Smallholder Agriculture in the Information Age: Limits and Opportunities

Mariya Zheleva
Computer Science
University at Albany SUNY
mzheleva@albany.edu

Petko Bogdanov
Computer Science
University at Albany SUNY
pbogdanov@albany.edu

Daphney-Stravoula Zois
Electrical and Computer Engineering
University at Albany SUNY
dzois@albany.edu

Wei Xiong
Computer Science
University at Albany SUNY
wxiong@albany.edu

Ranveer Chandra
Microsoft Research
Redmond, WA
ranveer@microsoft.com

Mark Kimball
Essex Farm
Essex, NY
essexfarm@gmail.com

ABSTRACT

Recent projections by the United Nations show that the food production needs to double by 2050 in order to meet the nutrition demand of the world's growing population. A key enabler of this growth are smallholder family farms, that form the backbone of agricultural (AG) production worldwide. To meet this increasing demand, smallholder farms need to implement critical advances in task management and coordination, crop and livestock monitoring and efficient farming practices. Information and Communication Technology (ICT) will play a critical role in these advances by providing integrated and affordable cyber-physical systems (CPS) that can longitudinally measure, analyze and control AG operations. In this paper we make headway towards the design and integration of such AG-CPS. We begin by characterizing the information and communication technology demand of smallholder agriculture based on traffic analysis of farm Internet use. Our findings inform the design and integration of an end-to-end AG-CPS called FarmNET that provides (i) robust control mechanisms for *multi-sensor AG data collection and fusion*, (ii) wide-area, heterogeneous wireless networks for *ubiquitous farm connectivity*, (iii) algorithms and models for farm data analytics that produce *actionable information* from the collected agricultural data, and (iv) control mechanisms for *autonomous, proactive farming*.

KEYWORDS

Smallholder agriculture, integrated AG-CPS, wireless networks, data analytics, control and automation.

ACM Reference format:
Mariya Zheleva, Petko Bogdanov, Daphney-Stravoula Zois, Wei Xiong, Ranveer Chandra, and Mark Kimball. 2017. Smallholder Agriculture in the Information Age: Limits and Opportunities. In *Proceedings of LIMITS '17, Santa Barbara, CA, USA, June 22-24, 2017,* 12 pages.
DOI: http://dx.doi.org/10.1145/3080556.3080563

LIMITS '17, Santa Barbara, CA, USA

DOI: http://dx.doi.org/10.1145/3080556.3080563

1 INTRODUCTION

Smallholder farms rely predominantly on single-family labor. Such farms form the backbone of agricultural (AG) production and are essential to eradicate hunger in the face of a changing climate, while preserving our natural resources [60]. Recent estimates indicate that 80% of the food produced in the developing world comes from smallholder farms [78]. This number is far exceeded in the U.S., where USDA estimates that 97.6% of farms are smallholder enterprises and they are responsible for 85% of the nation's AG production [88]. With the world's booming population, the United Nations foundation estimates that the farm production needs to double by the year 2050 for society to be able to eradicate hunger and secure nutrition [52]. This creates an appealing market opportunity for smallholder farms to proliferate while solving one of humanity's big challenges. Such growth mandates improved efficiency in current farm practices related to (i) task management and coordination, (ii) crop and livestock monitoring, data analytics and control, (iii) expansion of local markets and (iv) access and adoption of new farming practices [66].

Information and communication technology (ICT) will play a critical role in such advances, however, current technologies [9, 14, 17–19] are highly-specialized (e.g. focusing on soybean production), provide closed solutions in that farmers have no control over their data and corresponding analytics, and most importantly, are not affordable for smallholder enterprises. Thus, practical progress in ICT for smallholder agriculture hinges on the availability of technology that either does not exist or needs to be re-purposed from its predominantly urban context to fit the unique spatial, temporal and environmental characteristics of smallholder farming. Such technology includes (i) *robust sensing infrastructures* to measure farm state and operations, (ii) *ubiquitous wireless network connectivity* to transmit sensor and farmer data, (iii) *domain-specific data models and analytics* to extract actionable knowledge from the data, and (iv) *adaptive control algorithms* for efficient sensing and proactive control of farm processes towards autonomous farming. Beyond availability, it is critical for these technologies to be ***seamlessly integrated into an AG cyber-physical system (AG-CPS)*** that interacts with farmers, farm assets and processes to efficiently measure, analyze and control them, and inform decision-making, improved farming practices, distribution chains and consumer relations. To this end, ***there is a need of fundamental research and***

engineering of novel and integrated end-to-end mechanisms for farm sensing, network connectivity, data analytics and control.

Related research and technology integration can be subdivided into (i) wireless networks for under-served areas [37, 69, 72, 74, 75, 91, 107, 110, 116, 117, 121, 134, 135], (ii) data mining for AG applications [34, 48, 51, 59, 61, 64, 70, 80, 89, 90, 92, 101, 102, 104, 119, 122, 132, 137, 144] and (iii) estimation and control for decision making in uncertain environments under constraints [146–156]. Key limitations of existing solutions are that they are either developed for non-AG contexts or tackle connectivity, data analytics and control in isolation. To bring meaningful ICT innovation in agriculture, we need to adopt an integrated approach.

Our work informs such an integrated approach via systematic analysis of farm Internet use and ICT needs. Our analysis is based on a year-long and continuing collaboration with Essex Farm in upstate New York that has allowed us to learn first-hand about the ICT needs of farm operations. With farmers' permission, we have also been able to collect traces of farmer mobility and Internet use. Our analysis of these traces shows that farm traffic is a unique mix of farmer and IoT sensor activity with interlocking characteristics. In terms of volume and direction, the IoT sensor traffic is upload-intensive, while farmer traffic is download-intensive. While farmer traffic is bursty and unpredictable, that of IoT sensors is primarily periodic and, thus predictable. Finally, while farmer traffic is spatially-concentrated, IoT sensor traffic is distributed across the farm's territory.

These insights create a unique design space for (i) ubiquitous AG wireless network architecture and protocols, (ii) novel AG data analytics and (iii) AG control. We integrate these key components and present our vision of an end-to-end AG-CPS called FarmNET. FarmNET integrates four key components to collect longitudinal data, and utilize real-time data analytics, domain-specific models and control algorithms to enable increased quality and productivity of farm operations, with minimal footprint.

(1) **Sensing frontend.** Each FarmNET sensor will be wirelessly connected and highly-reconfigurable. The wireless capability of sensors will allow for seamless farm data offload, whereas the reconfigurability will enable adaptive sampling to efficiently manage the tradeoffs between volume and periodicity of farm measurements versus accuracy of data analytics and control algorithms.
(2) **Communication network.** To accommodate the unique farm traffic, we envision a heterogeneous wireless network that is comprised of a plug-and-play wireless backhaul realized over TV white spaces, and a three-modal last mile implemented over LTE or Wi-Fi. This architecture design poses fundamental challenges in (i) joint cross-layer optimization of last-mile and backhaul access that is informed by the properties of the heterogeneous farm traffic, and (ii) characterization, modeling and integration of power efficiency of AG-CPS.
(3) **Data analytics.** A key challenge in FarmNET is to enable robust knowledge discovery from noisy and sparse sensor measurements and to employ them for analytic tasks such as anomaly detection, root cause analysis, historical trend detection and prediction based on statistical data-driven models and simulation. This problem lends itself for a graph-theoretical formulation with nodes representing farm entities (pastures, animals, arable fields, etc.), connections modeling interactions among entities, and multi-variate graph signals modeling temporal entity states (e.g. soil moisture, animal health and milk production). This dynamic heterogeneous graph framework will enable a holistic understanding of all sensed farm operations and enable modeling and control-enabled optimization of the global farm health.
(4) **Control.** To enable high–output and efficient, controlled–environment AG technologies and systems, we envision a holistic controlled sensing framework that will integrate IoT–sensing capabilities with agriculture data collection, network structure and humans in the loop, to enable real–time accurate agricultural monitoring and control. To this end, we propose (i) a stochastic dynamic system model that fully describes the farm's state (i.e., health and footprint) over time, while incorporating the effect of both cyber (e.g., control signals) and physical components (e.g., agricultural variables), (ii) recursive, structured farm state estimators, and (iii) strategies that control sensing and farm processes by optimizing farm state estimation accuracy and different operation costs (e.g., sensing, farm processes).

In what follows, we first describe the current state of AG ICTs (§2). Next, we provide analysis of farm ICT needs that is based on empirical evaluation of farm Internet access (§3). Our analysis informs the design of an end-to-end AG-CPS dubbed FarmNET that is presented in §4. Finally, we conclude our paper in §5.

2 LIMITS AND OPPORTUNITIES

In this section, we first describe the current state of ICTs for agriculture and detail their limitations. We survey existing solutions that can be harnessed to address these limitations. We also discuss several technological needs that cannot be met by re-purposing of existing technologies and require fundamentally new design.

2.1 Current State of AG ICTs

Multiple solutions that target precision agriculture and AG decision support exist both in ***industry*** and ***academia***.

2.1.1 Industrial products. Industrial products can be largely subdivided in such that target *(i) sensing* [1, 2, 5, 7, 14], *(ii) data analytics* [4, 8, 10–13, 15–17, 17, 21] and *(iii) consumer relations* [3]. A large fraction of the sensing solutions perform single-modality sensing, i.e., only imagery [2, 5], moisture [7] or nitrates [1]. gThrive [14] is the only one that supports multi-modal sensing, and basic data fusion and analytics. The industry has largely focused on AG data analytics with a large number of start-ups [4, 8, 10–13, 15, 17, 21] and well-established AG corporations [16, 17] entering this business. All of these are closed-form, cloud-based solutions that do not allow flexibility in data management and do not provide farmers access to their own data. This, as found in recent research [82], has raised concerns around data privacy, security and control. A common limitation found by smallholder farmers with regards to existing industrial products is that they are prohibitively-expensive, and thus, not economically-feasible for smallholder farm operations. In addition, such systems are typically not open and thus do not provide opportunities for modular customizations for different smallholder farm operations.

Summary of limitations. Industrial products have several key limitations, related to cost, and flexibility of access and data management. Furthermore, none of the existing industrial products provide a wide-profile, end-to-end solution; instead, they focus either on sensing or data analytics and typically take network connectivity for granted. While decision support is the focus of some existing products, current efforts in control mechanisms for autonomous farming are limited, highly-specialized and out of financial reach of smallholder farms.

2.1.2 Academic research. Related research and technology integration can be subdivided into (i) wireless networks for under-served areas [37, 69, 72, 74, 75, 91, 107, 110, 116, 117, 121, 134, 135], (ii) data mining for AG applications [34, 48, 51, 59, 61, 64, 70, 80, 89, 90, 92, 101, 102, 104, 119, 122, 132, 137, 144] and (iii) estimation and control for decision making in uncertain environments under constraints [146–156]. Key limitations of existing solutions are that they are either developed out of the AG context or tackle connectivity, data analytics and control in isolation.

Recent technological advances in precision agriculture hinge on the availability of ***wireless network connectivity***, however, all of them take connectivity for granted (e.g. products described in §2.1.1). At the same time, smallholder farmlands, with their extremely-low population density, often provide the least-appealing business case for commercial network deployments. As a result, farmlands are characterized with spotty, inconsistent, intermittent or all together lacking network coverage. Advances in wireless networks for under-served areas [37, 69, 72, 74, 75, 91, 107, 110, 116, 117, 121, 134, 135] bring promise for improved farm connectivity. Unfortunately, existing solutions are designed exclusively for human-generated traffic and are not readily applicable for farm connectivity that needs to accommodate a mix of human and IoT sensors traffic with varying delay constraints (detailed description and preliminary results in § 3). A large volume of prior work focuses on wireless sensor network connectivity (WSN) [22–25, 29, 31, 56, 62, 142] with some specializing in WSN for agriculture [33, 58, 111, 139–141]. A key limitation of these works is that they only accommodate sensor data and will not scale well for an integrated AG-CPS such as FarmNET, that is optimized to handle heterogeneous farmer-sensor traffic.

In ***AG data analytics***, data mining and machine learning techniques have been applied to extract high-level knowledge of the farm state [119]. Of central interest are crop [122, 137] and animal health [48, 51, 59], soil properties [35, 100], animal tracking and behavior inference [59, 70, 80, 90, 102, 104, 122, 132, 137, 144]. These techniques are designed for offline processing of previously-collected data, however, they are not suitable for real-time tracking of multiple interacting entities (e.g. animals, pastures, feed and weather) as an evolving network. Tracking of such multi-entity, longitudinal interactions requires a dynamic-network-mining approach and is critical to enable anomaly detection, root-cause analysis and realistic simulation for "what-if" analysis. Dynamic network mining is an emerging research field that has produced scalable methods for anomalous temporal subnetwork detection [44, 98, 99, 130], prediction of the network's global state based on local properties [54, 55], information and disease propagation [42, 43] applied to transportation, biological and social networks. Such methods, however, are not readily-adoptable in the AG context, as AG processes incur different interaction dynamics. For example, the temporal interaction between grazing herd animals and pasture paddocks requires novel definitions of anomalies and novel predictive models for pasture productivity in the presence of grazing animals, varying weather and nutrients within the spatial network of paddocks.

Prior work on ***monitoring and control*** for precision agriculture and farm monitoring such as [28, 30, 36, 39, 40, 45, 46, 53, 58, 67, 81, 103, 143] has considered either ad–hoc or static optimization approaches. However, real–time and cost–efficient monitoring and control requires rigorous dynamic farm system modeling, which jointly considers the cyber and physical components and precisely defines their interactions, optimization and control system theories. Similar approaches have been successfully applied to other applications (e.g., environmental monitoring [20, 49, 94–97, 115, 127–129], target tracking [32, 47, 50, 65, 71, 73, 83, 86, 106, 124, 125], physical activity tracking [145, 147, 149, 152]), however, the proposed solutions are not readily applicable in the AG context, since AG-CPS (i) require more complex dynamic models, (ii) require functions that control both sensing and farm processes, and (iii) exploit the unique AG system characteristics (e.g., humans–in–the–loop, network–induced constraints, model structure).

Summary of limitations. While the academic community has made a substantial headway towards data-driven agriculture, no work focuses on providing an end-to-end solution to enable sustainable, proactive and autonomous agriculture. Furthermore, a key limitation of existing data-driven approaches is that they focus on single-time, offline analytics and are thus not suited for longitudinal, real-time and actionable analysis of AG data. Similarly, AG control has considered static optimization or ad hoc approaches, however, further development is necessary to enable real-time, dynamic control. Finally, in terms of wireless networks, existing solutions focus on accommodating human-generated traffic and will not scale well for the heterogeneous demand on farm networks.

2.2 Limits and Opportunities Specific to Smallholder Farming

Varying seasonal workforce; limited connectivity in remote rural locations; lack of accessible systems for quality control, planning and operational analytics; and maintaining close working relationships with (possibly multiple) end-customers are among the main challenges for smallholder family-operated farms. Thus, automation of the common monitoring tasks via low-cost sensing and connectivity solutions, data-driven planning and control as well as offering measurable sustainability/quality statistics for end customers open tangible opportunities for improving smallholder farming enterprises of varying scales. These challenges transcend to developing world smallholder farms, although the tradeoffs between cost and system utility in these scenarios need to be further considered.

In order to embrace the above opportunities, emerging AG-CPS need to employ longitudinal data, and utilize real-time data analytics, domain-specific models and control algorithms to enable increased quality and productivity with minimal footprint. *These challenges create an appealing research agenda for (i) robust multi-sensor data collection and fusion, (ii) wide-area, heterogeneous wireless networks for ubiquitous farm connectivity, (iii) algorithms and* *models for farm data analytics that produce actionable information from raw sensor data, and (iv) novel estimation and control mechanisms for autonomous and proactive farming.* These components will need to be *seamlessly-integrated in a holistic and modular AG-CPS.* In the design of such integrated AG-CPS, of central importance should be the trade-off between the economic feasibility of the sensing and communication infrastructure and the accuracy and efficiency of data analytics, monitoring and control to promote proactive farming practices. This tradeoff can be tackled by developing and adopting algorithms, models, hardware and software that leverage open-source and highly-reconfigurable components.

3 ANALYSIS OF FARM ICT NEEDS

Our analysis of farm ICT needs is based on a measurement campaign we executed in Essex Farm between July and December of 2016. In what follows we provide background on Essex Farm, our methodology and objectives, results, and design implications.

3.1 Essex Farm

Essex Farm[1] was established in 2004 by Mark and Kristin Kimball and is a unique diverse-profile family-operated farm in Upstate New York that spans an area of 1,100 acres. It operates as a farm-to-door CSA (Community Supported Agriculture), but unlike classical CSAs, it provides a full, all-you-can-eat diet, year-round to its members. The farm specializes in a diverse profile of agricultural activities from vegetable and fruit production to grains, eggs, dairy and wide spectrum of meats. Beyond production, the farm also collaborates with a local enterprise called The Hub on the Hill[2] to make preserves from the seasonal produce to maintain its supplies year-round. The farm employs anywhere between 5 and 20 additional farmers throughout the year. These farmers are typically young professionals, who come from different parts of the U.S. and Europe and are looking to get training and hands-on experience with farming. Thus Essex Farm provides them with a unique opportunity to (i) learn in a farm with a diverse activity profile and (ii) interact with cutting-edge IT innovation that is undergoing on the farm. Overall, the diverse activity profile of the farm, its farm-to-door operation, its employment of young farming professionals and its collaboration with other enterprises in AG sustainability makes for a unique ecosystem to understand AG ICT needs and opportunities.

3.2 Methodology

The farm Internet access is currently provided over a 5MBps microwave wireless terrestrial link that beams over lake Champlain to connect the farm with their ISP in Vermont. The gateway link is then locally-distributed through three Wi-Fi access points connected to the gateway via an Ethernet LAN. Besides the farm Wi-Fi, there is also limited coverage provided by commercial mobile carriers. ***The goals of our measurements*** were to (i) characterize the commercial network availability and quality on the farm, and (ii) understand the volume, direction and spatio-temporal characteristics of farm traffic demand. For the first task, we developed an Android application to collect geo-tagged network performance information every 30 seconds and submit it to our server for storage and analysis. Two phones running the app were carried by different farmers over the course of a week in high season (July). For the second task, we collected longitudinal pcap traces at the farm gateway. Once collected, we post-processed these pcap traces using tstat[3] in order to extract individual TCP and UDP flows and study the per-flow performance and inter-arrival rate.

3.3 Analysis Results

The farm Internet demand is generated by a mix of sensors and farmers. Farmers use applications that require real-time Internet access, whereas sensors are a mix of real-time access and delay-tolerant nodes. In the remainder of our analysis, we split the collected traces into farmer-generated and sensor-generated, and apply the same analysis methodology to the two trace subsets. This leads to unique insights into the characteristics of farm ICT demand.

We begin our analysis by focusing on the traces collected by our Android application. Using these traces, we study the spatial characteristics of traffic demand, and the network availability and quality. We find that farmer traffic is highly-localized, whereas sensor traffic is spatially-distributed. We also find that the current Wi-Fi network and commercial cellular network are not able to meet the offered demand. We then focus on pcap trace analysis. We find that farmer and sensor traffic have opposing characteristics in terms of traffic volume, direction and predictability. In what follows, we detail our results.

3.3.1 Spatial distribution of supply and demand. *– Network availability.* Figure 1 presents our results for RSSI measurements of Wi-Fi and one of the major U.S. mobile carriers[4]. Reliable Internet access is available only in the area around the farm office, house, shop and barn and is extremely poor (cellular) or lacking (Wi-Fi) in the rest of the farm.
– Spatial traffic characteristics. We split the farm territory (1,100 acres) into 10x10 meter squares and analyze the frequency of farmer visits of each square. In the course of a week ***only 2.5% of all the 10x10 grid squares were visited by farmers***. This indicates high spatial concentration of farmer traffic demand, which means that a majority of farmer Internet access can be accommodated with several stationary always-on access points. Unlike farmers' traffic, the IoT sensor traffic is spatially-distributed across the farm due to the need for ubiquitous farm measurements.

3.3.2 Traffic volume, direction and predictability. *– Temporal traffic characteristics.* We collected longitudinal pcap traces that capture all the farm traffic (IoT sensors and farmers). We compare the flow inter-arrival time (IAT) of farmer and sensor traffic. According to Figure 2, the farmer traffic arrives at a wide range of intervals (from 100ms to 100s) and is unpredictable. At the same time, 70% of the IoT sensor traffic is characterized with an IAT of 50s, thus, IoT farm traffic is predictable.
– Volume and direction of traffic. Lastly, we are interested in characterizing the intensity of traffic (volume) in the uplink and downlink

[1] http://www.essexfarmcsa.com/
[2] http://thehubonthehill.org/
[3] http://tstat.tlc.polito.it/
[4] Interactive maps available at https://goo.gl/yvEwZu

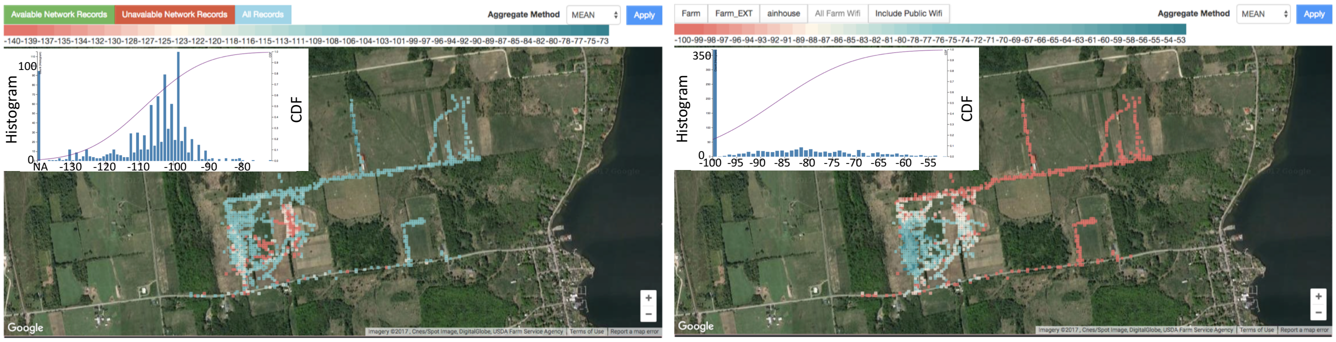

Figure 1: Cellular (left) and Wi-Fi (right) coverage on Essex Farm. No cellular access was detected in roughly 20% of the measurements. Even when available, the cellular network signal strength rarely exceeded -100dBm, which does not suffice for meaningful Internet access. Wi-Fi access is available only in the office area and is lacking anywhere else on the farm.

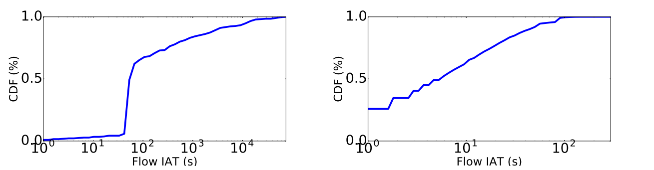

Figure 2: Flow inter-arrival time (IAT) of sensor (left) and farmer traffic (right). IoT sensor traffic is predictable with 70% of flows having IAT of 50s. Farmer traffic's IAT varies between 100ms and 100s and is not predictable.

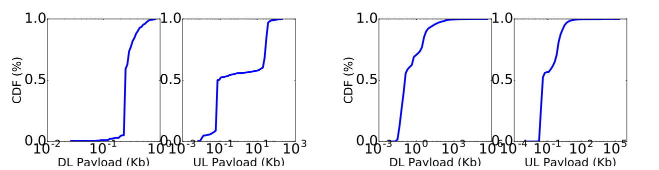

Figure 3: Volume and direction of sensor (left) and farmer (right) traffic. Sensor traffic is predominantly in uplink direction, averaging 100KB in size per flow. Farmer traffic is downlink-intensive with up to 500KB-sized flows.

	Farmer traffic	**IoT-Sensor traffic**
Spatial	Highly-localized	Distributed
Timeliness	Real-time	Real-time+delay-tolerant
Periodicity	Bursty, unpredictable	Mostly predictable
Volume	Bandwidth-intensive	Not bandwidth-intensive
Direction	Downlink	Uplink

Table 1: Summary of findings and design outlook.

direction. For this analysis, we again use pcap traces. Figure 3 presents our results for IoT sensors (left) and farmer traffic (right) demonsrtating that the sensor traffic is uplink-intensive, while the farmer traffic is downlink-intensive.

3.4 Design Implications

Our analysis, summarized in Table 1, shows that IoT sensor and farmers traffic have opposing characteristics across all evaluation criteria. This creates a unique design space for novel (i) AG wireless network architecture and protocols, (ii) AG data analytics and (iii) AG monitoring and control.

– Implications on wireless networking. The joint spatial and timeliness characteristics of farm traffic have direct *implications on network architecture design*, as they permit successful accommodation of majority of the farm traffic, without having to deploy a dense, always-on wireless network throughout the entire farm. Instead, we envision a plug-and-play wireless backhaul that maintains a number of always-on stationary hot-spots and allows opportunistic access for the swath of delay-tolerant sensor traffic. Our findings on traffic types, periodicity, volume and direction call for *a cross-layer protocol design*, that targets rapid transfer of heterogeneous farm traffic and is informed by the volume, direction and predictability of this traffic.

– Implications on data management and mining. The necessity to support farmer and sensor traffic will raise an important trade-off question for data analytics and mining using IoT sensor readings: What is the minimum temporal resolution for different sensing modalities that ensure high quality (e.g., correctly identified anomalies, accurate animal tracking, etc.), while minimizing the rate of sensor readings and thus, not overloading the backhaul with unnecessary sensor data? In addition, real time analytics will have to incorporate delay tolerance and possibly missing values, while still providing maximally useful results to the end users.

– Implications on monitoring and control. The statistical characteristics of IoT–sensor–generated traffic and the structure of the proposed network architecture suggests that measurement and control signal information will be communicated probabilistically. As a result, the dynamic farm state model will need to incorporate the specific network–induced constraints, while the monitoring and control processes models will need to account for the delayed arrival (or missing) of measurement and control signals. Real–time, cost–efficient controlled–environment agriculture requires the design of appropriate estimation and control strategies for under-served areas to ensure the unobstructed operation of the AG-CPS system.

4 FARMNET

FarmNET, as illustrated in Figure 4, is an integrated architecture for real-time agricultural data collection, analytics and farm control. The farm ecosystem consists of farmers, IoT sensors and farm operations pending optimization. FarmNET integrates sensing and communication with data analytics and control to facilitate longitudinal

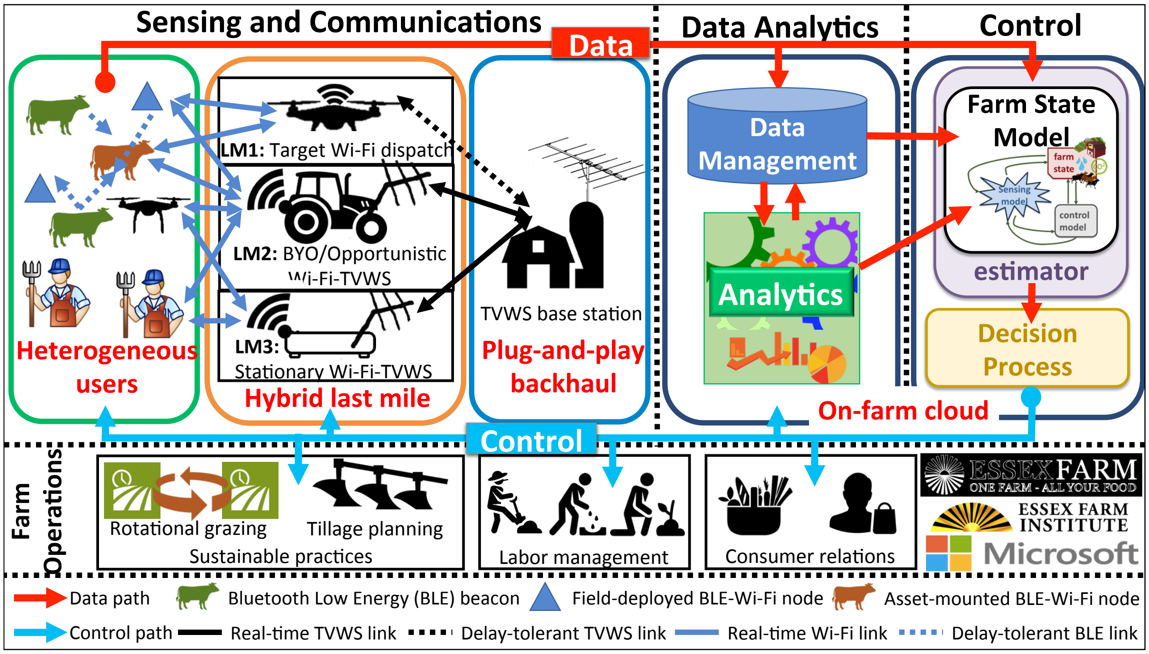

Figure 4: FarmNET architecture.

data collection, and analysis to produce actionable information and enable proactive farming. In what follows, we detail our vision of the desing of each of FarmNET's three key components and discuss their integration into an end-to-end AG-CPS.

4.1 Sensing and Communications

Background and motivation. Connectivity for IoT sensors and farmers is central to an AG-CPS, however, sparsely populated smallholder farmlands are not economically appealing for commercial network deployments, and thus commercial coverage is spotty, inconsistent or lacking altogether. We identify several unique characteristics of smallholder farmlands that lead to such poor commercial connectivity. First of all, smallholder farmlands are characterized with *extremely-low population density*. An average smallholder farm spans a large territory: between 87 and 148 acres [87], rendering its population density orders of magnitude smaller than classical rural scenarios. Our preliminary results (§3) also demonstrate that *network demand of smallholder farms is a unique* mixture of high-volume, spatially-localized and bursty human-generated traffic and low-volume, spatially distributed and predictable IoT sensor traffic. High-volume (and high-revenue) demand that comes from farmers' Internet access is often spatially-localized, however, low-volume demand of IoT sensors is spatially-distributed across the large farmland. To accommodate such demand for low population density over large territory, typical commercial cellular providers face *low return of investments* and thus have low economic incentives.

FarmNET network architecture. To accommodate the heterogeneous farm demand, we envision a hybrid, multi-modal last mile access and a plug-and-play wireless backhaul network architecture, illustrated in the Communications pane of Figure 4. The *hybrid last mile* features three operation modalities (i) stationary, always on hot-spots, (ii) bring-your-own (BYO) opportunistic hot-spots and (iii) drone-mounted hot spots for target dispatch. To enable resource sharing, extensibility and open design, farmers' devices and IoT sensors will be Wi-Fi capable. To provide on-demand Internet access anywhere on the farm, we envision that last mile modalities will connect to a plug-and-play, wide-area backhaul network. Such plug-and-play network can be supplied over TV White Spaces (TVWS) in order to bridge the farm traffic to the Internet and the on-farm cloud without requiring the establishment of on-farm LANs or mesh networks, each of which incurs high deployment and operational cost. To connect to the TVWS network, each last mile modality will need to be equipped with a TVWS client device (TVWS CPE). For the stationary and BYO last mile solutions, the CPE can be collocated with the access point, however, the target dispatch modality would require a delay-tolerant connection, since a TVWS CPE would exceed the drone's load constraints. FarmNET's target dispatch last-mile will be scheduled by control mechanisms (§4.3) to offload sensor data from areas without always-on coverage.

4.1.1 Research challenges. The architecture and operations of our sensing and communications need to be informed by the unique nature of farm traffic demand. The system design objective should be minimization of the traffic distribution time from the network users (i.e. sensors and farmers) to the network gateway, with simultaneous optimization of the power efficiency of key architecture components. These objectives pose fundamental research challenges in cross-layer protocol design, joint last-mile and wireless backhaul optimization and energy-efficient communications.

Cross-layer protocol design. Four major factors affect the traffic distribution time in our architecture: (i) medium access control in the Wi-Fi last mile, (ii) queueing at the Wi-Fi access point (Wi-Fi AP) and the TVWS client, (iii) medium access control in the

TVWS wireless backhaul, and (iv) transmission of redundant sensor data. Thus, we need to speed up the last-mile medium access, allow adaptive queueing and processing of heterogeneous traffic at the last-mile access points, and provide adaptive traffic scheduling over the plug-and-play TVWS backhaul.

Adaptive scheduling of last-mile traffic. Efficiently transmitted multi-user data from the last mile need to be further propagated over the TVWS CPE to the TVWS link and then to the local network gateway. Additional delay components we need to manage in this step are due to (i) data queueing and (ii) scheduling of transmissions over the TVWS link. Current TVWS network standards assume a point-to-multi-point (P2MP) architecture and include IEEE 802.11af [63] and IEEE 802.22 [133]. The former uses Wi-Fi-like contention-based schemes for random medium access control, limiting the effective range of this standard [108]. In the context of TVWS it has been demonstrated for ranges up to 1km [85]. IEEE 802.22 uses a TDD-based channel partitioning that achieves up to 100km communication range [85]. Our AG-CPS context calls for wide-area communication networks, and thus, AG-CPS wireless backhauls should be 802.22-centric.

Traditionally, packet scheduling in TDD networks takes a reactive approach: as queues fill up, the clients request uplink resources in order to transmit back-logged data. This approach incurs delays as traffic is queued waiting for the TDD link to cycle through its frame structure. Our prior observations of sensor data volume, direction and predictability, open space for design of *proactive and adaptive resource allocation schemes* that schedule uplink slots for sensor traffic at the time of traffic arrival. As a result, the TVWS backhaul will provide P2MP communications, while efficiently accommodating sensor traffic fluctuations and scaling back its resources to handle farmer traffic.

Power-efficient sensing and communication. Along with commercial network connectivity, power is the next most scarce resource on farmlands. While electricity may be available in some key locations on a farm, such as the farmhouse, office and barns, it is largely unavailable on the remainder of the farmland. This mandates that AG-CPS designs should *optimize power consumption* and provide *uniform services despite non-uniform availability of power* across the farmland. The unique integrated design of FarmNET posits novel research questions that will further our understanding of AG power consumption, and enable *power-efficient design of AG-CPS.* Power efficiency of communication systems has been studied in multiple domains from cellular network infrastructures [38, 74, 76, 93, 109] to end user devices [105, 112, 113, 118, 126]. Power efficiency has also been considered for wireless sensor networks for agriculture [57, 68], however, existing approaches are limited to the sensing infrastructure and do not consider the effects of system inter-dependencies of an integrated data-driven and control-enabled design. An overarching approach to power efficiency would rely on adaptive duty-cycling through the ON and OFF state of communicating devices, informed by communication demand and the inherent power consumption of the underlying hardware. To this end, two urgent questions emerge: *WHEN* to turn power "off" and "on" and *HOW* to do it. To answer these questions we need to (i) characterize the power needs of individual components both in isolation and jointly,and (ii) integrate components in an energy-aware manner.

4.2 Data Management, Mining and Analytics

Extracting actionable information from farm operations sensed data is key to enabling real-time control and decision making for AG-CPS [101]. FarmNET will sense a variety of data (e.g., animal location and behavior, farming machines and farmers' location and activity, soil state, crops and pasture growth, environmental conditions) and output, e.g., dairy, meat, grains, vegetables, etc. These diverse modalities open novel data management challenges for data acquired by noisy measurements at heterogeneous temporal scales. Moreover, it should enable analytics capabilities in a scalable manner, enabling prioritization according to the delay tolerance of specific farm operations. Another challenge is, that FarmNET will have to assist farmers in answering "what-if" questions by employing simple and accurate *data-driven* models that are incrementally re-trained based on new evidence. The novel aspects of our envisioned FarmNET system data analytics include (i) knowledge extraction that handles sensor uncertainty and sparse and economical sensor deployments suitable for smallholder farms (ii) support for anomaly detection, root-cause analysis and simulation using dynamic network models and (iii) a flexible middle layer between sensing and control that relies on high-level states as opposed to raw sensor readings, thus completing the architecture of an end-to-end CPS.

To manage the tracking data stream and provide analytics in a centralized on-the-farm data cloud, we will adopt Apache Spark [6]. It supports various data sources (SQL, streams, etc.) and fast writes for IoT sensors' streams, and works well with various analytics platforms (MLIB, GarphX). Beyond sensor data, we will import other data sources, such as real-time weather information and forecasts, operational data and manually collected ground truth data. The data repository will support (i) analytics tasks such as temporal trends, year-to-year comparison, anomaly detection and root cause analysis; (ii) simulation and prediction of the farm operations; and (iii) necessary data for the control component of FarmNET to enable optimized communications, sensing and operational outcomes.

Tracking uncertain mobile objects and their activity. Tracking and localization of moving objects using WSNs in various environments have seen tremendous advances, with a variety of sensors, data collection and algorithms proposed [26, 84]. In the agricultural context, such objects include animals, farmers and farming machinery. Successful and affordable tracking of the above will enable effective labor distribution and identification of bottlenecks, precision rotational grazing [80] and sick animal detection and improved farm fleet management. The case of tracking livestock presents its own challenges: large areas to cover and large number of individuals [77, 103]. There are, however, regularities in the herd behavior, e.g., shared preference for specific grass species and tendency to stay together [138], that allow for the cost-effective heterogeneous sensing WSN architecture of FarmNET. In particular, the majority of the herd members (*passive*) will be equipped with low-energy bluetooth beacons (BLE), whose presence can be sensed by a small sub-sample of *active* herd members equipped with Bluetooth transponders, GPS, and radio to offload sensed data to available sinks.

Beyond position, we need to also track the activity of mobile objects. Activity classification of farm animal behavior (standing, grazing, laying, (in)active, etc.) has been considered using accelerometers [59, 70, 89, 90, 132, 144], GPS sensors and signal strength to a gateway within a wireless sensor network [102, 104]. In particular, methods for classifying feeding behavior have achieved high accuracy: precision and sensitivity exceeding 80% for 5-10 minute windows [59]. Beyond instantaneous classification, FarmNET will incorporate predictions of the best performing techniques in a temporal smoothing framework for continuous behavioral smoothness and thus limit the effect of instantaneous mis-classification.

Tracking the state of crops and pastures. Crops, pasture grass species and their growth state comprise another important potential domain for precision farming [138]. FarmNET will predominantly rely on affordable UAV-image sensing as opposed to high-density individual plant sensors. Beyond arable land crop monitoring, which has received the highest attention from both industry and AG ICT research, FarmNET will support precision rotational grazing that requires quantifying the temporal interactions among the pasture biomass, grazing animals, water, nutrients and the effect of weather conditions [80]. Exact estimation of biomass and nutrients over time typically involves labor-intensive sampling and expensive laboratory testing [80]. Low-cost sensors have been predominantly applied to arable land [122, 137] and only recently employed for grassland [120]. Although less precise, the utility of sensing approaches involving images [34, 64, 92] laser [114] and ultrasonic [123] sensors have recently been studied for pasture measurement. Kabir et al. [79] compared the precision for grass growth estimation of ultrasonic, CCD and reflectance sensors mounted on a tractor and concluded that a camera sensor at 90 degrees angle had a dominating performance. Although capable of ensuring a good quality, continuous monitoring using tractor-mounted sensors is costly and infeasible for the temporal resolution (at least once a day) we are envisioning. Instead, FarmNET will utilize a combination of image sensors mounted on an UAV and the sensed behavior of the grazing animals, farmer activity and farming machine-mounted sensors to track the crops and pasture growth over time. Research challenges include: scheduling flight plans (altitude and spatially adaptive) to ensure a matching tracking quality to that of on-the-ground tractor-mounted sensors.

Anomaly detection and root-cause analysis. We will cast our tracking data as a dynamic graph of entities (e.g., fields, crops, animal herds, farmers, machines) with associated multivariate time series as their state over time, and edges corresponding to dependencies. Example entity signals will include: grazing over time, occupancy of animals over time, water due to precipitation and grass level over time. We will support various dynamic network analytics task building on our previous work on global network state classification [54, 55], detection of outlier regions in a single time snapshot [41, 130] and in time periods [44, 98, 99]. The multivariate nature of graph signals comprise the novel research challenges in this task. Unlike existing methods defining anomalies as simple additive score functions [44, 98, 99], we will consider dependencies among signals, leading to novel graph optimization formulations and corresponding algorithms. FarmNET will also support historical trend queries over subgraph regions of interest, de-trending from seasonal weather patterns and detection of unhealthy animals.

Anomaly detection in our dynamic graph setting [44, 98, 99] will direct attention to management zones that behave differently from the rest of the network (e.g., overgrazing, regions with soil erosion). Beyond the identification of such zones in time, we will also be able to perform root-cause analysis by examining anomalies preceding the anomalous outcome, e.g., was the weather abnormal, were animals left grazing longer than prescribed, or were they concentrated in only one location which got overgrazed? Once a suspected cause is at hand, farmers will be able to take the corresponding interventions to alleviate the situation, e.g. add fertilizer, decrease the grazing period in the enclosing paddock and others, schedule tillage or weed removal.

Supporting "What-if" questions via simulations. Detailed dynamic models for the interaction between animal herds, pastures and effects of nutrient cycles and climate have been proposed recently [131, 136]. Such models are characterized by dozens of parameters and typically make assumptions about the homogeneity of grazing behavior and grass and legume growth. Calibrating all parameters for a specific farm operation is not a trivial task. We will employ an alternative data-driven modeling based on high-resolution (both temporal and spatial) tracking data for grazing, crops and pasture growth, weather conditions and nutrient supply. We will model each field node as a queue [27] in which organic matter (grass, vegetables, legumes, etc.) arrive in the queue at a rate dependent on its current state (e.g., grass length), weather conditions and fertilizers. We will learn this independent growth rate using simple linear regression models from past observations. The herd's grazing or crop harvesting within the same node act as a server for the queue and the rate of those processing (i.e. consuming the queue) can also be regressed based on tracking observations. Equipped with (i) the rate functions of growth and (ii) consumption we can simulate the dynamic system under different configurations. For example, we can answer question such as: *What will be the state of a pasture if we doubled the herd; or shorten the grazing interval in the rotational schedule; or if we reconfigure the paddocks? What will happen to a specific crop if there is a consecutive draught period of 20 days?* Being able to simulate realistically such scenarios can enable better decisions as the model will be configured based on observations from the specific farm. In addition, our models can only get more accurate if we incorporate more sensing modalities and increase sensor density.

4.3 Real–Time Farm Monitoring and Control

Knowledge of the state of pastures and crop fields in agriculture is crucial for farmers. As weather patterns change, crops mature, and cattle graze pastures, farmers rely on a combination of experience, visual observation, and intuition as to when to irrigate, apply fertilizer, or move cattle to another pasture. Their decisions are far from optimal. To enable high–output and efficient controlled–environment agriculture technologies and systems, we need a novel, holistic controlled sensing framework that will integrate IoT–sensing capabilities with agriculture data collection, network structure and humans in the loop, to enable real–time accurate monitoring and control. This objective requires (i) a stochastic dynamic system model that completely describes the farm's state (i.e., health and footprint) over time, while incorporating the effect of both cyber (e.g., control signals) and physical components (e.g., agricultural variables), (ii) recursive structured state estimators, and (iii) control strategies that optimize estimation accuracy and costs. Prior work on control for precision agriculture and farm monitoring systems has separately considered the cyber and physical components and proposed ad–hoc [28, 30, 36, 39, 40, 45, 46, 53, 58, 81, 103, 143] or static optimization approaches [67]. An integrated approach will, instead, be based on rigorous **stochastic dynamic** farm system modeling, optimization and control system theories.

Model and Estimators. CPS-assisted farm management relies on realistic modeling of its constituent components and the availability of accurate farm state estimates. We propose to model the health and footprint evolution of a farm as a discrete–time stochastic dynamical system. We will define the farm health as the collective set of a variety of variables (e.g., soil moisture, brightness temperature, vegetation biomass, soil chemical composition, climate, production per unit of harvested area, product feed value) that characterize the production ability of the farm and the quality of its products. In contrast, we will define the farm footprint as a set of variables quantifying the effect on the environment (e.g., carbon dioxide, water consumption, demand for land to feed, soil nutrients quality). Our model will exploit (i) the physics–based models of farm variables' evolution from the literature and §4.2, (ii) the structural and statistical properties of the IoT sensor models, and (ii) the form of control mechanisms in various levels (i.e., adaptation of sensor rates, dispatch of network infrastructure, humans in the loop, farm processes) and the effect of network traffic's statistical properties (§3) and network architecture (§4.1). Beyond modeling, we will develop low-complexity recursive Bayesian estimators of the farm state that account for the network architecture characteristics and traffic's statistical properties, the model structure and humans in the loop.

Joint Optimization of Sensing, Estimation, and Control. FarmNET's system design requires a controlled sensing framework that considers all relevant information to maximize the farm's health, while minimizing its footprint in a dynamically changing environment. Leveraging the dynamic farm model, the system will automatically decide which sensor data to collect and network infrastructure to dispatch to continuously estimate and monitor the farm's health and footprint. Simultaneously, it will automatically control the different farm processes (e.g., tillage scheduling and configuration, herd location scheduling, fertilizer use and optimal lighting) to improve the farm's state. Determining and implementing the exact solution of the outlined optimization problem is a computationally-expensive task. Thus, we plan to design low-complexity, near-optimal sensing, estimation and control strategies that will exploit properties of the AG–CPS, optimal solution properties and farmers' domain knowledge and feedback to enable real-time farm monitoring and control.

5 DISCUSSION AND CONCLUSION

Today, one of the critical problems that humanity is facing is *how to secure nutrition in the face of a changing climate?*. To tackle this problem, we need to increase agricultural production, while dramatically reducing its environmental footprint. Smallholder farms that form the backbone of agricultural production, are thus faced with a tremendous opportunity to expand and proliferate, while solving one of the world's most pressing issues. This opportunity, however, comes with a list of challenges related to understanding and optimization of farm operations. Information and communication technology will play a critical role in such understanding by designing agricultural cyber-physical systems to measure, model, analyze, evaluate and dynamically control the farm state and operations.

Multiple solutions from industry and academia attempt to solve this problem. Each of these two categories suffers its own inherent limitations. Industrial products are for-profit and often strive to provide closed-form solutions. Such solutions are highly-specialized and are not compatible with each other. This requires farmers to purchase various non-extensible systems to provide full-profile AG ICT support, which quickly becomes intractable and financially-infeasible in the smallholder AG context. Academic solutions for AG ICT are often data- or sensor-centric. While such solutions make important progress towards modeling of AG processes, they typically address AG-CPS components in isolation and are not well-fitted for real-time, longitudinal analysis and farm simulations.

While our preliminary analysis and corresponding design is US-centric, we believe that the proposed solution has far-reaching implications on smallholder agriculture in the international context. A key factor to investigating the applicability of the proposed solutions to smallholder farms in the developing world is to follow a modular design approach both (i) horizontally: independent end-to-end infrastructure for various operations within a policulture farm such as animals, vegetables, grain, etc.; and (ii) vertically by ensuring interoperability of varying cost/quality sensing, connectivity and analytics solutions. We see such a design central to further understand the technological, farmer utility and cost challenges in various farmer enterprises. This tradeoff between cost and utility as well as in-depth understanding of the specific needs of farmers in rural areas of the developing world is of key importance when customizing AG-CPS systems for those settings. Our modular design and evaluation of the utility of individual modules in collaboration with Essex Farm will shed light to the above questions of feasibility to the developing world.

Our research makes important headway towards integrated, end-to-end AG-CPS by providing in-depth analysis of farm ICT demand. Our findings indicate that farm demand is a unique mixture of farmer- and sensor-generated traffic with interlocking characteristics. In terms of spatial distribution, farmer traffic is highly-localized, whereas sensor traffic is distributed. Furthermore, while farmer traffic is bursty, downlink-intensive and unpredictable, sensor traffic is periodic, uplink-intensive and predictable. These findings create a unique design space for an integrated AG-CPS dubbed FarmNET. FarmNET consists of three interdependent components including (i) sensing and communication, (ii) AG data analytics and (iii) AG monitoring and control. In this vision paper, we outlined key functional blocks of each component, surveyed the state of the art and outlined an agenda for future development.

An integrated system such as FarmNET is essential to enable practical ICT innovation in agriculture. Such system requires an interdisciplinary approach that brings expertise from sensing and wireless networks, data science, and estimation and control with a constant practitioner feedback in the loop.

REFERENCES

[1] 360 Yield Center. http://www.360yieldcenter.com/Nitrogen.
[2] AgEagle. http://ageagle.com/.
[3] AgGateway. http://www.aggateway.org/.
[4] Aglytix. http://www.aglytix.com/.
[5] Agribotix. http://www.agribotix.com/.
[6] Apache spark http://spark.apache.org/.
[7] AquaSpy. http://www.aquaspy.com/.
[8] aWhere. http://agfundernews.com/awhere-closes-7m-campaign-agfunder.html/.
[9] Climate Corporation (A Monsanto company). https://climate.com/company/.
[10] CropMetrics. http://www.virtualoptimizer.com/.
[11] FarmLogs. http://agfundernews.com/farmlogs-lands-10m-series-b-help-manage-farms.html/.
[12] Farmobile. http://www.farmobile.com/.
[13] Granular. http://www.granular.ag/.
[14] gThrive. http://www.gthrive.com/.
[15] iCropTrak. http://www.icroptrak.com/.
[16] John Deere. http://www.deere.com/.
[17] MyAgCentral. https://www.myagcentral.com/.
[18] OnFarm. http://www.onfarm.com/tag/saas/.
[19] PowWow Energy. https://www.powwowenergy.com/.
[20] SoilSCAPE: Soil Moisture Sensing Controller And oPtimal Estimator. http://soilscape.usc.edu/bootstrap/index.html.
[21] SST Software. http://www.sstsoftware.com/.
[22] A. A. Abbasi and M. Younis. A survey on clustering algorithms for wireless sensor networks. *Computer communications*, 30(14):2826–2841, 2007.
[23] I. F. Akyildiz, T. Melodia, and K. R. Chowdhury. A survey on wireless multimedia sensor networks. *Computer networks*, 51(4):921–960, 2007.
[24] I. F. Akyildiz, W. Su, Y. Sankarasubramaniam, and E. Cayirci. Wireless sensor networks: A survey. *Comput. Netw.*, 38(4):393–422, Mar. 2002.
[25] J. N. Al-Karaki and A. E. Kamal. Routing techniques in wireless sensor networks: a survey. *IEEE wireless communications*, 11(6):6–28, 2004.
[26] T. A. Alhmiedat and S.-H. Yang. A survey: localization and tracking mobile targets through wireless sensors network. 2007.
[27] A. O. Allen. *Probability, statistics, and queueing theory.* Academic Press, 2014.
[28] C. M. Alves-Serodio, J. Monteiro, and C. Couto. An integrated network for agricultural management applications. In *Industrial Electronics, 1998. Proceedings. ISIE'98. IEEE International Symposium on*, volume 2, pages 679–683. IEEE, 1998.
[29] G. Anastasi, M. Conti, M. Di Francesco, and A. Passarella. Energy conservation in wireless sensor networks: A survey. *Ad hoc networks*, 7(3):537–568, 2009.
[30] D. Anurag, S. Roy, and S. Bandyopadhyay. Agro-sense: Precision agriculture using sensor-based wireless mesh networks. In *Innovations in NGN: Future Network and Services, 2008. K-INGN 2008. First ITU-T Kaleidoscope Academic Conference*, pages 383–388. IEEE, 2008.
[31] T. Arampatzis, J. Lygeros, and S. Manesis. A survey of applications of wireless sensors and wireless sensor networks. In *Intelligent Control, 2005. Proceedings of the 2005 IEEE International Symposium on, Mediterrean Conference on Control and Automation*, pages 719–724. IEEE, 2005.
[32] G. K. Atia, V. V. Veeravalli, and J. A. Fuemmeler. Sensor Scheduling for Energy-Efficient Target Tracking in Sensor Networks. *IEEE Trans. Signal Process.*, 59(10):4923–4937, Oct. 2011.
[33] A. Baggio. Wireless sensor networks in precision agriculture. In *ACM Workshop on Real-World Wireless Sensor Networks (REALWSN 2005), Stockholm, Sweden*. Citeseer, 2005.
[34] G. Bareth, A. Bolten, J. Hollberg, H. Aasen, A. Burkart, and J. Schellberg. Feasibility study of using non-calibrated uav-based rgb imagery for grassland monitoring: case study at the rengen long-term grassland experiment (rge), germany. *DGPF Tagungsband*, 24(2015):1–7, 2015.
[35] E. Ben-Dor and A. Banin. Near-infrared analysis as a rapid method to simultaneously evaluate several soil properties. *Soil Science Society of America Journal*, 59(2):364–372, 1995.
[36] L. Bencini, F. Chiti, G. Collodi, D. Di Palma, R. Fantacci, A. Manes, and G. Manes. Agricultural monitoring based on wireless sensor network technology: Real long life deployments for physiology and pathogens control. In *Sensor Technologies and Applications, 2009. SENSORCOMM'09. Third International Conference on*, pages 372–377. IEEE, 2009.
[37] P. Bhagwat, B. Raman, and D. Sanghi. Turning 802.11 inside-out. *SIGCOMM Comput. Commun. Rev.*, 34(1):33–38, Jan. 2004.
[38] S. Bhaumik, G. Narlikar, S. Chattopadhyay, and S. Kanugovi. Breathe to stay cool: Adjusting cell sizes to reduce energy consumption. In *Proceedings of the First ACM SIGCOMM Workshop on Green Networking*, Green Networking '10, pages 41–46, New York, NY, USA, 2010. ACM.
[39] G. Bishop-Hurley, D. Swain, D. Anderson, and P. Corke. Animal Control-What constitutes a reliable cue to stop animal movement? 2006.
[40] G. Bishop-Hurley, D. L. Swain, D. Anderson, P. Sikka, C. Crossman, and P. Corke. Virtual fencing applications: Implementing and testing an automated cattle control system. *Computers and Electronics in Agriculture*, 56(1):14–22, 2007.
[41] P. Bogdanov, B. Baumer, P. Basu, A. Bar-Noy, and A. K. Singh. As strong as the weakest link: Mining diverse cliques in weighted graphs. In *Proceedings of Machine Learning and Knowledge Discovery in Databases - European Conference (ECML/PKDD)*. Springer, 2013.
[42] P. Bogdanov, M. Busch, J. Moehlis, A. K. Singh, and B. K. Szymanski. The social media genome: Modeling individual topic-specific behavior in social media. In *Proceedings of the IEEE/ACM International Conference on Advances in Social Networks Analysis and Mining (ASONAM)*. ACM, 2013.
[43] P. Bogdanov, M. Busch, J. Moehlis, A. K. Singh, and B. K. Szymanski. The social media genome: Modeling individual topic-specific behavior in social media. In *Journal of Social Network Analysis and Mining (SNAM)*. Springer, 2014.
[44] P. Bogdanov, M. Mongiovi, and A. K. Singh. Mining heavy subgraphs in time-evolving networks. In *Proceedings of the IEEE International Conference on Data Mining (ICDM)*, 2011.
[45] Z. Butler, P. Corke, R. Peterson, and D. Rus. Virtual fences for controlling cows. In *Robotics and Automation, 2004. Proceedings. ICRA'04. 2004 IEEE International Conference on*, volume 5, pages 4429–4436. IEEE, 2004.
[46] Z. Butler, P. Corke, R. Peterson, and D. Rus. From robots to animals: virtual fences for controlling cattle. *The International Journal of Robotics Research*, 25(5-6):485–508, 2006.
[47] N. Cao, S. Choi, E. Masazade, and P. K. Varshney. Sensor selection for target tracking in wireless sensor networks with uncertainty. *IEEE Transactions on Signal Processing*, 64(20):5191–5204, 2016.
[48] A. Chedad, D. Moshou, J.-M. Aerts, A. Van Hirtum, H. Ramon, and D. Berckmans. Ap—animal production technology: recognition system for pig cough based on probabilistic neural networks. *Journal of agricultural engineering research*, 79(4):449–457, 2001.
[49] Y. Chen, Q. Zhao, V. Krishnamurthy, and D. Djonin. Transmission scheduling for optimizing sensor network lifetime: A stochastic shortest path approach. *IEEE Transactions on Signal Processing*, 55(5):2294–2309, 2007.
[50] E. K. Chong, C. M. Kreucher, and A. O. Hero. Partially observable Markov decision process approximations for adaptive sensing. *Discrete Event Dynamic Systems*, 19(3):377–422, 2009.
[51] Y. Chung, S. Oh, J. Lee, D. Park, H.-H. Chang, and S. Kim. Automatic detection and recognition of pig wasting diseases using sound data in audio surveillance systems. *Sensors*, 13(10):12929–12942, 2013.
[52] U. S. Committee. New Cooperation for Global Food Security. http://www.un.org/press/en/2009/gaef3242.doc.htm.
[53] S. Cugati, W. Miller, and J. Schueller. Automation concepts for the variable rate fertilizer applicator for tree farming. In *The Proceedings of the 4th European Conference in Precision Agriculture, Berlin, Germany, June*, pages 14–19, 2003.
[54] X.-H. Dang, A. K. Singh, P. Bogdanov, H. You, and B. Hsu. Discriminative subnetworks with regularized spectral learning for global-state network data. In *Proceedings of the 25th European Conference on Machine Learning / 18th European Conference on Principles and Practice of Knowledge Discovery in Databases (ECML/PKDD)*. Springer, 2014.
[55] X.-H. Dang, H. You, P. Bogdanov, and A. Singh. Learning predictive substructures with regularization for network data. In *Proceedings of the IEEE International Conference on Data Mining (ICDM)*, 2015.
[56] I. Demirkol, C. Ersoy, and F. Alagoz. Mac protocols for wireless sensor networks: a survey. *IEEE Communications Magazine*, 44(4):115–121, 2006.
[57] I. Demirkol, C. Ersoy, and E. Onur. Wake-up receivers for wireless sensor networks: benefits and challenges. *IEEE Wireless Communications*, 16(4):88–96, Aug 2009.
[58] S. E. Díaz, J. C. Pérez, A. C. Mateos, M.-C. Marinescu, and B. B. Guerra. A novel methodology for the monitoring of the agricultural production process based on wireless sensor networks. *Computers and Electronics in Agriculture*, 76(2):252–265, 2011.
[59] J. A. V. Diosdado, Z. E. Barker, H. R. Hodges, J. R. Amory, D. P. Croft, N. J. Bell, and E. A. Codling. Classification of behaviour in housed dairy cows using an accelerometer-based activity monitoring system. *Animal Biotelemetry*, 3(1):15, 2015.
[60] A. Dobermann and R. Nelson. Opportunities and solutions for sustainable food production. 2013.
[61] C. Eastwood, S. Kenny, et al. Art or science?: Heuristic versus data driven grazing management on dairy farms. *Extension Farming Systems Journal*, 5(1):95, 2009.
[62] D. Estrin, L. Girod, G. Pottie, and M. Srivastava. Instrumenting the world with wireless sensor networks. In *Acoustics, Speech, and Signal Processing, 2001. Proceedings.(ICASSP'01). 2001 IEEE International Conference on*, volume 4, pages 2033–2036. IEEE, 2001.
[63] A. B. Flores, R. E. Guerra, E. W. Knightly, P. Ecclesine, and S. Pandey. Ieee 802.11af: a standard for tv white space spectrum sharing. *IEEE Communications Magazine*, 51(10):92–100, October 2013.
[64] C. Francone, V. Pagani, M. Foi, G. Cappelli, and R. Confalonieri. Comparison of leaf area index estimates by ceptometer and pocketlai smart app in canopies with different structures. *Field Crops Research*, 155:38–41, 2014.

[65] J. A. Fuemmeler, G. K. Atia, and V. V. Veeravalli. Sleep control for tracking in sensor networks. *IEEE Transactions on Signal Processing*, 59(9):4354–4366, 2011.

[66] 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*, 327(5967):812–818, 2010.

[67] J. C. Groot, G. J. Oomen, and W. A. Rossing. Multi-objective optimization and design of farming systems. *Agricultural Systems*, 110:63–77, 2012.

[68] L. Gu and J. A. Stankovic. Radio-triggered wake-up capability for sensor networks. In *Proceedings. RTAS 2004. 10th IEEE Real-Time and Embedded Technology and Applications Symposium, 2004.*, pages 27–36, May 2004.

[69] S. Guo, M. H. Falaki, E. A. Oliver, S. U. Rahman, A. Seth, M. A. Zaharia, U. Ismail, and S. Keshav. Design and implementation of the kiosknet system. In *2007 International Conference on Information and Communication Technologies and Development*, pages 1–10, Dec 2007.

[70] D. Hanson and C. Mo. Monitoring cattle motion using 3-axis acceleration and gps data. *Journal of Research in Agriculture and Animal Science*, 2(10):1–8, 2014.

[71] K. A. Harris and V. V. Veeravalli. Implementing energy-efficient tracking in a sensor network. In *Acoustics, Speech and Signal Processing (ICASSP), 2013 IEEE International Conference on*, pages 4608–4612. IEEE, 2013.

[72] S. Hasan, K. Heimerl, K. Harrison, K. Ali, S. Roberts, A. Sahai, and E. Brewer. Gsm whitespaces: An opportunity for rural cellular service. In *Dynamic Spectrum Access Networks (DYSPAN), 2014 IEEE International Symposium on*, pages 271–282, April 2014.

[73] Y. He and E. K. Chong. Sensor scheduling for target tracking: A Monte Carlo sampling approach. *Digital Signal Processing*, 16(5):533–545, 2006.

[74] K. Heimerl, K. Ali, J. Blumenstock, B. Gawalt, and E. Brewer. Expanding rural cellular networks with virtual coverage. In *Proceedings of the 10th USENIX Conference on Networked Systems Design and Implementation*, nsdi'13, pages 283–296, Berkeley, CA, USA, 2013. USENIX Association.

[75] K. Heimerl et al. Local, sustainable, small-scale cellular networks. ICTD13, Cape Town, South Africa, 2013.

[76] K. Heimerl, S. Hasan, K. Ali, T. Parikh, and E. Brewer. An experiment in reducing cellular base station power draw with virtual coverage. In *Proceedings of the 4th Annual Symposium on Computing for Development*, ACM DEV-4 '13, pages 6:1–6:9, New York, NY, USA, 2013. ACM.

[77] J. I. Huircán, C. Muñoz, H. Young, L. Von Dossow, J. Bustos, G. Vivallo, and M. Toneatti. Zigbee-based wireless sensor network localization for cattle monitoring in grazing fields. *Computers and Electronics in Agriculture*, 74(2):258–264, 2010.

[78] I. F. P. R. Institute. Smallholder farming. http://www.ifpri.org/topic/smallholder-farming.

[79] M. S. N. KABIR, S.-O. CHUNG, K. Yong-Joo, L. Geung-Joo, Y. Seung-Hwa, L. Kyeong-Hwan, T. OKAYASU, and E. INOUE. Sensor comparison for grass growth estimation. *J. Fac. Agr., Kyushu Univ*, 61(2):367–374, 2016.

[80] R. L. Kallenbach. Describing the dynamic: Measuring and assessing the value of plants in the pasture. *Crop Science*, 55(6):2531–2539, 2015.

[81] Y. Kim, R. G. Evans, and W. M. Iversen. Remote sensing and control of an irrigation system using a distributed wireless sensor network. *IEEE transactions on instrumentation and measurement*, 57(7):1379–1387, 2008.

[82] C. Krintz, R. Wolski, N. Golubovic, B. Lampel, V. Kulkarni, B. Roberts, and B. Liu. Smartfarm: Improving agriculture sustainability using modern information technology. *ACM SIGKDD DSFEW*, 2016.

[83] V. Krishnamurthy and D. V. Djonin. Optimal threshold policies for multivariate POMDPs in radar resource management. *IEEE transactions on Signal Processing*, 57(10):3954–3969, 2009.

[84] N. D. Larusso and A. Singh. Efficient tracking and querying for coordinated uncertain mobile objects. In *Data Engineering (ICDE), 2013 IEEE 29th International Conference on*, pages 182–193. IEEE, 2013.

[85] D. Lekomtcev and R. Marsalek. Comparison of 802.11af and 802.22 standards – physical layer and cognitive functionality. In *Electro Revue, VOL. 3, NO. 2, June 2012*.

[86] Y. Li, L. W. Krakow, E. K. Chong, and K. N. Groom. Approximate stochastic dynamic programming for sensor scheduling to track multiple targets. *Digital Signal Processing*, 19(6):978–989, 2009.

[87] S. K. Lowder, J. Skoet, and T. Raney. The number, size, and distribution of farms, smallholder farms, and family farms worldwide. *World Development*, 87:16–29, November, 2016.

[88] J. M. MacDonald. Family Farming in the United States. https://www.ers.usda.gov/amber-waves/2014/march/family-farming-in-the-united-states/.

[89] P. Martiskainen, M. Järvinen, J.-P. Skön, J. Tiirikainen, M. Kolehmainen, and J. Mononen. Cow behaviour pattern recognition using a three-dimensional accelerometer and support vector machines. *Applied animal behaviour science*, 119(1):32–38, 2009.

[90] A. Mason and J. Sneddon. Automated monitoring of foraging behaviour in free ranging sheep grazing a biodiverse pasture. In *Sensing Technology (ICST), 2013 Seventh International Conference on*, pages 46–51. IEEE, 2013.

[91] K. W. Matthee, G. Mweemba, A. V. Pais, G. van Stam, and M. Rijken. Bringing internet connectivity to rural zambia using a collaborative approach. In *2007 International Conference on Information and Communication Technologies and Development*, pages 1–12, Dec 2007.

[92] P. McEntee, R. Belford, R. Mandel, J. Harper, and M. Trotter. Sub-paddock scale spatial variability between the pasture and cropping phases of mixed farming systems in australia. In *Precision agriculture'13*, pages 389–394. Springer, 2013.

[93] N. Mishra, K. Chebrolu, B. Raman, and A. Pathak. Wake-on-wlan. In *Proceedings of the 15th International Conference on World Wide Web*, WWW '06, pages 761–769, New York, NY, USA, 2006. ACM.

[94] M. Moghaddam. A Robust Wireless Sensor Network Architecture for the Large-scale Deployment of the Soil Moisture Sensing Controller and Optimal Estimator (SoilSCaPE). In *American Geophysical Union Fall Conference*, 2011.

[95] M. Moghaddam, D. Entekhabi, Y. Goykhman, K. Li, M. Liu, A. Mahajan, A. Nayyar, D. Shuman, and D. Teneketzis. A wireless soil moisture smart sensor web using physics-based optimal control: Concept and initial demonstrations. *IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing*, 3(4):522–535, 2010.

[96] M. Moghaddam, M. Liu, X. Wu, M. Burgin, Y. Goykhman, Q. Wang, D. Shuman, A. Nayyar, D. Teneketzis, and D. Entekhabi. Soil Moisture Sensing Controller and Optimal Estimator (SoilSCaPE): An in-situ Wireless Sensor Network for Validation of Spaceborne Soil Moisture Estimates. In *American Geophysical Union Fall Conference*, 2010.

[97] M. Moghaddam, A. Silva, R. Akbar, D. Clewley, M. Burgin, A. Castillo, and D. Entekhabi. SoilSCAPE In situ Network for Multiscape Validation of SMAP Data Products. In *2013 IEEE International Geoscience and Remote Sensing Symposium*, 2013.

[98] M. Mongiovi, P. Bogdanov, R. Ranca, A. K. Singh, E. E. Papalexakis, and C. Faloutsos. Netspot: Spotting significant anomalous regions on dynamic networks. In *Proceedings of SIAM International Conference on Data Mining (SDM)*, 2013.

[99] M. Mongiovi, P. Bogdanov, and A. K. Singh. Mining evolving network processes. In *Proceedings of the IEEE International Conference on Data Mining (ICDM)*. IEEE, 2013.

[100] I. D. Moore, P. Gessler, G. Nielsen, and G. Peterson. Soil attribute prediction using terrain analysis. *Soil Science Society of America Journal*, 57(2):443–452, 1993.

[101] A. Mucherino, P. J. Papajorgji, and P. M. Pardalos. *Data mining in agriculture*, volume 34. Springer Science & Business Media, 2009.

[102] E. S. Nadimi and H. Søgaard. Observer kalman filter identification and multiple-model adaptive estimation technique for classifying animal behaviour using wireless sensor networks. *Computers and Electronics in Agriculture*, 68(1):9–17, 2009.

[103] E. S. Nadimi, H. Søgaard, T. Bak, and F. W. Oudshoorn. ZigBee-based wireless sensor networks for monitoring animal presence and pasture time in a strip of new grass. *Computers and electronics in agriculture*, 61(2):79–87, 2008.

[104] E. S. Nadimi, H. T. Søgaard, and T. Bak. Zigbee-based wireless sensor networks for classifying the behaviour of a herd of animals using classification trees. *Biosystems engineering*, 100(2):167–176, 2008.

[105] S. Nedevschi, J. Chandrashekar, J. Liu, B. Nordman, S. Ratnasamy, and N. Taft. Skilled in the art of being idle: Reducing energy waste in networked systems. In *Proceedings of the 6th USENIX Symposium on Networked Systems Design and Implementation*, NSDI'09, pages 381–394, Berkeley, CA, USA, 2009. USENIX Association.

[106] S. Nitinawarat, G. K. Atia, and V. V. Veeravalli. Efficient target tracking using mobile sensors. In *Computational Advances in Multi-Sensor Adaptive Processing (CAMSAP), 2011 4th IEEE International Workshop on*, pages 405–408. IEEE, 2011.

[107] M. Z. others. Kwiizya: local cellular network services in remote areas. ACM MobiSys13, Taipei, Taiwan, 2013.

[108] R. Patra, S. Nedevschi, S. Surana, A. Sheth, L. Subramanian, and E. Brewer. Wildnet: Design and implementation of high performancewifi based long distance networks. In *Proceedings of the 4th USENIX Conference on Networked Systems Design and Implementation*, NSDI'07, pages 7–7, Berkeley, CA, USA, 2007. USENIX Association.

[109] C. Peng, S.-B. Lee, S. Lu, H. Luo, and H. Li. Traffic-driven power saving in operational 3g cellular networks. In *Proceedings of the 17th Annual International Conference on Mobile Computing and Networking*, MobiCom '11, pages 121–132, New York, NY, USA, 2011. ACM.

[110] A. Pentland, R. Fletcher, and A. Hasson. Daknet: rethinking connectivity in developing nations. *Computer*, 37(1):78–83, Jan 2004.

[111] F. Pierce and T. Elliott. Regional and on-farm wireless sensor networks for agricultural systems in eastern washington. *Computers and electronics in agriculture*, 61(1):32–43, 2008.

[112] F. Qian, Z. Wang, A. Gerber, Z. M. Mao, S. Sen, and O. Spatscheck. Characterizing radio resource allocation for 3g networks. In *Proceedings of the 10th ACM SIGCOMM Conference on Internet Measurement*, IMC '10, pages 137–150, New York, NY, USA, 2010. ACM.

[113] F. Qian, Z. Wang, A. Gerber, Z. M. Mao, S. Sen, and O. Spatscheck. Top: Tail optimization protocol for cellular radio resource allocation. In *The 18th IEEE International Conference on Network Protocols*, pages 285–294, Oct 2010.

[114] P. J. Radtke, H. T. Boland, and G. Scaglia. An evaluation of overhead laser scanning to estimate herbage removals in pasture quadrats. *Agricultural and forest meteorology*, 150(12):1523–1528, 2010.
[115] M. Rahimi, M. Hansen, W. J. Kaiser, G. S. Sukhatme, and D. Estrin. Adaptive sampling for environmental field estimation using robotic sensors. In *Intelligent Robots and Systems, 2005.(IROS 2005). 2005 IEEE/RSJ International Conference on*, pages 3692–3698. IEEE, 2005.
[116] B. Raman and K. Chebrolu. Experiences in using wifi for rural internet in india. *IEEE Communications Magazine*, 45(1):104–110, Jan 2007.
[117] Rhizomatica. Rhizomatica — Mobile Communications for All. rhizomatica.org/. [Online; accessed 4-November-2016].
[118] S. Rosen, H. Luo, Q. A. Chen, Z. M. Mao, J. Hui, A. Drake, and K. Lau. Discovering fine-grained rrc state dynamics and performance impacts in cellular networks. In *Proceedings of the 20th Annual International Conference on Mobile Computing and Networking*, MobiCom '14, pages 177–188, New York, NY, USA, 2014. ACM.
[119] G. Ruß and A. Brenning. Data mining in precision agriculture: management of spatial information. In *International Conference on Information Processing and Management of Uncertainty in Knowledge-Based Systems*, pages 350–359. Springer, 2010.
[120] J. Schellberg and E. Verbruggen. Frontiers and perspectives on research strategies in grassland technology. *Crop and Pasture Science*, 65(6):508–523, 2014.
[121] P. Schmitt, D. Iland, M. Zheleva, and E. Belding. HybridCell: Cellular connectivity on the fringes with demand-driven local cells. In *IEEE INFOCOM '16*, San Francisco, CA, USA, April 2016.
[122] E. W. Schuster, S. Kumar, S. E. Sarma, J. L. Willers, and G. A. Milliken. Infrastructure for data-driven agriculture: identifying management zones for cotton using statistical modeling and machine learning techniques. In *Emerging Technologies for a Smarter World (CEWIT), 2011 8th International Conference & Expo on*, pages 1–6. IEEE, 2011.
[123] D. K. Shannon, J. Lory, R. Kallenbach, T. Lorenz, J. Harper, G. Schmitz, W. Rapp, B. Carpenter, and D. England. Initial results utilizing a commercially available ultrasonic sensor for forage yield measurements. In *Proc. Am. Soc. Agric. Biol. Eng*. Am. Soc. Agric. Biol. Eng, St. Joseph, MI, 2013.
[124] X. Shen, S. Liu, and P. K. Varshney. Sensor selection for nonlinear systems in large sensor networks. *IEEE Transactions on Aerospace and Electronic Systems*, 50(4):2664–2678, 2014.
[125] X. Shen and P. K. Varshney. Sensor selection based on generalized information gain for target tracking in large sensor networks. *IEEE Transactions on Signal Processing*, 62(2):363–375, 2014.
[126] E. Shih, P. Bahl, and M. J. Sinclair. Wake on wireless: An event driven energy saving strategy for battery operated devices. In *Proceedings of the 8th Annual International Conference on Mobile Computing and Networking*, MobiCom '02, pages 160–171, New York, NY, USA, 2002. ACM.
[127] D. I. Shuman. *From sleeping to stockpiling: Energy conservation via stochastic scheduling in wireless networks*. PhD thesis, The University of Michigan, 2010.
[128] D. I. Shuman, M. Liu, and O. Q. Wu. Energy-efficient transmission scheduling with strict underflow constraints. *IEEE Transactions on Information Theory*, 57(3):1344–1367, 2011.
[129] D. I. Shuman, A. Nayyar, A. Mahajan, Y. Goykhman, K. Li, M. Liu, D. Teneketzis, M. Moghaddam, and D. Entekhabi. Measurement scheduling for soil moisture sensing: From physical models to optimal control. *Proceedings of the IEEE*, 98(11):1918–1933, 2010.
[130] A. Silva, P. Bogdanov, and A. Singh. Hierarchical in-network attribute compression via importance sampling. In *Proceedings of the 31st IEEE International Conference on Data Engineering (ICDE)*, 2015.
[131] J.-F. Soussana, A.-I. Graux, and F. N. Tubiello. Improving the use of modelling for projections of climate change impacts on crops and pastures. *Journal of experimental botany*, 61(8):2217–2228, 2010.
[132] A. Spink, B. Cresswell, A. Kölzsch, F. van Langevelde, M. Neefjes, L. Noldus, H. van Oeveren, H. Prins, T. van der Wal, N. de Weerd, et al. Animal behaviour analysis with gps and 3d accelerometers. In *Precision livestock farming, 10-12 September, 2013, Leuven, Belgium*, pages 229–239, 2013.
[133] C. R. Stevenson, G. Chouinard, Z. Lei, W. Hu, S. J. Shellhammer, and W. Caldwell. Ieee 802.22: The first cognitive radio wireless regional area network standard. *IEEE Communications Magazine*, 47(1):130–138, January 2009.
[134] L. Subramanian, S. Nedevschi, M. Ho, E. Brewer, and A. Sheth. Rethinking wireless for the developing world. In *In Hotnets-V*, 2006.
[135] S. Surana, R. Patra, S. Nedevschi, M. Ramos, L. Subramanian, Y. Ben-David, and E. Brewer. Beyond pilots: Keeping rural wireless networks alive. In *Proceedings of the 5th USENIX Symposium on Networked Systems Design and Implementation*, NSDI'08, pages 119–132, Berkeley, CA, USA, 2008. USENIX Association.
[136] J. J. Tewa, A. Bah, and S. C. Oukouomi Noutchie. Dynamical models of interactions between herds forage and water resources in sahelian region. In *Abstract and Applied Analysis*, volume 2014. Hindawi Publishing Corporation, 2014.
[137] J. Torres-Sánchez, J. Peña, A. De Castro, and F. López-Granados. Multi-temporal mapping of the vegetation fraction in early-season wheat fields using images from uav. *Computers and Electronics in Agriculture*, 103:104–113, 2014.
[138] D. Undersander, B. Albert, D. Cosgrove, D. Johnson, and P. Peterson. *Pastures for profit: A guide to rotational grazing*. 1997.
[139] E. R. Vivoni and R. Camilli. Real-time streaming of environmental field data. *Computers & Geosciences*, 29(4):457–468, 2003.
[140] N. Wang, N. Zhang, and M. Wang. Wireless sensors in agriculture and food industry—recent development and future perspective. *Computers and electronics in agriculture*, 50(1):1–14, 2006.
[141] T. Wark, P. Corke, P. 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*, 6(2), 2007.
[142] J. Yick, B. Mukherjee, and D. Ghosal. Wireless sensor network survey. *Computer networks*, 52(12):2292–2330, 2008.
[143] Q. Zhang, X.-l. Yang, Y.-m. Zhou, L.-r. Wang, and X.-s. Guo. A wireless solution for greenhouse monitoring and control system based on ZigBee technology. *Journal of Zhejiang University-Science A*, 8(10):1584–1587, 2007.
[144] H. Zhou, L. Yin, and C. Liu. Dairy cattle movement detecting technology using support vector machine. In *International Conference on Wireless Communications and Applications*, pages 23–32. Springer, 2011.
[145] D. Zois and U. Mitra. Active State Tracking with Sensing Costs: Analysis of Two–States and Methods for n–States. *Transactions on Signal Processing*, 2017.
[146] D. S. Zois, U. Demiryurek, and U. Mitra. A POMDP approach for active collision detection via networked sensors. In *Asilomar Conference on Signals, Systems, and Computers*, Nov. 2016.
[147] D.-S. Zois, M. Levorato, and U. Mitra. A POMDP Framework for Heterogeneous Sensor Selection in Wireless Body Area Networks. In *Proc. 31st IEEE International Conference on Computer Communications (INFOCOM)*, pages 2611–2615, March 2012.
[148] D.-S. Zois, M. Levorato, and U. Mitra. Heterogeneous Time-Resource Allocation in Wireless Body Area Networks for Green, Maximum Likelihood Activity Detection. In *Proc. IEEE International Conference on Communications (ICC)*, pages 3448–3452, June 2012.
[149] D.-S. Zois, M. Levorato, and U. Mitra. Energy–Efficient, Heterogeneous Sensor Selection for Physical Activity Detection in Wireless Body Area Networks. *IEEE Tansactions on Signal Processing*, 61(7):1581–1594, April 2013.
[150] D.-S. Zois, M. Levorato, and U. Mitra. Kalman-like state tracking and control in POMDPs with applications to body sensing networks. In *IEEE International Conference on Acoustics, Speech, and Signal Processing (ICASSP)*, May 2013.
[151] D.-S. Zois, M. Levorato, and U. Mitra. Non–linear smoothers for discrete–time, finite–state Markov chains. In *IEEE International International Symposium on Information Theory (ISIT)*, July 2013.
[152] D. S. Zois, M. Levorato, and U. Mitra. Active Classification for POMDPs: A Kalman-Like State Estimator. *IEEE Transactions on Signal Processing*, 62(23):6209–6224, Dec. 2014.
[153] D.-S. Zois and U. Mitra. On the properties of nonlinear POMDPs for active state tracking. In *IEEE Global Conference on Signal and Information Processing (GlobalSIP)*, pages 193–196, December 2013.
[154] D.-S. Zois and U. Mitra. A Weiss–Weinstein Lower Bound Based Sensing Strategy for Active State Tracking. In *IEEE International International Symposium on Information Theory (ISIT)*, July 2014.
[155] D.-S. Zois and U. Mitra. Controlled Sensing: A Myopic Fisher Information Sensor Selection Algorithm. In *IEEE Globecom*, Dec. 2014.
[156] D. S. Zois and M. Raginsky. Active object detection on graphs via locally informative trees. In *2016 IEEE 26th International Workshop on Machine Learning for Signal Processing (MLSP)*, pages 1–6, 2016.