,c,d) Drone Flight Planning: (b) FarmBeats’s flight planning algorithm minimizes the number of waypoints to cover a region. (c) Depending on the aspect ratio of the field, flights without FarmBeats’s algorithm take upto 42% more time. This improves the time by a factor of 1.26 in the average case for our farms. (d) In addition, the yaw control algorithm described in Section 5.1 achieves a gain of up to 5% based on the wind velocity.

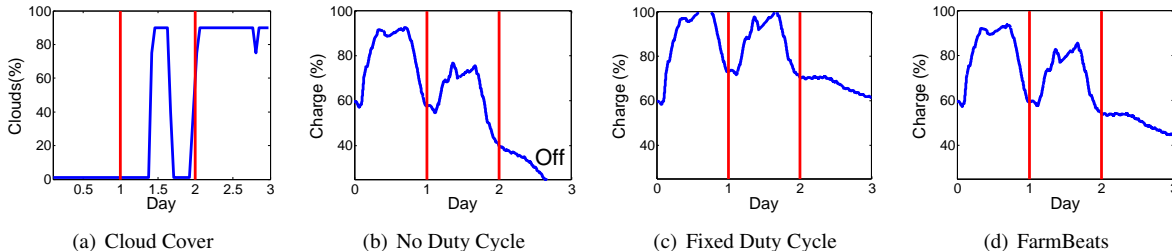


Figure 7: **Power-aware Base Station:** The cloudiness percentage over 3 days. (b) With no duty-cycling, the base station shuts down on a cloudy day. (c) A fixed conservative duty cycle can prevent the base station from going down, but it collects 15 times less sensor data. (d) FarmBeats’s Power-aware basestation can keep the base station on by reducing the duty-cycling on days are expected to be cloudy.

tries defined by their height to width ratio, where height is the distance along the North-South direction and width is measured along East-West. As expected, the gain achieved by FarmBeats increases as the height-width ratio increases. This is because FarmBeats algorithm generates fewer waypoints to cover the same area. In general, for the average case of our deployments, FarmBeats reduced the time taken to cover an area by 26%.

Finally, we evaluate the impact of our yaw control algorithm under different wind conditions. The maximum speed was set to 10m/s and the altitude was set to 30m. For every flight, we fully charged the battery. We measure the percentage of time saved by FarmBeats’s yaw control algorithm for each flight and plot it in Figure 6(d). As seen in the figure, FarmBeats can save up to 5% time depending on the wind velocity. Moreover, as the north-south component (the principal direction of motion for this set of experiments) of the wind increases, FarmBeats can leverage it better.

7.3 Orthomosaic Generation

The novel orthomosaic generation algorithm proposed in this paper advances the state-of-the-art on two fronts. First, our approach of combining sparse 3D reconstruction techniques from video with image stitching techniques is more robust than existing techniques based on either aerial 3D mapping or aerial image stitching. In addition, our approach is computationally more efficient and runs considerably faster than Pix4D [42], an aerial 3D mapping-based tool catering to Precision Agriculture. **Qualitative Results:** We show two representative ortho-

mosaics constructed by FarmBeats and Microsoft ICE in Figure 4 and 8(b) respectively. Figure 8(a) shows what the farm looked like in Google Earth in the past. The orthomosaic generated by Microsoft ICE failed in this case, while our result is consistent and accurate. Our geo-referenced image covers about 5 acres of farmland and provides a detailed visual summary to the farmer. By visually inspecting the high-resolution image, they can discover anomalies such as the water puddle that can render a part of the field unsuitable for agriculture for a couple of seasons. Moreover, the farmer can see where cows are grazing during the day and make a decision about whether they want to move them to another spot for the next day. The decision is based on how much grass they want to leave on the field to be converted into manure.

Processing Time: As shown in Figure 8(c), our implementation is 2.2 times faster than Pix4D on average. Specifically, our method took 14 minutes to construct an orthomosaic on average whereas Pix4D took 32 minutes on average on a set of videos captured by our drones at 1080p resolution at 30 frames per second. This demonstrates the improved running time of our method.

Finally, the orthomosaic generated by our system are approximately 5 times smaller than the original video size at full resolution (in .png format) before applying lossy compression. A single pixel in the geo-referenced orthomosaic is about 2 cm in size which is equivalent to a single penny on the ground. The image resolution and compression quality are parameters that can be tuned to meet any target file size.



Figure 8: **Orthomosaic Generation:** (a) The Google Earth image for the farm in Figure 4. (b) Microsoft ICE image stitching pipeline fails to reconstruct it accurately. (c) Pix4D takes about 2.2x longer on average compared to our approach.

7.4 Generating Precision Maps

As described in Section 5.3, FarmBeats uses the visual features from the orthomosaic overview to extrapolate the sensor values and generate precision maps for soil temperature, soil moisture and pH.

Qualitative Evaluation: We show a representative set of these precision maps in Figure 5. As shown in the figure, based on sensor values in the rest of the farm, the moisture prediction pipeline can estimate that the top left part of the farm has high moisture content even though that part has no sensor there. Similarly, the pH map generates an actionable input in the sense that the bottom left and center of the farm have very low pH and are highly acidic. As a result of this map, the farmer applied lime to enhance the pH and make the soil more neutral.

Note that the pH of the farm varies within the farm at fine granularity. As seen in Figure 5, within a couple of acres, the pH can vary from 4 (very acidic) to 7 (neutral). Soil moisture variance is even higher, with variance seen within a few meters. Precision maps generated by FarmBeats capture this variance accurately, by using the drone videos to extrapolate the sensor data.

Quantitative Evaluation: In order to evaluate the accuracy of the precision maps generated by FarmBeats using the approach described in Section 5.3, we evaluated our system on 5 datasets constructed from the drone videos and sensor data. Each dataset corresponds to a drone flight over the farm (covering 2 acres) and one set of sensor measurements from the sparse sensor deployment. The hyperparameters are learned by doing 5 fold cross validation. As an accuracy metric, we measure the correlation between the predicted sensor values and the ground truth sensor values to see how well the variations in the field are captured by FarmBeats. We compare against two techniques, which do not use the drone video based extrapolation of the sensor values:

- **Nearest Neighbor (SensorsNN):** We assign the value from the nearest sensor to each point in the field.
- **Inverse Distance based Interpolation (SensorsInterp):** We linearly interpolate known sensor values in the field, by using inverse distance as a weight. This technique has been previously been proposed in the context of precision agriculture [14, 55].

For all the analysis, we use leave-one-out evaluation, i.e., we generate a precision map after leaving one of the sen-

sors out of the training set and evaluate the map on the left out sensor. We repeat this process for all the ground sensors and report the averaged results.

The comparison of correlation across the different schemes is shown in Figure 9(a). As shown in the figure, FarmBeats outperforms existing sensor based interpolation techniques. In particular, FarmBeats can accurately estimate the variations of the different sensor values in the field. While sensor based methods do not mirror the variations and hence have nearly zero correlation with the sensor values, FarmBeats’s estimates have high positive correlation with the true sensor values, thus indicating the utility of using the drone video in conjunction with the drone estimates. Finally, the precision maps generated by FarmBeats are 3 orders of magnitude smaller in size on average than the video and can be easily shipped to the cloud during periods of connectivity.

7.5 Other Applications

Figure 9 highlights two other applications that the farmers used FarmBeats for. First, the farmer in NY used FarmBeats sensors to monitor his storage freezers. The temperature in these freezers is carefully regulated below 10° F to prevent produce from going bad. As shown in Figure 9(b), an employee leaving the door open could lead to this temperature going up causing loss to the farmer. This problem is solved by FarmBeats by enabling automated notifications based on these sensor readings in the FarmBeats phone application.

Second, the farmers plugged in cameras at different locations like cow sheds and connected them to the nearest FarmBeats base station. One frame of the camera is shown in Figure 9(c). While the intent of the current application is to manually monitor the cows, one can potentially build an application that can detect anomalies in cow behavior or use cow motion to track animal health [38]. As a preliminary result, we ran a deep neural network based cow detector on the data. The identification boxes are overlaid on the figure.

7.6 End-to-end Deployment Statistics

Data Aggregation: FarmBeats’s deployments at both farms have been running for over six months. Over these deployments, FarmBeats interfaced with around 10 different sensor types, three different camera types, three versions of drones and the farmers’ phones. It collected

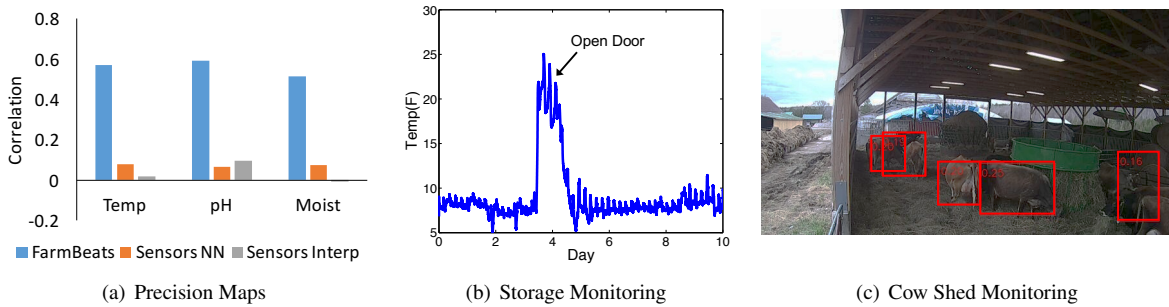


Figure 9: **FarmBeats Applications** (a) FarmBeats’s precision maps are more accurate than standard sensor based interpolation techniques. (b) Temperature (measured in F) in a storage unit can raise an alarm when an employee leaves a door open. (c) Cows being monitored in a cow shed. The red boxes indicate a standard cow detector output.

more than 10 million sensor measurements, half million images and 100 drone surveys.

Resilience to Outages: FarmBeats’s deployments faced one week-long Internet outage due to a thunderstorm and several smaller term Internet outages. The FarmBeats gateway continued to be available during these times.

Cost: The TVWS client radios cost \$200,⁴ and there are no additional data charges, than the farmer’s existing internet connection. The Particle Photons cost about \$20 and can add Wi-Fi support to each sensor. Thus, use of the hybrid networking approach reduces the system cost by an order of magnitude as compared to existing systems which cost over \$1000 in equipment cost per sensor and over 100\$ annual subscription fee (see Table 1).

Applications: Farmers used FarmBeats’s precision agriculture system to guide their precision irrigation units. The precision pH maps generated were used by farmers to apply lime in the more acidic regions. As mentioned before, farmers also used FarmBeats for storage monitoring with sensors and animal shelter monitoring, selling station monitoring with cameras. Beyond that, farmers also used FarmBeats base stations to access Wi-Fi while on the farm to run productivity applications like Trello.

8 RELATED WORK

FarmBeats builds on past work in wireless sensor networks, precision agriculture and ICTD.

Wireless Sensor Networks: Past work has used multi-hop networks [5, 19, 23, 26, 39, 45, 57, 60] to gather data from sensors in the farm. However, all these systems suffer from bandwidth constraints that make them unable to support sensors, cameras and drones. Further, these systems do not account for constraints imposed by weak cloud connectivity and weather related power and Internet outages. The same is true for recent advances in LP-WAN technologies [29, 52]. In contrast, FarmBeats includes support for sensors, cameras and drones; is backed

⁴With the standardization of IEEE 802.11af [1] standard, we expect the price to the client and base station to be similar to Wi-Fi, of less than 10\$. We are testing one such multi-mode TVWS/Wi-Fi chip from a major Wi-Fi vendor.

by cloud connectivity and has mechanisms to adapt to weather variability.

Agriculture: Agronomists have studied various aspects of precision agriculture, from defining more accurate management zones [31], to improving prescription [37], to leveraging soil science [54] and plant physiology [8] techniques. Prior work has also looked at applications of precision agriculture to irrigation, variable seeding, nutrient application, and others. There has been prior work on developing technology for enabling precision agriculture. Researchers have built specialized sensors for measuring nutrients [25], water levels [21], and other such sensors, and we build on top of this work. FarmBeats’s work is complementary to this body of work as it facilitates the automation of data collection using these sensors and enables the precision agriculture systems.

ICTD: ICTD solutions focus on user interfaces to make existing technologies more accessible [13] enhanced access to information [9] and better communications. The mechanisms of data collection is manual in most scenarios. The few attempts at automated data collection, like [9], fall into the same pitfalls as discussed before. We believe FarmBeats is complimentary to this work and will aid the proliferation of ICTD by enabling end-to-end IoT connectivity in weakly connected scenarios.

9 CONCLUSIONS & FUTURE WORK

FarmBeats is a low-cost, highly available IoT platform for agriculture. It supports high bandwidth sensors using TVWS, which is a low-cost, long range technology. FarmBeats uses a weather-aware solar-powered IoT base station, and an intelligent Gateway that ensure that services are available in the Cloud and offline. It also incorporates new path-planning algorithms that extend drone battery life. We have deployed the system in two farms, and the farmers are already using it for three applications: precision agriculture, animal monitoring, and storage monitoring. Moving forward, we are working with the farmers to develop several other applications on top of FarmBeats. Further, we plan to make anonymized data available for researchers to enable more agricultural applications.

REFERENCES

- [1] IEEE 802.11af: <https://standards.ieee.org/findstds/standard/802.11af-2013.html>.
- [2] Adaptrum. <http://www.adaptrum.com/>.
- [3] M. H. Almarshadi and S. M. Ismail. Effects of Precision Irrigation on Productivity and Water Use Efficiency of Alfalfa under Different Irrigation Methods in Arid Climates. *Journal of Applied Sciences Research*, 2011.
- [4] AutoPano. [kolor.com](http://www.kolor.com).
- [5] A. Baggio. Wireless sensor networks in precision agriculture. *ACM Workshop on Real-World Wireless Sensor Networks*, 2005.
- [6] P. Bahl, R. Chandra, T. Moscibroda, R. Murty, and M. Welsh. White Space Networking with Wi-Fi like Connectivity. *ACM SIGCOMM Computer Communication Review*, 2009.
- [7] M. Brown and D. G. Lowe. Automatic Panoramic Image Stitching Using Invariant Features. *International Journal of Computer Vision*, 2007.
- [8] K. G. Cassman. Ecological Intensification of Cereal Production Systems: Yield Potential, Soil Quality, and Precision Agriculture. *Proceedings of the National Academy of Sciences (PNAS)*, 1999.
- [9] J. chun Zhao, J. Zhang, Y. Feng, and J. Guo. The Study and Application of the IOT Technology in Agriculture. *IEEE International Conference on Computer Science and Information Technology*, 2010.
- [10] DataMapper. <http://www.precisionhawk.com/datamapperinflight>.
- [11] Decagon Devices. Decagon Devices Cellular Logger. <https://www.decagon.com/en/data-loggers-main/data-loggers/em50g-wireless-cellular-data-logger/>.
- [12] DJI. <http://developer.dji.com>.
- [13] J. Doerflinger and T. Gross. Sustainable ICT in Agricultural Value Chains. *IT Professional*, 2012.
- [14] T. A. Doerge. In *International Plant Nutritional Institute*, 1999.
- [15] DroneDeploy. dronedeploy.com.
- [16] Federal Communications Commission. <https://www.fcc.gov/general/white-space-database-administration>.
- [17] C. D. Franco and G. Buttazzo. Energy-aware coverage path planning of uavs. In *International Conference on Autonomous Robot Systems and Competitions (ICARSC)*, , 2015.
- [18] FreeWave. <http://www.freewave.com/>.
- [19] A.-J. Garcia-Sanchez, F. Garcia-Sanchez, and J. Garcia-Haro. Wireless Sensor Network Deployment for Integrating Video-surveillance and Data-monitoring in Precision Agriculture over Distributed Crops. *Computers and Electronics in Agriculture*, 2011.
- [20] H. C. J. Godfray, J. R. Beddington, I. R. Crute, L. Haddad, D. Lawrence, J. F. Muir, J. Pretty, S. Robinson, S. M. Thomas, and C. Toulmin. Food Security: The Challenge of Feeding 9 Billion People. *Science*, 2010.
- [21] B. R. Hanson and S. Orloff. Monitoring soil moisture for irrigation water management. Technical report, UC Davis, 2007.
- [22] J. Hsu, S. Zahedi, A. Kansal, M. Srivastava, and V. Raghunathan. Adaptive Duty Cycling for Energy Harvesting Systems. In *International Symposium on Low Power Electronics and Design*, 2006.
- [23] O. V. K. Langendoen, A. Baggio. Murphy loves Potatoes: Experience from a Pilot Sensor Network Deployment in Precision Agriculture. In *International Parallel and Distributed Processing Symposium (IPDPS)*. IEEE, 2006.
- [24] Kathleen McLaughlin . Gaps in 4G Network Hinder High-tech Agriculture: FCC Prepares to Release 500 Million to Improve Coverage.
- [25] H.-J. Kim, K. A. Sudduth, and J. W. Hummel. Soil Micronutrient Sensing for Precision Agriculture. *Journal of Environmental Monitoring, The Royal Society of Chemistry*, 2009.
- [26] W. Lee, V. Alchanatis, C. Yang, M. Hirafuji, D. Moshou, and C. Li. Sensing Technologies for Precision Specialty Crop Production. *Computers and Electronics in Agriculture*, 2010.
- [27] Z. Li and V. Isler. Large Scale Image Mosaic Construction for Agricultural Applications. *IEEE Robotics and Automation Letters*, 2016.
- [28] D. Logger. PoE Switch. <http://www.digital-loggers.com/poe48.html>.
- [29] LoRa Technology. <https://www.lora-alliance.org/what-is-lora/technology>.
- [30] J. Lowenberg-DeBoer. The Precision Agriculture Revolution: Making the Modern Farmer. <https://www.foreignaffairs.com/articles/usa/2015-04-20/precision-agriculture-revolution>.
- [31] A. McBratney and M. Pringle. Estimating average and proportional variograms of soil properties and their potential use in precision agriculture. *Precision Agriculture*, 1999.
- [32] A. McBratney, B. Whalen, T. Ancev, and J. Bouma. Future Directions of Precision Agriculture. *Precision Agriculture*, 2005.
- [33] Microseven. Ip cameras. <http://www.microseven.com/product/IP-Cameras.html>.
- [34] Microsoft. Azure IoT Hub. <https://azure.microsoft.com/en-us/services/iot-hub/>.
- [35] Microsoft. Azure IoT Suite. <https://www.microsoft.com/en-us/cloud-platform/internet-of-things-azure-iot-suite>.
- [36] Microsoft Research. Image Composite Editor. <http://research.microsoft.com/en-us/um/redmond/projects/ice/>.
- [37] N. D. Mueller, J. S. Gerber, M. Johnston, D. K. Ray, and J. A. F. Navin Ramankutty. Closing Yield Gaps through Nutrient and Water management. *Nature*, 2012.
- [38] L. Nagl, R. Schmitz, S. Warren, T. S. Hildreth, H. Erickson, and D. Andresen. Wearable Sensor System for Wireless State-of-health Determination in Cattle. In *Engineering in Medicine and Biology Society*, 2003.
- [39] T. Ojha, S. Misra, and N. S. Raghuvanshi. Wireless sensor networks for agriculture. *Computers and Electronics in Agriculture*, 2015.
- [40] OpenWeatherMap, Inc. OpenWeather API. <http://openweathermap.org/>.
- [41] B. Ortiz, J. Shaw, J. fulton, and A. Winstead. Management Zones II – Basic Steps for Delineation. *Alabama Precision Ag Extension: Precision Agriculture Series*, 2011.
- [42] PIX4D. Pix4d. <https://pix4d.com/>.
- [43] C. E. Rasmussen and C. K. I. Williams. *Gaussian Processes for Machine Learning (Adaptive Computation and Machine Learning)*. The MIT Press, 2005.
- [44] S. Roberts, P. Garnett, and R. Chandra. Connecting Africa Using the TV White Spaces: From Research to Real World Deployments. In *IEEE LANMAN*, 2015.
- [45] L. Ruiz-Garcia, L. Lunadei, P. Barreiro, and J. I. Robla. A Review of Wireless Sensor Technologies and Applications in Agriculture and Food Industry: State of the Art and Current Trends. *Sensors*, 2009.

- [46] A. R. Schepers, J. F. Shanahan, M. A. Liebig, J. S. Schepers, S. H. Johnson, and A. Luchiari. Appropriateness of Management Zones for Characterizing Spatial Variability of Soil Properties and Irrigated Corn Yields across Years. *Agronomy Journal*, 2004.
- [47] D. Schimmelpfennig. Cost Savings From Precision Agriculture Technologies on U.S. Corn Farms. *USDA Amber Waves*, 2016.
- [48] SenseFly. Sensefly. sensefly.com.
- [49] Sentera. <https://sentera.com/>.
- [50] M. Shamos. *Computational Geometry*. Yale University, 1978.
- [51] N. Sharma, J. Gummeson, D. E. Irwin, and P. J. Shenoy. Cloudy Computing: Leveraging Weather Forecasts in Energy Harvesting Sensor Systems. In *IEEE SECON*, 2010.
- [52] SIGFOX. <http://www.sigfox.com/>.
- [53] S. N. Sinha and M. Pollefeys. Pan-tilt-zoom Camera Calibration and High-resolution Mosaic Generation. *Computer Vision and Image Understanding*, 2006.
- [54] M. Soderstrom, G. Sohlenius, L. Rodhe, and K. Piikki. Adaptation of Regional digital soil mapping for precision agriculture. *Precision Agriculture*, 2016.
- [55] J. V. Stafford. *Precision Agriculture Book*. 2005.
- [56] R. Szeliski. Image Alignment and Stitching: A Tutorial. Technical report, Microsoft Research, 2004.
- [57] W. T., P. Csiro, P. Corke, Sikka, L. Klingbeil, Y. Guo, C. Crossman, P. Valencia, D. Swain, and G. Bishop-Hurley. Transforming Agriculture through Pervasive Wireless Sensor Networks. *IEEE Pervasive Computing*, 2007.
- [58] United Nations General Assembly. Food Production Must Double by 2050 to Meet Demand from Worlds Growing Population, Innovative Strategies Needed to Combat Hunger, Experts Tell Second Committee, 2009. <http://www.un.org/press/en/2009/gaef3242.doc.htm>.
- [59] C. M. Vigorito, D. Ganesan, and A. G. Barto. Adaptive Control of Duty Cycling in Energy-harvesting Wireless Sensor Networks. In *IEEE SECON*, 2007.
- [60] G. Werner-Allen, K. Lorincz, J. Johnson, J. Lees, and M. Welsh. Fidelity and Yield in a Volcano Monitoring Sensor Network. In *Operating Systems Design and Implementation*, 2006.

10 SUPPLEMENTARY MATERIAL

10.1 FarmBeats Applications and Usage

The goal of FarmBeats is to serve as a substrate for multiple sensing modalities on the farm. Sensors with varying data requirements can plug-in to FarmBeats and operate seamlessly. This allows farmers to use FarmBeats for various applications. Our primary target has been to deliver a class of applications that fall under the category of precision agriculture. Precision agriculture is a technique to improve yield by treating the farm as heterogeneous land, and uses variable treatment throughout the farm, such as variable seeding, fertilizer application, lime application, irrigation, and many other agricultural services. In principle, precision agriculture is good for the overall farming ecosystem [37]. It improves yield, reduces the operating expenses for the farmer [47], and is also good for the environment.

In contrast to existing systems which divide the farm into large static management zones [14, 41] and fail to




Figure 10: **Leveraging Wind Assistance:** When the effective velocity, \vec{v}_{eff} , aids the UAV motion, then the yaw is set to be perpendicular to the UAV motion (left), otherwise the yaw is aligned with the effective wind velocity to minimize air drag (right).

capture temporal and climatic variations [32, 46], FarmBeats delivers near real-time and fine-grained precision maps to farmers for soil characteristics like pH, moisture, etc. As described in section 5.3, FarmBeats uses a combination of aerial imagery and sparse sensor deployment to deliver these maps. They are currently being used by farmers to monitor and amend irrigation, and lime application practices in their farms.

In addition to precision agriculture applications, farmers have used FarmBeats for monitoring cattle using cameras in barns and for using sensors to monitor temperature in storage units. This functionality was suggested by farmers and later added to the system. Future applications suggested by farmers include monitoring the net carbon footprint of the agricultural production cycle, net nutrient usage for each crop cycle, crop suggestions using long-term data and flood monitoring.

10.2 Leveraging Wind to Assist UAV Path Planning

Algorithm 1 Pseudo-code for determining UAV yaw based on wind speed

▷ Input: Wind velocity (\vec{v}_W), UAV velocity (\vec{v}_{UAV}), intended acceleration (\vec{a}_{UAV})

▷ where all inputs are measured with respect to the earth reference frame

▷ Output: UAV yaw (y)

Compute wind velocity w.r.t the UAV: $\vec{v}_{eff} = \vec{v}_W - \vec{v}_{UAV}$

if $\vec{v}_{eff} \cdot \vec{a}_{UAV} > 0$ **then**

Set the yaw perpendicular to \vec{a}_{UAV}

$y = \angle \vec{a}_{UAV} + \frac{\pi}{2}$

else

Align the yaw with \vec{v}_{eff} to minimize drag

$y = \angle \vec{v}_{eff}$

end if

Here, we describe the algorithm to leverage wind to as-

sist in UAV path planning. On a high level, the algorithm works in two steps:

- Compute the effective wind velocity, \vec{v}_{eff} . The effective wind velocity is the wind velocity in the reference frame of the UAV. Specifically, if the velocity of the UAV with respect to the ground is \vec{v}_{UAV} and the wind velocity with respect to the ground is \vec{v}_W , then $\vec{v}_{eff} = \vec{v}_W - \vec{v}_{UAV}$. An example of this vector computation is shown in figure 10.
- If \vec{v}_{eff} has a component that can aid the UAV motion, then we make the yaw perpendicular to the direction of intended acceleration, otherwise, we align it with the direction of \vec{v}_{eff} to minimize air drag. For instance, as shown in figure 10, if the UAV wants to accelerate and the UAV velocity is large such that the effective wind velocity has no component that aids acceleration, then the algorithm aligns the UAV yaw with the effective wind velocity, minimizing the air drag.

Algorithm 1 describes the pseudo-code for FarmBeats's approach.