



Article

The Design and Deployment of an End-To-End IoT Infrastructure for the Natural Environment

Vatsala Nundloll ^{1,*}, Barry Porter ^{1,*}, Gordon S. Blair ^{1,*}, Bridget Emmett ², Jack Cosby ², Davey L. Jones ^{3,4}, Dave Chadwick ³, Ben Winterbourn ², Philip Beattie ¹, Graham Dean ¹, Rory Shaw ³, Wayne Shelley ⁵, Mike Brown ⁶ and Izhar Ullah ¹

- ¹ School of Computing and Communications, Lancaster University, Lancaster LA1 4YW, UK; philbeattie2@gmail.com (P.B.); g.dean1@lancaster.ac.uk (G.D.); izhar.ullah@uetpeshawar.edu.pk (I.U.)
- ² Centre for Ecology & Hydrology, Environment Centre Wales, Bangor LL57 2UW, UK; bae@ceh.ac.uk (B.E.); jaccos@ceh.ac.uk (J.C.); benwinterbourn@gmail.com (B.W.)
- ³ School of Natural Sciences, Bangor University, Deiniol Road, Bangor, Gwynedd LL57 2UW, UK; d.jones@bangor.ac.uk (D.L.J.); d.chadwick@bangor.ac.uk (D.C.); roryshaw@hotmail.com (R.S.)
- ⁴ UWA School of Agriculture and Environment, University of Western Australia, Crawley WA 6009, Australia
- ⁵ British Geological Survey, Nottingham NG12 5GG, UK; wael@bgs.ac.uk
- ⁶ Centre for Ecology & Hydrology, Lancaster Environment Centre, Lancaster LA1 4AP, UK; mjbr@ceh.ac.uk
- * Correspondence: vatsala@lancaster.ac.uk (V.N.); b.f.porter@lancaster.ac.uk (B.P.); gordon.s.blair@googlemail.com (G.S.B.)

Received: 11 April 2019; Accepted: 4 June 2019; Published: 7 June 2019



Abstract: Internet of Things (IoT) systems have seen recent growth in popularity for city and home environments. We report on the design, deployment, and use of the IoT infrastructure for environmental monitoring and management. Working closely with hydrologists, soil scientists, and animal behaviour scientists, we successfully deployed and utilised a system to deliver integrated information across these two fields in the first such example of real-time multidimensional environmental science. We describe the design of this system; its requirements and operational effectiveness for hydrological, soil, and ethological scientists; and our experiences from building, maintaining, and using the deployment at a remote site in difficult conditions. Based on this experience, we discuss key future work for the IoT community when working in these kinds of environmental deployments.

Keywords: wireless sensor networks; Internet of Things; environmental science; experiences

1. Introduction

The Internet of Things (IoT) represents the next major step for the Internet as it evolves from a communication substrate that connects computers to one that connects and embraces everyday objects (things). This has the potential to revolutionize many different sectors of the economy and society more generally, e.g., enabling smart cities, smart transport systems, intelligent management of energy supplies, etc., all enabled by data collection from sensors. Most research in the Internet of Things has been carried out in cities and urban areas more generally (see Section 3).

We see great potential and, indeed in many ways, an even bigger potential for transformation in applying Internet of Things concepts to rural areas to both gain a better understanding of the natural environment and to manage this environment through appropriate interventions. However, the deployment of Internet of Things technology in the natural environment poses new challenges around, for example, the limitations of power and Internet connectivity and hostile operational conditions.

The focus of this paper is on the Environmental IoT project, a first attempt to instrument and manage an environmental catchment in all its facets, across different geographical locations and at all its scales. We also take this a stage further, building on our experiences in the Environmental Virtual Observatory pilot project [1]. This project utilised cloud computing to offer a shared repository of data, models, and other tools and artifacts to allow a range of stakeholders to visualise data, to run models, to feed data into models, to deal with uncertainty in models, and to discuss the results with other stakeholders and communities. The Internet of Things and cloud computing are strongly complementary technologies that, when combined, can provide a complete, end-to-end technological infrastructure for a step change in our understanding of complex environmental factors.

This research has the potential to have major impact on many aspects of rural life—on farmers and associated agricultural businesses, the water industry, tourists and tourism related businesses, and society more generally. This has the potential to completely transform these associated businesses, enabling critical areas such as integrated land and water management, coastal zone protection, and precision agriculture. We are particularly interested though in new kinds of science and associated management strategies that stem from bringing real-time data sets together and from observing related interdependencies.

Current practices in the environmental and earth sciences focus heavily on standalone data logger systems, with some early initiatives starting to embrace wireless sensor networks for environmental modelling (e.g., References [2–5]). These initiatives, though, tend to focus on particular environmental facets at particular scales, e.g., focusing on habitat monitoring [6,7], glaciology [8], permafrost [9], and volcanoes [10]. While wireless networks have become quite sophisticated in many cases, they need a step change in their interoperability, scope, and usability to become functioning Environmental IoT's, which also need to encompass a widespread deployment of spatially distributed devices with embedded identification, sensing, and/or actuation capabilities [11]. The development of a full Environmental IoT would also provide analytical tools to understand the functioning of natural systems based on real-time networks of sensors deployed widely across the landscape. Combined with other existing environmental data (maps of geology, topography, soils, etc.) and model outputs (rain and flood forecasts, etc.), the Environmental IoT can provide the basis for decision and support systems for an adaptive management of natural resources and for raising alerts. One of the major factors in the slow uptake of the technology from the lab to the field is because most of these systems rely on nonstandard, custom-designed elements that need specialised expertise [5].

This paper reports in detail on experiences from the design and live deployment of an Environmental IoT targeting specific local and regional environmental applications (hillslope to river catchment) using inexpensive off-the-shelf technologies and deploying the system for a particular environmental issue: flood and pollution monitoring and alerts in rural environments. Our central hypothesis is that our combination of IoT technology coupled with Cloud Computing enables a paradigm shift in our understanding and management of the natural environment, especially related to understanding ecosystem interdependencies, in times of unprecedented environmental change.

The key contribution of this paper is an evaluation of this hypothesis in terms of the technological, scientific, and methodological aspects:

- Technological: the presentation of a complete end-to-end systems architecture embracing cloud and IoT infrastructure designed to operate effectively in the target operational environment;
- Scientific: the assessment of the potential impact of IoT/cloud technology on environmental and earth sciences;
- Methodological: a reflection on the methodological approaches required to facilitate the desired cross-disciplinary conversation to develop and utilise an Environmental IoT infrastructure.

The paper is structured as follows. Section 2 discusses the target deployment environment for the Environmental IoT infrastructure, highlighting the key challenges being addressed by the infrastructure. Section 3 then provides a more in-depth assessment of the related work, supporting

our view that the application of IoT principles in the natural environment has been limited in scope and vision. Section 4 presents in detail the design and deployment of the Environmental IoT platform, and Section 5 then discusses the experiences of this work. Section 6 provides a more high-level discussion around the technological, scientific, and methodological aspects discussed above. Finally, Section 7 presents concluding remarks.

2. The Target Deployment Environment

We focus on one specific geographic region around Conwy in North Wales (see Figure 1), a typical rural area supporting important industries including agriculture, forestry, tourism, and fishing but facing huge challenges brought about by climate change and conflicting demands on land/water resources. The Conwy river rises in the Snowdonia National Park in the Migneint range which has high rainfall providing water resources and important carbon storage and biodiversity in peatland habitats. The Conwy flows through intensive agricultural land out to an estuary important for the shellfish industry, conservation, and tourism. The lower Conwy was recently in the national news due to intensive flooding of several larger towns. Such extreme events are consistent with climate change predictions. The major shellfisheries and tourist beaches close to the Conwy estuary are continually subject to contamination, to risk of closure, or to loss of blue flag status due to microbial pollution from rural wastewater treatment plants and agricultural runoff. These pollution events are becoming increasingly frequent and are also exacerbated by floods, storms in general, and periods of drought.

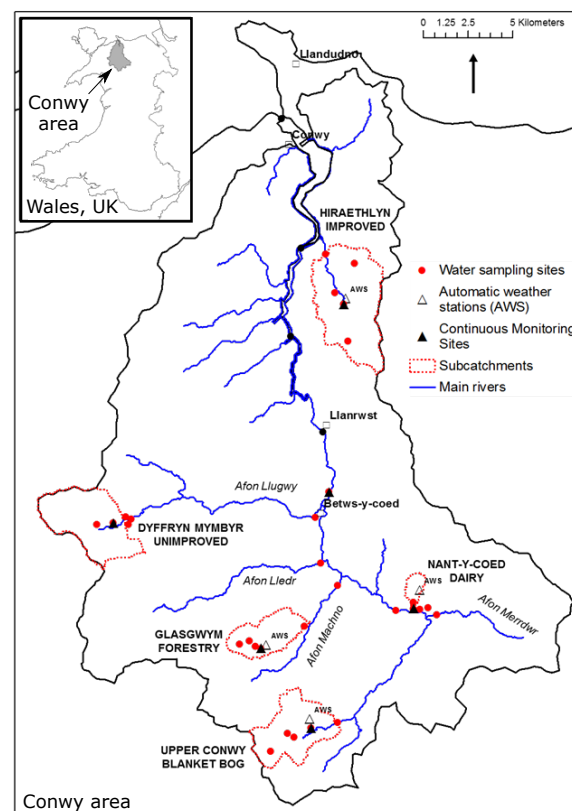


Figure 1. A wide-area view of the geographical region of our deployment. The areas marked in dotted lines are those actively studied by the environmental science team using various types of traditional (non-IoT (Internet of Things)) measurement equipment. The primary site used for our IoT project is that marked as “Hiraethlyn Improved”.

One of the main challenges in the Conwy and many other catchments/landscapes is the potential conflict arising from the needs of different industries, e.g., agriculture, water, tourism, and urban development, and the need for decision support tools to future-proof against climate change. The

GDP associated with Wales and the UK indicate that something is missing in local and national economies to varying degrees, and balancing their requirements is a major challenge for both government and industry alike. In the past, “silo management” has often resulted in a development by one industry negatively impacting on a second (e.g., the intensification of agriculture reducing water quality), and there is a move towards a more holistic consideration of these challenges (cf. Natural Environment White Paper [12]). However, a major problem identified in delivering this vision is the fragmented and inaccessible nature of the resources (i.e., data, models, and tools) needed for a more integrated approach.

This problem is illustrated by the following scenario:

- A *Representatives from the Conwy shellfish industry contacted Prof. Emmett to ask why their shellfish products had just been rejected by Europe based on their high microbiological pollutant load and risk to public health. The industry wanted to know where the pollution was coming from, what caused the pollution to occur in the last month, and what safeguards they could put in place for the future. The professor used the Environmental IoT to identify any recent anomalous events in the Conwy catchment, which she suspected was the source of the problem. She determined that the recent movement of livestock into lowland areas combined with high rainfall and spring tides had caused a significant transfer of nutrients and faecal bacteria into coastal waters. She advised that the use of a mobile APP would allow them to predict when these events might happen in the future and when their waters would be safe to harvest shellfish again.*

An Environmental IoT provides a major opportunity to fill this gap, bringing together data and assets from across different domains (soil, water, plant ecology, and animals) and organisations (e.g., public, regulatory, and industry) to enable integrated problem solving for the agriculture and water industries but with a clear potential to link to other critical industries in the future.

3. Related Work

As mentioned in the introduction, the vast majority of Internet of Things research to date has been focused on smart city and home environments [13]. The instrumentation of the natural environment using IoT technology has received less attention. The key practical differences here are the lack of wide-area connectivity, lack of mains power sources, physical locations that are hard to access, and the often harsh environmental conditions. We survey the most closely related work in this area below. Our deployment is the first to deliver a holistic IoT solution for the natural environment, enabling a real-time study of multiple dimensions of that environment simultaneously.

One of the first experimental testbeds of a multi-tiered architecture for habitat monitoring was proposed by Cerpa et al. [14] in 2001. Using the mote technology developed by Kris Pister at the University of California, Berkeley [15], the authors introduced the principles of building long-lived wireless sensor networks, adaptive self-configuration, and its application to habitat monitoring. Their goal was to integrate these building blocks into a deployable system.

Researchers at UC Berkeley in collaboration with biologists at the College of the Atlantic deployed the first-documented tiered wireless sensor network architecture for habitat monitoring, such as detecting the presence of birds on Great Duck Island (GDI), Maine [16]. They deployed two networks for a total number of about 150 nodes in both single-hop and multi-hop configurations. The initial deployment comprised of 32 nodes using UC Berkeley motes, including temperature, photoresistor, barometric pressure, humidity, and infrared sensors. Two years later, the researchers analysed a rich set of data captured from these sensors from a systems perspective regarding the lifetime of nodes and network performance, reliability, routing, and yield [17].

Researchers from Princeton University deployed a 30-node wireless sensor network with mobile base stations called ZebraNet at the Mpala Research Center in central Kenya which aimed at monitoring and tracking wildlife, particularly zebras, for biology researchers [18]. The ZebraNet system included nodes (tracking collars carried by animals) operating as a peer-to-peer network for transferring data. The collar, working as a wireless computing device, was comprised of a GPS, flash memory, a wireless

transceiver, and a CPU. ZebraNet was the first project to design protocols for sensor networks having a mobile base station.

The University of California, Los Angeles's (UCLA) Center for Embedded Network Sensing (CENS) deployed the Extensible Sensing System (ESS) for habitat sensing at the James Reserve in California. CENS was a multidisciplinary collaborative project to develop and deploy wireless sensing systems to monitor the biological, ecological, seismological, and different environmental phenomena.

Researchers at the University of Southampton deployed an environmental sensor network called GlacsWeb [19] to monitor and study the dynamics of glaciers in hostile environments. The system comprised of 9 nodes (probes) inserted in the glacier, a Base Station on the surface of the glacier having a differential GPS, and a Reference Station working as a gateway for transferring data. Each node was equipped with pressure, temperature, and orientation (tilt) sensors for monitoring the melting behaviour of glaciers. GlacsWeb developed a low-power sensor network supporting the transfer of sensor readings to a web accessible database.

Cardell-Oliver et al. [20] designed and implemented a reactive, event-driven wireless sensor network for the environmental monitoring of soil moisture. A novelty of their implementation is how their network reacts to extreme events such as intensive rainfall to send frequent readings and to collect the most interesting data, reverting to a less frequent collection in periods of lower rainfall. They used Mica2 433 MHz motes with MDA300 sensor boards comprising of soil moisture and rainfall sampling nodes, a base node connected to a Superlite GSM gateway and routing nodes. They evaluated their network by taking into account three performance parameters, i.e., responsiveness, robustness, and longevity of the network.

Researchers at UC Berkeley deployed a wireless sensor network, termed "Macroscopic in the Redwoods", to monitor the complex spatiotemporal variations of the microclimate of a redwood tree [21]. The sensor node was a Mica2Dot, a repackaging of the Mica2 manufactured by Crossbow. The mote consisted of an Atmel ATmega 128 microcontroller operating at 4 MHz, a 433-MHz radio, and 512 KB of flash memory. The measured biological parameters were temperature, relative humidity, and light levels. Although Macroscopic was deployed at a small scale, it had the potential to measure the microclimate phenomena of a redwood tree at a spatiotemporal large scale. Further, it provided useful knowledge regarding deployment methodology and data analysis.

LOFAR-agro (Low Frequency Array) is the first deployment of a large-scale sensor network in precision agriculture to protect potato crops from *Phytophthora*, a fungal disease, with close monitoring of the microclimate [22]. In this project, around 100 nodes were deployed in a potato field to measure the temperature and relative humidity of the environment, known to be the main factors in causing the disease. The researchers experienced a range of unanticipated technical challenges due to poor design, improper application of software engineering principles, semantic mismatches, programming bugs, and insufficient time for testing.

Researchers at Harvard University in collaboration with researchers at the University of North Carolina and the University of New Hampshire deployed a sensor network on an active volcano in Volcan Reventador, northern Ecuador to measure seismic and infrasonic signals [10]. The network comprised of 16 stations with seismic and acoustic sensors. Each node consisted of a Moteiv TMote Sky, a descendant of the UC Berkeley's Mica mote, an 8-dBi 2.4GHz external omnidirectional antenna, a seismometer, a microphone, and a custom hardware interface board. The network was deployed for more than three weeks and collected data from several events. The model was based on event triggering which detected 230 eruptions and captured about 107 Mbytes of data.

In a collaborative project called "Life Under Your Feet", researchers at John Hopkins University and with Microsoft deployed an experimental wireless sensor network in a Baltimore urban forest to monitor soil ecology [23]. They used MicaZ from Crossbow Inc. with a MTS101 data acquisition board comprising of ambient light and temperature sensors. They attached Watermark soil moisture sensors and a soil thermistor to the board. They reported a mixture of previously known and new observations from the operation of their testbed, such as reprogramming the motes, sensor nodes and calibration

cost, the importance and challenge of low-level programming, the need for fast and efficient routing infrastructure and network design, and reliable delivery and high quality of data.

Bishop-Hurley et al. [24] designed and tested a prototype of a neck-collar-based virtual fencing system to monitor and control livestock movement and behaviour. They conducted an experiment on 25 steers at the Belmont Research Station, Australia for over a month. The electronics used in the experiment consisted of the Fleck WSN (Wireless Sensor Network) device in combination with a control board which could produce an electrical pulse. Two 6-volt batteries placed in the collar's compartments, were used to charge and power the Fleck. For radio communication, they used a 433 MHz wave flex whip antenna. A statistical analysis was done on observational and videotaped data to see how animals reacted to the stimuli. The experiment showed that the sensory cue has the potential to be used for monitoring livestock in large grazing systems.

Sensor networks have also been used to detect and monitor river flooding [25], with the capability of integrating remote fixed-network grids for computationally intensive flood-modelling and also performing on-site grid computation. Also in the natural environment, the GreenOrbs [26] project is an extremely large deployment of sensor nodes in a forest environment to help understand canopy coverage over time.

SensorScope was an outdoor wireless sensor network deployed on top of a rock glacier in Switzerland for environmental monitoring in harsh weather conditions [27]. The researchers deployed 16 sensing stations on a 500×500 m area for a couple of months to measure air temperature and humidity, surface temperature, incoming solar radiation, wind speed and direction, and precipitation. They opted for a Shockfish TinyNode sensor mote platform composed of a Texas Instruments MSP430 microcontroller running at 8MHz and a Semtech XE1205 radio transceiver. They designed and implemented the communication stack from scratch for their stations with TinyOS. SensorScope mainly aimed at providing a low-cost, efficient, and reliable WSN with an emphasis on improved data collection techniques and, to a great extent, deployed successfully in outdoor harsh environments.

Corke et al. [28] reported experiences learned from their wireless sensor networks deployed over six years to understand the natural and agricultural environments and the major challenges they pose in Australia. They developed and deployed sensor networks for different applications such as cattle monitoring, ground water and lake water quality monitoring, virtual fencing, and rainforest monitoring, highlighting the technological challenges they faced, which technology they used for each application, and what lessons they learned from each deployment.

All of the above work suggests a strong need for Environmental Internet of Things developments, with high interest from multiple different domains of science. However, work to date remains relatively tightly focused on single dimensions of the environment, lacking a broader view that can fuse data across multiple scientific domains to build a holistic environmental perspective. Further to this, we note that a large proportion of the above works are technology-oriented and often fail to deeply involve scientists in the domain of study throughout the design process. In contrast to these two points, our work takes a multidimensional perspective on an environmental IoT deployment and has deeply involved soil scientists, hydrologists, and animal behaviour scientists from the first stages of design through to actual deployment and maintenance.

4. Design and Deployment

4.1. Architecture and Deployment Overview

Our environmental Internet of Things architecture spans elements in the field and in the cloud, providing an end-to-end solution that is exposed as an Internet service. In designing the system, we had three goals: (i) to use off-the-shelf hardware and software as much as possible; (ii) to choose solutions that minimised the necessary amount of software development and system maintenance; and (iii) to provide a system that could be reused by the environmental science team without expert help.

Our resulting architecture is shown in Figure 2. In the field, we use a micro-server as an Internet gateway (a Raspberry Pi) which is situated at a location with plentiful energy, either from a mains supply or from a large battery. This device is fitted with an Internet link technology, in our case 3G, as well as a long-range radio to communicate with the rest of the IoT deployment in the field. Our remaining IoT devices are then deployed into the actual field of study, which may be several kilometres away from the Raspberry Pi. One of these IoT devices (I_A) is itself equipped with a long-range radio and communicates with the Raspberry Pi; the remaining IoT devices have short-range radios and communicate only with I_A . This design has two major benefits: (i) the use of a separate Internet link node, which can be at a distant location, provides high flexibility because a wide area coverage (i.e., 3G) may be sporadic or nonexistent in the actual area of study, and (ii) the use of just one long-range radio “pair” reduces the cost of the overall deployment, as this higher class of radio is more expensive to buy and also draws more power, needing more expensive energy harvesting and storage technology at I_A . On the cloud side, we then have one main server which receives data from the deployment, storing that data into a database. Other cloud servers (or end users) use web services to extract data from this machine and feed it into models or visualisation systems.

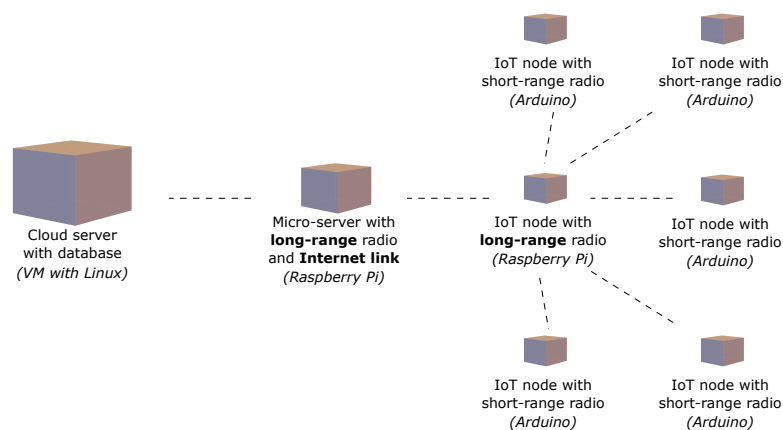


Figure 2. The architecture of our IoT system: Further Internet services connect to the cloud server, providing scientific modelling or lightweight applications to visualise data.

In our particular case study, our system architecture is projected onto the deployment area shown in Figure 3, providing both a geographical view and a detailed hydrology view. The Raspberry Pi is located in a small utility shed offered by a local farmer; our field of study is then 1.5 km away from this location. The terrain between the Pi and the field of study also involves significant elevation changes and dense tree lines. Besides the long-range IoT node, which communicates with the Pi, we then have two types of node deployed in the field: fixed nodes, measuring soil moisture levels, air temperature and humidity, and ground flow, and mobile nodes attached to livestock, measuring animal location and behaviour via GPS and accelerometers. Together, these nodes provide an unprecedented real-time insight into multiple facets of the natural environment and their interrelationships.

On the cloud side, data has been stored on MongoDB and is available for analysis by the stakeholders involved. The fast and easy availability of the data through the cloud means that the scientists can draw early insight from the data and can make appropriate decisions. The other benefit of this is that it can stimulate conversations and possibly collaborations among environmental scientists. The data has also been used in a follow-on project around the semantic enrichment of data in order to tag the provenance of the data and also to enable spatiotemporal queries.

In the remainder of this section, we discuss the details of the nodes used in our deployment, the cloud infrastructure, and the user-facing elements of our architecture.




Figure 3. The deployment location of our IoT system, showing a geographical view of the entire deployment as well as a detailed technical view of the field of study in hydrological terms: Mobile (livestock) nodes are not shown on the geographical view, but tracked animals are located within the treeline enclosing the area marked “field of study” (courtesy of Satellite image from Google Earth).

4.2. Node Technology

As mentioned above, we use two types of sensing nodes—soil sensing and livestock tracker nodes. The soil sensing nodes are used to take measurements related to the soil, whilst the livestock trackers are used to monitor the location of farm animals (specifically sheep during this project) in the area concerned. Both types of nodes have different physical requirements but need to be able to communicate with one another and so have a common link technology. In addition to this, we use a long-range link node and a Raspberry Pi.

In this section, we present and justify our design choices in building these nodes, comparing with alternative existing technologies where available. Our overriding motivation was to develop a platform that was cheap and flexible. Images of each node type are shown in Figure 4.




Figure 4. IoT nodes (from left to right: relay, soil node, and livestock node on a sheep).

4.2.1. Fixed Nodes (Soil)

For both the fixed and mobile nodes, we used Arduino technology as the basis of our solution. We made this choice after examining a wide range of IoT platforms, including commercial off-the-shelf solutions. Our choice here was motivated by three major concerns: (i) The environmental scientists wanted a large range of sensing options, including custom-built sensory inputs; (ii) we needed a common technology that could support both fixed and mobile nodes working in the same network; and (iii) we were limited in our equipment budget. Looking at other technologies, IoT solutions tend either to be too specific in their application domains, lacking the crucial configurability that we needed; to lack support for both fixed and mobile technology in a common network; or to come in at a higher price point per node than we could afford. We, therefore, settled on Arduino-based nodes

across our deployment. This was itself not without problems, however. We needed to develop more software ourselves and also had to deal with the relatively high power consumption of standard Arduino boards.

Our soil monitoring nodes cost £150 each, including all sensing inputs, casings, and power supplies. We chose to use the Arduino Mega2560 due to its relatively high processing power and larger collection of ready-built libraries. We used the Grove input extension board, allowing a high number of sensing inputs to be connected, with an Xbee Pro module for radio communication with other nodes. Each node also had two 7800 mAh batteries connected to it, providing sufficient power for almost a week. Each such battery had a built-in solar panel for trickle charging.

There are two types of soil nodes deployed in this experiment—standard and super soil nodes (stations). The standard soil node consists of the cheap Grove sensors, as shown in Figure 5, whereas the super soil node consists of the Acclima TDT Soil Moisture sensor, as shown in Figure 6 together with the Grove sensors.




Figure 5. Standard soil node.




Figure 6. Super soil node.

In the remainder of this section, we discuss the issues of sensory inputs, power supplies, and weather proofing for our soil nodes in particular.

Sensor Inputs

Our sensor input strategy, guided by the environmental scientists on our team, was to deploy a large number of very cheap sensor inputs at most of our nodes, combined with a small number of high-resolution (but very expensive) sensors of known scientific quality. For our cheap sensors, we used Grove hardware as listed in the upper part of Figure 7. All of these are off-the-shelf units except for the physical design of the surface flow detector, which was custom-built by the environmental science team.

| Sensor | Quantity | Readings | Type of Soil Node |
|----------------------------------|----------|--------------------------------|--------------------|
| Soil Moisture Sensors | 3 | Soil Moisture | Standard Soil Node |
| Soil Temperature Sensors | 3 | Soil Temperature | |
| Air Temperature/Humidity Sensors | 3 | Air Temperature/Air Humidity | |
| Surface Flow Water Sensor | 1 | Water Detection at the surface | |
| Acclima Sensor | 1 | Volumetric Water Content | Super Soil Node |
| | | Soil Temperature | |
| | | Bulk Relative Permittivity | |
| | | Bulk Electric Conductivity | |

Figure 7. Sensor inputs on our soil nodes, mixing a large quantity of cheap Grove sensors with a small number of high-quality Acclima sensors.

As a high-quality counterpart to these sensors, we used an Acclima soil moisture sensor (Digital TDT SDI-12 (Landscape Technologies, Sydney, Australia)), as shown in Figure 6. This sensor is considered one of the most accurate and stable soil moisture sensors available today. The Acclima sensor is more expensive than the Grove sensors; however, it is a preferred tool used by environmental scientists in order to measure different parameters in the soil. In typical (non-IoT) deployments, an Acclima is attached to a data logger unit and scientists periodically visit the site to download the data from this logger by physically attaching a USB drive. Instead, we connected the Acclima directly to our Arduino board and wrote a simple driver to periodically take readings from it and send them through our IoT to the cloud.

In total, each soil node has 3 air temperature/air humidity sensors, 3 Grove soil moisture sensors, 3 soil temperature sensors, 1 surface flow sensor, and (optionally) 1 Acclima soil moisture sensor. The reason behind replicating most of the Grove sensors was to get a more accurate and reliable value of the sensor readings in case one of them fails. Another reason was because the environmental scientists wanted to obtain a range of values at different spatial points for comparison (for example, the air temperature sensors were mounted at three different heights). Moreover, although Acclima is a preferred sensor to Grove by the environmental scientists, the scientists wanted to experiment with the latter as a cheap alternative, as the Acclima sensor is comparatively more expensive. Hence, the experiment was carried out using the configuration of Grove sensors in order to see if such sensors do provide equally reliable readings to the Acclima sensor. The readings that have been compared between Grove and Acclima are mainly the soil temperature and soil moisture readings. As for the remaining Grove readings such as surface flow, air temperature, and air humidity, these have been included in the experiment to get additional readings related to the soil and air. This has also been possible due to the fact that the Grove manufacturers provide a vast array of sensors which can easily be used and interfaced with the Arduino board.

A major question of the environmental scientists was whether the cost/accuracy tradeoff was a good one, i.e., despite usually wanting to have the highest quality measurement equipment available, in this case, they were willing to compromise on quality for quantity of real-time data points, as long as we still had some high-quality inputs in the field as a known reference point.

Power Supply

Powering the Arduino Mega2560 board, with all of its attached sensors, proved to be a major challenge. We found that the most reliable way to power the system was via its barrel connector, rather than the USB line. However, with a recommended voltage range of 7–12 V via this connector, we required very high power batteries to achieve a sensible lifetime.

To help understand this lifetime issue, we measured the current draw of a complete sensor node, which was found to be 130 mA when active and 45 mA when in low-power sleep mode. Both values are extremely high compared to more bespoke sensor node solutions that use chips such as the TI MSP430. Additionally, the power draw is actually dominated by the Arduino board itself, which takes 90 mA in active mode and 30 mA in sleep mode. Using aggressive power cycling, if we assume that a soil node sleeps for 5 min and wakes for 5 s to send data, then we can say that the node is awake for 1% of the time and asleep for the rest. With this in mind, we can calculate an average current draw of 46 mA with 5 min of sleep, and with a 7800 mAh battery we achieve $7800/46 \times 0.7 = 118$ h or just under 5 days.

To meet the voltage input requirements and to extend the lifetime, each node was equipped with three 7800 mAh lithium ion batteries, offering 10 days of operations. For the environmental science team, this operating lifetime was around the same amount of time for which they would have regularly visited the field of study to take data from data loggers, and so having this as a worst-case battery replacement interval was acceptable. Operationally, the environmental science team kept a second set of batteries charged at their lab so that they could quickly replace any drained batteries in the field.

Weather Proofing and Protection

Each sensor node was housed in a 1.9-L plastic box enclosure, with space for batteries inside. The boxes were internally packed out with polystyrene to prevent movement of the components inside, and dessicant sachets were included in each node to counter the build-up of condensation that could damage the electronics. For each external sensor, holes were drilled in the enclosure through which to pass the cables; these holes were then sealed with adhesive. Each node included a 4" antenna (RN-SMA4-RP, Microchip Technology, Chandler, AZ, USA) mounted externally using the RPSMA connector's (MC000800, Multicomp by Newark, Chicago, IL, USA) nut and a sealing ring.

The individual Grove sensors, being generally designed for use in protected environments, were modified for field deployment. All were fitted with 50-cm cables to the Arduino. Grove temperature sensors (Seeed Studio, Shenzhen, China, 101020015) were fully enclosed in adhesive-lined heat-shrinkable tubing to include the cable termination. Hot melt glue was injected into the top of the tube prior to shrinking around the cable, creating a rigid section and protecting the connection. Three sensors per node were installed as soil-temperature replicates at a 5-cm depth. Grove moisture sensors (Seeed Studio, 101020008) and Grove water sensors (Seeed Studio 101020018) were similarly enclosed but left the probes bare. Three soil moisture sensors per node were installed in the soil surface. Each node also had one Grove water sensor, which was used to prepare a prototype overland-flow detector. A plastic box was used to exclude rain from the sensor, with a wide cutout on the base of the upslope side and a drain on the other. Inside the box, a V-shaped funnel was constructed in plastic to direct any overland flow over the sensor and out of the drain. Water flowing through the bottom of the V-shape would pass over the water sensor, completing the electrical connection to send a signal to the Arduino indicating over-land water flow. The Grove temperature and humidity sensors (Seeed Studio, 101020011) were shielded from rain under an upturned plastic cup. For each node, two sensors were installed at 35 cm above the soil surface and one sensor was installed at 5 cm above the soil surface.

The entire soil node assembly was placed inside a steel mesh dome to protect it from curious animals—these domes were already at the field deployment site and had been well-tested by the environmental science teams in their pre-IoT data logger deployments.

4.2.2. Mobile Nodes (Livestock)

Our livestock nodes were designed to inform the environmental scientists of the location of animals and to correlate this with rainfall events that could lead to over-land flow causing contamination of the nearby river. We again based the design of this node on Arduino technology, this time using an Arduino Pro Mini as the main board, connected to a GPS module, a 3-axis accelerometer, and an Xbee radio module. Our choice of Arduino in this case was entirely down to cost: Commercial off-the-shelf solutions for animal tracking cost over £1000 per unit, whereas our solution costs just £100. However, due to the need to minimise the scale of our unit to make it sufficiently lightweight to attach to an animal, we needed to hand-craft the design of this node much more closely, involving a detailed electronic design of the board connecting the individual components together and also 3-D printing of the enclosure.

Sensor Inputs

As animal tracking is, in general, an area of scientific study that requires less fine detail than soil science, we had fewer technological constraints here coming from the requirements of the animal behaviour scientists. However, this was still a highly challenging design due to the interaction of technology with animals. We considered a large range of options for sensing animal location, from RFID tags on the animals with readers positioned around the field of study to GPS units that would require more power and so a heavier overall node for animals to carry, and the use of web cameras sited at high points around the field of study (e.g., in trees) combined with image processing in the cloud as an approach that would completely avoid attaching equipment to animals. We settled on the GPS option as it would provide the most flexibility in what we could potentially do with our data streams, including the easy identification of individual animal behaviour. We chose to augment this with accelerometers as a way to potentially understand behaviour in a more detailed way.

Power Supply

Our animal tracking nodes were powered with a 10,000 mAh battery, designed to be replaced once per week. Similar to the soil nodes, we used aggressive power cycling with nodes in sleep mode most of the time. However, the amount of time that the animal tracker nodes needed to be awake was significantly extended by the duration of time that the GPS unit took to start up and then to acquire a high-quality GPS fix.

Weather Proofing

The casing design for the animal tracking nodes was complicated by the fact that they needed to be both weather-proof and physically suitable to the animals without causing irritation. To provide a very close fitting enclosure for our electronics, we 3-D printed the casing and then designed a harness to secure the device to livestock (in this case, sheep). After several trials, we settled on a harness made of soft but durable material which secured the sensor node on the top of the animal's back. Figure 4 shows the soil node and the animal tracker in their respective enclosures.

4.2.3. Link Nodes

Link nodes are used to collect data received from our sensor nodes and to deliver them to the cloud. In our IoT system, we have two kinds of link node: a long-range relay and a gateway node.

Relay Node

The relay node is a Raspberry Pi model A+, sited in the field of study. We note that our initial network design used an Arduino as the relay; however, it quickly became clear that the Arduino's relatively low-power CPU was unable to cope with forwarding the volume of data from the rest of the IoT nodes in the field. The relay node serves to receive data from the sensor nodes in the field and to send them to the gateway node via a long-range local area radio transmitter. This design gives us much higher flexibility in the physical locations of the field of study and the (Internet link) gateway node. The relay node receives data from both the soil nodes and the livestock tracking nodes, being in the local radio range of both kinds of node. The location of the relay was carefully chosen to maximise its radio signal propagation profile both in terms of the other IoT nodes in the field of study and the gateway node half a kilometre away.

The relay node is equipped with two different radios: a short-range Xbee and a long-range Xbee, each with their own antenna. The first Xbee (Xbee1) is used to receive data from soil/livestock nodes. The second Xbee (Xbee2) is then used to send data from the relay to the gateway. As we were aware that various weather and other local interference effects could disrupt the communication between the relay and the gateway, we implemented a simple store-and-forward system at the relay. Each data packet received by the gateway node from the relay is acknowledged, and any unacknowledged packets are stored locally at the relay node (using a TinyDB database) until their reception is confirmed by the gateway. Figure 8 shows how the relay works. The radio xbee1 (step 1) receives data packets from the sensor network, and these are stored in a queue (step 2) at the relay node. They are then processed in a FIFO (First In First Out) method (step 3), where each data packet is forwarded by the radio xbee2 (steps 5 and 6). If the relay does not receive any acknowledgement from the gateway about a received data packet, then this data gets stored at the TinyDB database (step 4).




Figure 8. The relay node.

The relay node (Raspberry Pi A+) was powered via its USB socket. To facilitate this, a DC/DC regulator of 12 V/5 V (CPT, CPT-UL-1) was installed inside the enclosure (see Figure 4 for an image of the relay node in its enclosure). The output was attached to a USB charging cable and the input to the buccanier connector. This enabled a 12-V, 200-Ah deep cycle battery to be connected externally. This could be maintained with a solar panel, but this was unnecessary due to the high capacity and the availability of these batteries in our stores.

Gateway Node

The gateway node is a Raspberry Pi model B and works in a similar way to the relay node, using a store-and-forward approach to push data into the cloud. Specifically, whenever a data packet is received from the relay node, the gateway node stores that data in a local TinyDB database. The gateway then periodically opens a connection to our cloud server (via a 3G cellular connection) and sends all data from its TinyDB database that has not yet been acknowledged as received by the cloud server.

We use TinyDB databases on the relay and gateway nodes as they are extremely lightweight, occupying minimal storage space on the relatively limited Raspberry Pi secondary storage. The location for the gateway provided 3G signal coverage and access to the mains power. Therefore, the gateway was powered via USB using a mains adapter plug for continuous power.

4.3. Cloud Infrastructure

The cloud side of our IoT architecture is relatively simple, using just one server as a database to store all data received from the deployment—data which arrives as periodic packages from the gateway node over its 3G link. Other cloud resources can then access this server to use the data for specific purposes, such as scientific modelling or visualisations. We chose to use MongoDB as our data storage technology due to its schema-free design that is highly flexible to different kinds of data—as is highly likely to be the case when using many different sensor modalities in the way that we have. Modelling and visualisation services are run on elastic cloud computing resources, scaling up the amount of processing power needed as their complexity increases.

Besides receiving and storing data, the main purpose of our cloud infrastructure is to act as a management service for the deployment. This allows nodes to be registered as they get deployed (including their position, sensing attributes, and a logical group membership for each node) or for their details to be edited if they are later moved or changed and also for the sampling rate of deployed nodes to be changed if desired. The latter is done using logical node groups, such that the sampling rate affects a selected group of nodes (e.g., “soil nodes” or “riverside soil nodes”). The management service also records the last time it heard any data from each node and the last-reported battery level of each node to aid in the day-to-day running of the deployment. As part of this service, we have developed a mobile app through which environmental scientists can add and update details about deployed nodes as they make changes in the field of study (see Figure 9).

Figure 9. Node registration.

While this architecture is sufficient for our current deployment, in the future, we would explore more efficient solutions for processing and data distribution, particularly for modelling and visualisation services—such as allowing some data processing to occur on the database server itself to avoid shipping large quantities of data to other locations.

4.4. User Interface and Data Presentation

An important part of providing a truly end-to-end IoT solution is the user interface to the system. In this respect, we have developed a significant “dashboard” for users to explore. We have also begun a

semantic annotation initiative to connect our data to the semantic web for wider usability. We describe both of these elements below.

4.4.1. Dashboard Development

Our dashboard application is an important part of “in the wild” research, gaining insights into what it means to deploy in the natural environment with all the complexities that this brings (different geographies, exposure to the elements, and extreme events). The dashboard application makes use of the data streams collected through the IoT infrastructure that have been stored on our MongoDB database. The purpose of this dashboard is to enable environmental scientists to quickly make sense of the data collected, informing them of the status of the sensed environment and of potential related hazards detected through the system and providing room for decision making.

The data in the MongoDB comprises of readings from the soil sensor nodes and the sheep trackers, all deployed in the Hiraethlyn field, as mentioned earlier. Whilst the soil sensor readings contribute towards environmental parameters such as soil temperature, soil moisture, air humidity/air temperature, and surface flow, the sheep trackers provide the GPS locations of sheep in the said area. A user-driven application has been developed in order to present the readings in a dashboard-style, as shown in Figure 10.




Figure 10. Environmental IoT dashboard application.

The dashboard consists of different windows, namely an operations dashboard, a soil dashboard, and a sheep dashboard, as indicated in the left pane of Figure 10. The operations dashboard is the main window and shows the number of sites where the IoT infrastructure has been deployed, the number of

soil nodes and sheep trackers used in each site, the status of the nodes (whether they are working), and the kinds of readings being received from each type of node. The application also enables the user to select a node and to analyse the readings obtained from this particular node, as shown in Figure 11. This figure shows the readings from the air temperature/humidity sensors from one soil node over a one-week period. Each of these windows is populated using pluggable “widgets” which can be individually developed to add to the dashboard. This is another key element of configurability to support a highly flexible set of current and future sensor modalities. An example of an open-source IoT dashboard can be found here: <https://thingsboard.io/>.

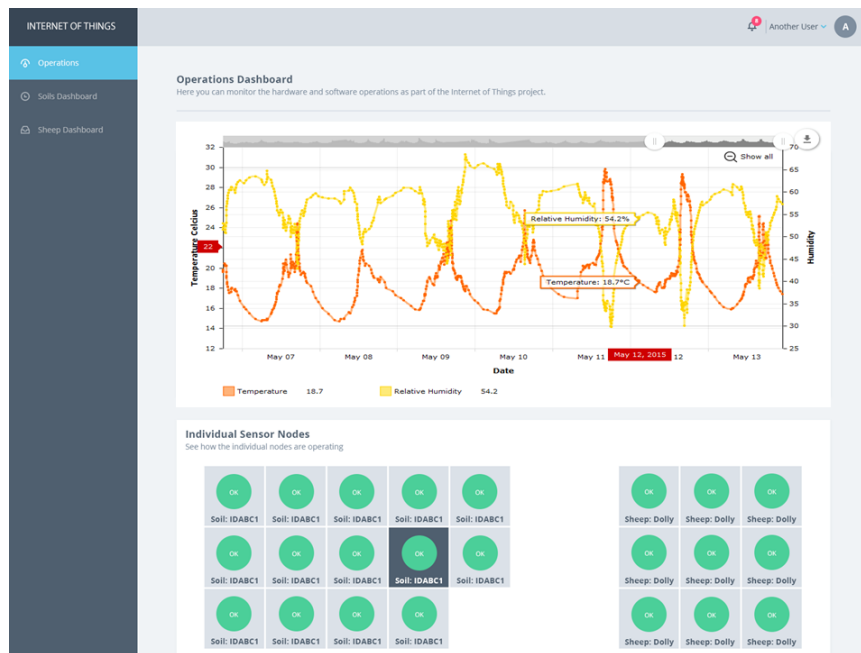


Figure 11. Readings for one soil node.

4.4.2. Semantic Data Enrichment

Whilst the dashboard application development focused on presenting the “raw data” collected through our IoT network in a significant way, we are also venturing further with this development by exploring how to semantically enhance our system. The MongoDB data can be used to plot the sensor readings obtained over a time period on the dashboard, thus providing a medium to the environmental scientists to carry out in-depth analyses over these sets of data. However, the data stored on MongoDB does not provide any other additional information to the scientists, other than the sensor readings, that may improve the analysis of the sensor readings. This entails information such as the types of sensors used to collect the readings, the measurement capabilities of such sensors, the related conditions under which the readings have been taken, etc.

The enrichment of the sensing data collected with such metadata can be made possible through the use of semantic technologies, namely ontologies. In general terms, an ontology is defined as “a formal, explicit specification of a shared conceptualisation”, as proposed by Tom Gruber (<http://semanticweb.org/wiki/Ontology.html>). It is a data model that enables us to define a context or domain through a set of formal, interrelated vocabulary terms.

An ontology design requires a sound knowledge of the domain to be modelled, which is normally obtained through domain experts. In this case, semi-structured interviews were conducted with the environmental scientists, who provided their domain knowledge and also narrated their experiences with dealing with scientific data. This amounted to heterogeneity arising in data formats and a lack of provenance information about the measurement readings. This semantic enrichment task aims at tackling these shortcomings in terms of providing adequate metadata information about the data

and a platform enabling the integration of heterogeneous data through a semantic approach. The ontology created also captures the features of the IoT network, partly based on extending several existing vocabularies such as SSN/DUL (https://www.w3.org/2005/Incubator/ssn/wiki/DUL_ssn) to describe the kinds of sensors used and their measurement capabilities, the Time ontology (<https://www.w3.org/TR/owl-time/>) to describe the timestamp of the values read, MUO (http://www.semantic-web-journal.net/sites/default/files/swj177_7.pdf) to describe the units of the values read by the sensors, and Geo/Geosparql (<http://www.opengeospatial.org/standards/geosparql>) to describe the location of the deployment site. We have written a separate full paper on the semantic enrichment work, and it will shortly be published.

The ontology also enables us to reason more deeply over the data model and to infer additional knowledge that pertains to the domain of interest. At a high level, we therefore hope to be able to answer complex and multidimensional queries of the following form: Can we correlate the soil moisture readings to the type of weather/or rainfall that occurred during that time period? Can we infer that the nearby river was polluted due to sheep wandering in the Hiraethlyn region when there was heavy rainfall? Did the rainfall trigger a run-off of water (with pollutants/FIO) due to the soil having reached a high level of saturation?

The exact way in which these queries can be expressed, particularly in a form that suits domain scientists, remains an open research topic. The addition and use of semantic annotations to data is also an important part of provenance, particularly in the hard science use case that we have here; we return to the issue of provenance more broadly in Section 6.

5. Deployment Experiences of an Environmental IoT

In this section, we report on some of the main deployment and operations stage of the project. Our deployment site is situated in a field in a region called Hiraethlyn, which is found in Conwy, North Wales. Figure 12 shows the number of the nodes in the field, and Figure 3 shows the layout of the soil nodes according to the topology of the field (swale, midslope, and foot of slope). Figure 13 shows the deployment area in Hiraethlyn, which is in a strategic location running alongside a tributary of the Conwy river. The river runs through the Snowdonia National Park in the Migneint range and drifts through intensive agricultural land out to an estuary important for the shellfish industry. A collection of images from this deployment experience is shown in Figure 14. One of the objectives of the IoT experiment is to enable the stakeholders to make informed decisions regarding any possibility of flooding of the Conwy river. The fact that this field slants towards the river tributary is one of the key contributors towards achieving this objective. As shown in Figure 13, the soil nodes are scattered around the Hiraethlyn field and also around the river bank. These soil nodes help to detect the soil moisture content of this field. This value coupled with the met office weather data can help to inform the environmental scientists about any possibility of flooding of the river.

Another objective is to determine whether there is a danger of pollution or contamination to the river water due to the run-off of pathogens from the soil to the river. Given that the sheep roaming in this field contribute to the accumulation of pathogens that wash down to the river, it was decided to fit few sheep with a collar reporting their every movement through the IoT infrastructure. Therefore, the deployment procedure entails setting up this IoT infrastructure, consisting of the soil nodes in the field, the tracker collars placed round the neck of sheep roaming in this area, as well as installing the relay and gateway nodes. The following section gives more insight into this deployment process.

| Type of Node | Quantity | |
|--|----------------------------------|-----------------------------|
| No. of soil nodes | 15 | 12 Standard + 3 Super Nodes |
| Layout of the soil nodes in Hiraethlyn field | Swale | 3 |
| | Midslope | 7 |
| | Foot of slope | 3 |
| | Stream side | 5 |
| No. of sheep trackers | 5 | |
| No. of relay nodes | 1 on swale (hill top) | |
| No. of gateway nodes | 1 (1 mile from Hiraethlyn field) | |

Figure 12. Number of nodes in the field.




Figure 13. Hiraethlyn field deployment site.




Figure 14. Images from our deployment site. Clockwise from top left: initial site reconnaissance with environmental scientists; the site as we initially found it; a protective dome over a soil sensing node; configuring the relay node; constructing hardware at the site; and the location of the Internet gateway node.

We first note that our deployment timing coincided (by chance) with one of the most severe storm and flooding periods in UK history, during November–December 2015, creating extraordinarily challenging conditions for operating technology in the field. The most significant problems caused by this were lower than anticipated solar charging levels and terminal water damage to some nodes. Nevertheless, our deployment was broadly successful, resulting in large volumes of real-time data collection from both the soil nodes and the animal trackers.

5.1. Design of the Soil Node

The basic enclosures were sufficient to exclude rain from the relay and soil nodes. There was some difficulty in excluding heavy rain from the hardware during soil node battery changes, which could be resolved with an external battery connected via a waterproof connector, avoiding the need to open the housing of the IoT node itself. Additionally, soil nodes near to the watercourse, in particular, were damaged by water ingress when submerged by flooding.

5.2. Design of the Animal Tracking Node

This node performed well and was unaffected by rain issues, though some build-up of condensation inside the node housing was observed, likely a particular problem as these nodes were in prolonged contact with warm animals and an unusually cold and wet environment. Reports from the local farmers indicated that the sheep to which the nodes were attached appeared behaviourally unaffected by the deployment.

5.3. Battery Usage (Interval for Changing Batteries)

The runtime of all of the nodes was shorter than anticipated, requiring attendance at the field site at 5-day intervals by the environmental science team to attend to nodes which had ceased reporting. We consider four reasons for this: (1) Additional consumption of power by the attached sensors may have been greater than estimated. (2) The 3.6-V battery packs connected in series, each include an internal circuit protecting against over-charge and discharge. Any discrepancy in charge between the packs comprising a unit would result in one component pack reaching upper/lower protection thresholds ahead of others, interrupting the circuit for all component packs. With this experience, we consider that 12-V, deep-cycle, sealed lead acid batteries would offer a more cost-effective and suitable power source despite a somewhat larger and heavier format. These could be connected externally and maintained by a solar panel depending on node situation. (3) Later in the project, we detected a problem with the sleep/wake cycle of the XBee, which resulted in nodes ceasing to report and which may have been mistaken for a discharged battery. The frequency of this occurrence during field deployment is not clear, but in hindsight, the observation that nodes could often be restarted by cycling the power rather than replacing the battery, and observation that a healthy voltage was often measured in batteries recovered for charging suggests that this error may account for a significant number of early drop-outs. (4) Exceptionally poor weather limited solar recharge effectiveness across the deployment. In addition, we note that we were never able to operate the accelerometers on the animal tracker nodes as the volume of radio traffic generated was too high a drain on their batteries. Independent trials suggested that on-node classification algorithms could have been used to process data locally before sending a simple summary status of the animal such as “standing”, “sleeping”, etc. As a general comment on our platform, we found that the Arduino hardware that we chosen had a surprisingly high power draw even in sleep mode; a modified solution that uses a low-power micro-controller to wake up the main Arduino chip when needed could have significantly helped with this but was beyond our capability to implement in the time constraints that we had.

An early guide from the environmental science team indicated a desire to try using a large number of cheap sensors rather than fewer more expensive sensors. The motivation behind this was to reduce deployment cost, sacrificing data quality for quantity in some cases, while still maintaining selected high-quality sensors in the field as calibration points. For our particular deployment conditions

(particularly the extreme weather), this does not appear to have paid off, with a high frequency of anomalies in sensor data being apparent. For the prototype overland flow detector and the air temperature/humidity sensors, problems with the constructed housings are likely to have contributed to anomalous data. Overland flow detectors are not commercially available due to difficulties in designing a reliable instrument, and our prototype suffered problems with false positives due to internal condensation and blocked drainage. Subsequently the Grove water sensor, incorporated to detect the presence of water in the channel, became corroded due to the relatively cheap copper alloy used in its construction. Research-grade air temperature and humidity sensors are usually enclosed in radiation shields to minimise heating by solar radiation and direct exposure to precipitation. These may be passively or actively aspirated. These shields cost much more than our chosen sensors. We used white plastic cups as shields which excluded precipitation but were not vented and may have trapped heat and condensation. The Grove soil moisture sensors also generated anomalous data due to the same copper corrosion noted above; while we were able to protect the control circuitry, we could not defend against electrolytic corrosion of the moisture probes. Figure 15 shows a corroded Grove soil moisture sensor. We stress that the Grove soil moisture sensors are not intended for this kind of “heavy duty” use with long-term submersion in standing water—had the weather been closer to general trends for this time of year, we would likely have had less problems in corrosion as most locations in the field of study are less moist.



Figure 15. A corroded soil sensor.

5.4. Calibration of the Sensors (Grove vs. Acclima)

As we had collected a significant volume of data from the deployment, at the end of the field deployment period, we devised a lab-based experiment to seek a calibration of the Grove Soil Moisture sensor data against that of an Acclima TDT. We powered 4x nodes at a constant 12 V using a laboratory power supply. We attached Acclima TDT sensors to the nodes and assigned each node to one of four 35-litre planters filled with topsoil. We installed the sensors as per the field deployment and placed the relay and gateway nearby. We then saturated the planters with water before drilling holes and leaving the planters to dry out to compare the drying curves revealed by the two soil moisture sensor types. This phase of our work (which is ongoing) was carried out without direct assistance from the computer science team and is a positive indication that we are able to reuse the field technology in different settings without expert knowledge. In the future, we hope to be able to redeploy the network in other fields of interest.

5.5. Data

Figures 16 and 17 respectively show the type of data gathered from the soil sensor nodes and the animal trackers. The soil nodes sense the environmental parameters such as soil moisture, soil temperature, humidity, air temperature, etc., whilst the animal tracker reports the location of the animal and accelerometer data. On the other hand, Figure 18 shows an example of a semantic application using the IoT data. The application reports nodes that have been deployed in the field, and upon a click of one such node, it opens up a dialog showing not only the reading at this given node (as shown in Figures 16 and 17) but also the provenance information of this node such as the location, the types of sensors found on this particular node, etc.

| Key | Value | Type |
|--|----------------------------|----------|
| √ (14080) [_id : 55f525260dd8e4084d- (15 fields) | | Document |
| acclima_soil_temp | Nil | String |
| timestamp | 2015-09-13 07:26:30.688781 | String |
| soil_temp_1 | -273.15 | String |
| address | A4 | String |
| acclima_buld_elec_conductiv | Nil | String |
| soil_moisture_2 | 620 | String |
| soil_moisture_3 | 579 | String |
| seq_id | 2262 | String |
| acclima_vol_water_content | | String |
| soil_moisture_1 | 786 | String |
| soil_temp_2 | -273.15 | String |
| soil_temp_3 | -273.15 | String |
| sensor | soil2 | String |
| acclima_bulk_rel_permittivity | Nil | String |

Figure 16. Soil data.

| Key | Value | Type |
|--|----------------------------|----------|
| √ (83213) [_id : 55f2466e26a3288e09fc- (10 fields) | | Document |
| _id | 55f2466e26a3288e09fc820 | Objectid |
| rsi | -76 | Int32 |
| timestamp | 2015-09-11 03:11:42.561085 | String |
| address | 0xd0 | String |
| lat | 53.238857 | Double |
| seq_id | 4779 | Int32 |
| sensor | gps | String |
| gps_quality | 1 | Int32 |
| batt_v | 0.65 | Double |
| lon | -4.015807 | Double |

Figure 17. Sheep data.




Figure 18. Map showing semantically enriched IoT data.

5.6. Simple Network Design

Our final network design was effectively a simple star network with one long-range hop to the Internet gateway. This was the result of significant field testing of a range of radio technology at the deployment site. Our initial plan was to make use of a multi-hop data collection protocol, as is common in the literature, but field testing indicated that this would not be required. In fact, we observed extremely good levels of point-to-point connectivity across long distances with treelines and other obstacles between nodes. As an anecdotal example of this, one of our nodes was transported in an off-road vehicle to a distant location for range testing; this node remained in contact with our base station for the entire journey despite being inside an enclosed steel vehicle. As a result, we adopted our current simple network design with one node sitting at the top of the hill in the field of study, acting as a relay for all other nodes in the field. The relay then sends its data to a gateway node which sits in the barn at the site. This is, in reality, a very common architecture for IoT deployments when working in the relative constraints of a real deployment at a remote location, where mains-powered sites are very limited. The logic here is that, instead of relying on redundancy, we rely on design heuristics in which the long-range node is sited at a “safe location” and is built from more reliable components relative to the end nodes. As long as these properties are true, this gives the long-range node significantly more reliability than each of the end-nodes and so failover support is less of a concern. In our case, the characteristics of the long-range node’s location were (1) indoor, entirely sheltered from the environment, and (2) on a reliable mains power, avoiding reliance on batteries or energy harvesting; the node itself was a Raspberry Pi which is a mass-produced VLSI-style device with excellent reliability. These factors compensate for the lack of redundancy at a deployment location containing no other mains-powered locations besides the one that we used. If future deployments were to need a node dispersal over a larger physical radius (more than 2 km apart), a more complex network design may then be required—our most likely solution would be to adopt a design using further relay nodes as cluster-heads in more distant parts of the network, such that only the relays communicated over longer distances.

6. Discussion

This section provides a more high-level discussion of the contributions of the project, revisiting in particular the three key aspects of the project: the technological, scientific, and methodological aspects. We look at each area in turn, reflecting on the key contributions and insights in each area.

6.1. Technological Contributions

This project was by definition an “in the wild” deployment in a complex and remote area of North Wales, and this dictated a lot of the design choices. In addition, the work was in support of real scientific enquiry and hence had to provide a complete and comprehensive solution from a sensor to data portal. This resulted in a pragmatic and end-to-end architecture that achieves seamless integration across both Internet of Things and cloud domains. The pragmatism is most evident in the network architecture (as presented in Section 4). For example, we initially planned to deploy a mesh network architecture for added coverage and resilience. With the range and effectiveness of the XBee Pro radios, this was however deemed unnecessary, and instead, we opted for the simpler point to point architecture, supplemented by the additional bridging nodes with long range radios (effectively better antennae). With careful placement of the bridging node(s), this solution proved successful. We also initially planned to achieve end-to-end integration through IP, as advocated by many researchers in the IoT community [29,30]. This however was not implemented, and instead, we opted for a more “systems-of-systems” architecture, reasoning about individual systems and the boundaries between them. This again proved to be a successful strategy and allowed us to optimise the protocols in the IoT deployment site for their particular purpose. The micro-server was then particularly important in acting as an Internet gateway in our systems-of-systems approach.

Further investigation of systems-of-systems architectures is being carried out in the Dionasys project (<https://www.dionasys.eu/>); and one of their publications is [31].

The work highlighted the importance of a flexible and adaptive placement of computation in the resultant systems-of-systems approach. Our experiences indicate that it is far too simplistic to bring the data back to the cloud for processing, and this is particularly important in this domain. For example, in many circumstances, the data is rather static and predictable, but at times of extreme events, such as storms/heavy rainfall, the data then becomes critical to understand flooding behaviours of potential water quality issues. It is also too draining in resources (bandwidth and battery) to communicate all data and also sensors such as our accelerometers (on the sheep) generated vast volumes of data. These insights have led to future work in (i) the dynamic enhancement of the micro-server to be an edge device looking at constructing a micro-cloud out of a number of Raspberry Pis, (ii) strategies looking at the dynamic placement of computation on the resultant complex distributed environment (including the pushing of behaviour to end or edge devices, and (iii) self-organisation strategies for the above based on machine learning.

The combination of IoT technology and cloud computing is powerful—with IoT providing a source of rich, ubiquitous data coupled with an ability to intervene in the real world through means of actuation and cloud computing offering massive and elastic capacity in terms of storage and processing, coupled with an associated service architecture to interpret, manipulate, and visualise this data. There is however one missing piece—data science—and it is clear from our work that more effort is required to make sense of the rich volumes of data generated by an Environmental IoT. It is also clear that this domain represents different challenges for data science than many areas of big data; the data in this area tends to be a larger collection of heterogeneous data elements requiring complex analysis around, for example, interdependency. Researchers refer to this as the “long tail of science” [32]. As a further follow-up, we are investigating this eternal braid of the three synergistic technologies (as displayed in Figure 19) in the understanding, mitigating, and adapting to climate change.




Figure 19. The complementarity of IoT, cloud, and data science is a key field of future work for integrative research efforts each needing to understand the role and capabilities of the others.

6.2. Scientific Contributions

Arguably the biggest success of the project was that we provided a working and fully functional Environmental IoT that enabled new modes of scientific discovery. The key to this was the deployment of a wide range of sensors covering hydrology, soil science, and animal behaviour. This enabled a new kind of integrative scientific enquiry, and also by bringing the data together, this also brought the associated scientific communities together to collaboratively understand the interdependencies in the complex environmental catchment under consideration. This acts as a beacon for future integrative and collaborative science in this area. The examination of this multi-faceted data across different geographic areas and (importantly) at different scales also enables this broader and deeper understanding of a catchment in all its complexities. The more obvious but equally profound contribution is providing real-time data to scientists. As mentioned in related work, environmental and earth sciences are overly reliant on data logging which require trips out to (often remote and inaccessible) locations to download data. The accessibility of rich, real-time data therefore represents a real paradigm shift for this area

and a major opportunity to support improved response to issues such as, for example, extreme events and natural disasters.

One key issue that emerged was the need to understand the semantics of the data, most notably related to the provenance of data (what type of sensors, their perceived accuracy, and so on). For example, we deployed a range of soil sensors covering a range of prices and qualities. This is important as scientists are interested in what can be achieved from dense deployment of very cheap sensors but equally need more accurate technologies for calibration. Hybrid technologies are therefore important. We have thus devoted significant effort to the resultant (semantic) data model, covering provenance but also other meta-data descriptions and supporting more conceptual reasoning about the data. This data model is briefly mentioned in Sections 4 and 5 and is also the subject of a sister paper which will be published shortly.

A final important design decision was the provision of data in the cloud as a key means of enabling integration and collaboration but also crucially in supporting open access to data (available through MongoDB). The importance of open data has already been recognised by many observers, most notably by the Royal Society [33]. This is particularly important in the field of environmental and earth sciences as the communities respond to the requirements of a new kind of “big” science as required by the challenges of environmental change. Open data also encourages innovation around both uses of the data and (importantly) associated services. Access to the Environmental IoT data is available through the associated data portal, but the open API also encouraged consideration of a range of apps including a “MySheep” app and also the potential for integration into the existing MySoils app [34].

6.3. Methodological Contributions

It is often the case that IoT developments are carried out by the computer science community on behalf of end users, and this can result in technology for its own sake. In this case, it was crucial that the project was a truly equal collaboration between technologists and environmental scientists. This was necessary to allow the scientists to appreciate what is possible from the underlying technology, often to enable a broader vision to develop, and also for the technologists to fully understand what is needed to achieve a true paradigm shift in scientific understanding. This was achieved, representing arguably the biggest contribution to the project dwarfing, for example, any technological insights. The key to this was an agile development method and the associated use of flexible technologies such as Arduinos and Raspberry Pis. While these technologies are not necessarily optimal in terms of final development, they are ideal for rapid prototyping and exploring a space and indeed as a basis for having a conversation between the different constituent disciplines involved. It was notable that, within a short time, the scientists in the project were adapting the initial designs to better match their needs and also to extend the functionality. In effect, the technological and methodological approaches enabled a conversation on the role of Internet of Things technology in the environment, and there is also evidence emerging through follow-on projects that this has resulted in a deeper and more sustainable level of understanding.

7. Conclusions

This paper has examined the potential role of Internet of Things technology in supporting a deeper understanding of the natural environment, reporting in particular in the design and “in the wild” deployment of an exemplar Environmental IoT examining a range of environmental factors in a catchment in North Wales.

The project has utilised field deployment of sensors and cloud computing to offer a shared repository of data to stakeholders. This data can be easily accessible and can be used to feed into models, to deal with uncertainty in models, and to discuss the results with other stakeholders and communities. We have also reported our experiences based on these challenges through our IoT system.

The one barrier is the lack of maturity of Internet of Things technology based partly on the lack of standardization in this domain and also the many open research questions to be resolved, for example

around minimization of energy and self-management. The slow uptake of the technology from the lab to the field is because most of these systems rely on nonstandard, custom-designed elements that need specialised expertise. The other challenges faced in setting up an Internet of Things in the natural environment are power limitations, poor Internet connectivity, and hostile operational conditions. Nonetheless, our aim through our Environment IoT system is to provide a major opportunity to fill the gap between different environmental science disciplines. The idea is to bring together data and assets from across different domains (soil, water, plant ecology, and animals) and organisations (public, regulatory, and industry) to enable integrated problem solving for the agriculture and water industries but with a clear potential to link to other critical industries in the future. We are particularly interested in new kinds of collaborations and associated management strategies that can stem from bringing real-time data sets together and observing related interdependencies.

Prior to investing in expensive hardware, it is good to know if an IoT system can operate under the given conditions of an environment. We believe that the insight gained through our Environmental IoT project can serve as a basis for validating challenging environments when setting up an IoT system for such environments. Our overriding motivation was to develop a platform that was cheap and flexible and because we were limited in our equipment budget. Looking at other technologies, IoT solutions tend to be either too specific in their application domains or came at a higher price point per node than we could afford. On the other hand, Acclima was a preferred sensor to Grove by the environmental scientists; however, they wanted to experiment with a cheaper alternative as the Acclima sensor is very expensive. They were willing to compromise on quality for quantity of real-time data points as long as there were a few of the high-quality inputs in the field as a known reference point.

We have outlined some key areas of future research following on from this work most notably on areas such as systems of systems architectures, the role of edge devices in setting up a micro-cloud and bringing some computation with techniques such as machine learning on the micro-cloud, and self-organisation in complex distributed systems.

In concluding the paper, we return to our central hypothesis: *“A combination of IoT technology coupled with Cloud Computing enables a paradigm shift in our understanding and management of the natural environment, especially related to understanding ecosystem interdependencies, in times of unprecedented environmental change.”* This is clearly a rather bold and open-ended claim, but nevertheless our experiences indicate that this is an area of huge potential and that our experiments have shown that new styles of scientific enquiry can be supported by such emerging technology. We also flag this as an important area of future development of IoT technology. While it is often the case that IoT deployments feel like a technology in search of a problem, it is absolutely the case that, in the environmental and earth sciences, it is the other way round—a problem in search of a solution where at least part of the solution lies in the combination of cloud and IoT technologies, coupled with emerging techniques from data science as discussed above.

Author Contributions: Conceptualization, B.P., G.S.B., B.E., D.L.J., D.C., and W.S.; funding acquisition, B.P. and G.S.B.; investigation, V.N., B.E., J.C., D.L.J., D.C., R.S., and G.D.; methodology, V.N., B.W., P.B., and G.D.; resources, B.E., J.C., B.W., W.S., and M.B.; software, V.N., P.B., and G.D.; supervision, B.P. and G.S.B.; validation, B.E. and J.C.; writing—original draft, V.N. and B.P.; writing—review & editing, V.N., B.P., G.S.B., B.W., and I.U.

Funding: This research was funded by EPSRC (EP/L023636/1).

Conflicts of Interest: The authors declare no conflict of interest. The funders had no role in the design of the study; in the collection, analyses, or interpretation of the data; in the writing of the manuscript; or in the decision to publish the results.

References

1. Elkhatab, Y.; Blair, G.S.; Surajbali, B. Experiences of using a hybrid cloud to construct an environmental virtual observatory. In Proceedings of the 3rd International Workshop on Cloud Data and Platforms (CloudDP '13), Prague, Czech Republic, 14 April 2013; ACM: New York, NY, USA, 2013; pp. 13–18.

2. Delin, K.A.; Jackson, S.P.; Johnson, D.W.; Burleigh, S.C.; Woodrow, R.R.; McAuley, J.M.; Dohm, J.M.; Ip, F.; Ferré, T.P.; Rucker, D.F.; et al. Environmental Studies with the Sensor Web: Principles and Practice. *Sensors* **2005**, *5*, 103–117. [\[CrossRef\]](#)
3. Martinez, K.; Hart, J.K.; Ong, R. Environmental Sensor Networks. *Computer* **2004**, *37*, 50–56. [\[CrossRef\]](#)
4. Martinez, K.; Padhy, P.; Elsaify, A.; Zou, G.; Riddoch, A.; Hart, J.K.; Ong, H.L.R. Deploying a sensor network in an extreme environment. In Proceedings of the IEEE International Conference on Sensor Networks, Ubiquitous, and Trustworthy Computing (SUTC'06), Taichung, Taiwan, 5–7 June 2006; IEEE Computer Society: Washington, DC, USA, 2006; Volume 1, pp. 186–193.
5. Hart, J.K.; Martinez, K. Towards an Environmental Internet of Things [IoT]. *Earth Space Sci.* **2015**, *2*, 1–7. [\[CrossRef\]](#)
6. Rey-Moreno, C.; Bebea-Gonzalez, I.; Foche-Perez, I.; Quispe-Tacas, R.; Liñán Benitez, L.; Simo-Reigadas, J. A Telemedicine WiFi network optimized for long distances in the Amazonian jungle of Peru. In Proceedings of the 3rd Extreme Conference on Communication: The Amazon Expedition (ExtremeCom'11), Manaus, Brazil, 26–30 September 2011; ACM: New York, NY, USA, 2011; pp. 9:1–9:6.
7. Szewczyk, R.; Osterweil, E.; Polastre, J.; Hamilton, M.; Mainwaring, A.; Estrin, D. Habitat Monitoring with Sensor Networks. *Commun. ACM* **2004**, *47*, 34–40. [\[CrossRef\]](#)
8. Martinez, K.; Hart, J.K.; Ong, R. Deploying a wireless sensor network in Iceland. In Proceedings of the 3rd International Conference on GeoSensor Networks (GSN'09), Oxford, UK, 13–14 July 2009; Springer: Berlin/Heidelberg, Germany, 2009; pp. 131–137.
9. Hasler, A.; Talzi, I.; Beutel, J.; Tschudin, C.; Gruber, S. Wireless sensor networks in permafrost research: Concept, requirements, implementation, and challenges. In Proceedings of the 9th International Conference on Permafrost (NICOP'2008), Fairbanks, AK, USA, 29 June–3 July 2008.
10. Werner-Allen, G.; Lorincz, K.; Ruiz, M.; Marcillo, O.; Johnson, J.; Lees, J.; Welsh, M. Deploying a wireless sensor network on an active volcano. *IEEE Internet Comput.* **2006**, *10*, 18–25. [\[CrossRef\]](#)
11. Miorandi, D.; Sicari, S.; Pellegrini, F.D.; Chlamtac, I. Internet of things: Vision, applications and research challenges. *Ad Hoc Netw.* **2012**, *10*, 1497–1516. [\[CrossRef\]](#)
12. The UK Government Environment, Food and Rural Affairs Committee. *Natural Environment White Paper*; The UK Government Environment, Food and Rural Affairs Committee: London, UK, 2012. Available online: publications.parliament.uk/pa/cm201213/cmselect/cmenvfru/492/492.pdf (accessed on 7 June 2019).
13. Atzori, L.; Iera, A.; Morabito, G. The Internet of Things: A survey. *Comput. Netw.* **2010**, *54*, 2787–2805. [\[CrossRef\]](#)
14. Cerpa, A.; Elson, J.; Estrin, D.; Girod, L.; Hamilton, M.; Zhao, J. Habitat Monitoring: Application Driver for Wireless Communications Technology. *SIGCOMM Comput. Commun. Rev.* **2001**, *31*, 20–41. [\[CrossRef\]](#)
15. Kahn, J.M.; Katz, R.H.; Pister, K.S.J. Next Century Challenges: Mobile Networking for “Smart Dust”. In Proceedings of the 5th Annual ACM/IEEE International Conference on Mobile Computing and Networking (MobiCom'99), Seattle, WA, USA, 15–19 August 1999; ACM: New York, NY, USA, 1999; pp. 271–278.
16. Mainwaring, A.; Culler, D.; Polastre, J.; Szewczyk, R.; Anderson, J. Wireless sensor networks for habitat monitoring. In Proceedings of the 1st ACM International Workshop on Wireless Sensor Networks and Applications (WSNA'02), Atlanta, GA, USA, 28 September 2002; ACM: New York, NY, USA, 2002; pp. 88–97.
17. Szewczyk, R.; Mainwaring, A.; Polastre, J.; Anderson, J.; Culler, D. An Analysis of a large scale habitat monitoring application. In Proceedings of the 2nd International Conference on Embedded Networked Sensor Systems (SenSys'04), Baltimore, MD, USA, 3–5 November 2004; ACM: New York, NY, USA, 2004; pp. 214–226.
18. Juang, P.; Oki, H.; Wang, Y.; Martonosi, M.; Peh, L.S.; Rubenstein, D. Energy-efficient Computing for Wildlife Tracking: Design Tradeoffs and Early Experiences with ZebraNet. *SIGOPS Oper. Syst. Rev.* **2002**, *36*, 96–107. [\[CrossRef\]](#)
19. Martinez, K.; Padhy, P.; Riddoch, A.; Ong, R.; Hart, J. Glacial environment monitoring using sensor networks. In Proceedings of the Workshop on Real-World Wireless Sensor Networks (REALWSN'05), Stockholm, Sweden, 20–21 June 2005.
20. Cardell-Oliver, R.; Smettem, K.; Kranz, M.; Mayer, K. Field testing a wireless sensor network for reactive environmental monitoring [soil moisture measurement]. In Proceedings of the Intelligent Sensors, Sensor Networks and Information Processing, Melbourne, Australia, 14–17 December 2004; pp. 7–12.

21. Tolle, G.; Polastre, J.; Szewczyk, R.; Culler, D.; Turner, N.; Tu, K.; Burgess, S.; Dawson, T.; Buonadonna, P.; Gay, D.; et al. A MacroScope in the Redwoods. In Proceedings of the 3rd International Conference on Embedded Networked Sensor Systems (SenSys'05), San Diego, CA, USA, 2–4 November 2005; ACM: New York, NY, USA, 2005; pp. 51–63.
22. Langendoen, K.; Baggio, A.; Visser, O. Murphy Loves Potatoes: Experiences from a pilot sensor network deployment in precision agriculture. In Proceedings of the 20th International Conference on Parallel and Distributed Processing (IPDPS'06), Rhodes Island, Greece, 25–29 April 2006; IEEE Computer Society: Washington, DC, USA, 2006; p. 174.
23. Musaloiu-E, R.; Terzis, A.; Szlavecz, K.; Szalay, A.; Cogan, J.; Gray, J. Life under your feet: A wireless soil ecology sensor network. In Proceedings of the Third Workshop on Embedded Networked Sensors, Cambridge, MA, USA, 30–31 May 2006.
24. Bishop-Hurley, G.J.; Swain, D.L.; Anderson, D.M.; Sikka, P.; Crossman, C.; Corke, P. Virtual Fencing Applications: Implementing and Testing an Automated Cattle Control System. *Comput. Electron. Agric.* **2007**, *56*, 14–22. [CrossRef]
25. Hughes, D.; Greenwood, P.; Blair, G.; Coulson, G.; Grace, P.; Pappenberger, F.; Smith, P.; Beven, K. An experiment with reflective middleware to support grid-based flood monitoring. *Concurr. Comput. Practice Exp.* **2008**, *20*, 1303–1316. [CrossRef]
26. Mo, L.; He, Y.; Liu, Y.; Zhao, J.; Tang, S.J.; Li, X.Y.; Dai, G. Canopy closure estimates with GreenOrbs: Sustainable sensing in the forest. In Proceedings of the 7th ACM Conference on Embedded Networked Sensor Systems (SenSys'09), Berkeley, CA, USA, 4–6 November 2009; ACM: New York, NY, USA, 2009; pp. 99–112.
27. Ingelrest, F.; Barrenetxea, G.; Schaefer, G.; Vetterli, M.; Couach, O.; Parlange, M. SensorScope: Application-specific Sensor Network for Environmental Monitoring. *ACM Trans. Sen. Netw.* **2010**, *6*, 17:1–17:32. [CrossRef]
28. Corke, P.; Wark, T.; Jurdak, R.; Hu, W.; Valencia, P.; Moore, D. Environmental Wireless Sensor Networks. *Proc. IEEE* **2010**, *98*, 1903–1917. [CrossRef]
29. Colitti, W.; Steenhaut, K.; Caro, N.D.; Buta, B.; Dobrota, V. REST enabled wireless sensor networks for seamless integration with web applications. In Proceedings of the IEEE Eighth International Conference on Mobile Ad-Hoc and Sensor Systems, Valencia, Spain, 17–22 October 2011; pp. 867–872.
30. Ko, J.; Eriksson, J.; Tsiftes, N.; Dawson-Haggerty, S.; Durvy, M.; Vasseur, J.; Terzis, A.; Dunkels, A.; Culler, D. Beyond interoperability: Pushing the performance of sensor net IP stacks. In Proceedings of the ACM Conference on Networked Embedded Sensor Systems (SenSys'2011), Seattle, WA, USA, 1–4 November 2011; ACM: New York, NY, USA, 2011.
31. Blair, G.; Bromberg, Y.D.; Coulson, G.; Elkhatib, Y.; Réveillère, L.; Ribeiro, H.B.; Rivière, E.; Taïani, F. Holons: Towards a systematic approach to composing systems of systems. In Proceedings of the 14th International Workshop on Adaptive and Reflective Middleware, Vancouver, Canada, 7–11 December 2015; ACM: New York, NY, USA, 2015; pp. 5:1–5:6.
32. Heidorn, P.B. Shedding light on the dark data in the long tail of science. *Libr. Trends* **2008**, *57*, 280–299. [CrossRef]
33. Boulton, G.; Campbell, P.; Collins, B.; Elias, P.; Hall, W.; Laurie, G.; O'Neill, O.; Rawlins, M.; Thornton, J.; Vallance, P.; et al. Science as an Open Enterprise. *R. Soc.* **2012**. Available online: https://openaccess.sdum.uminho.pt/wp-content/uploads/2013/02/1_GeoffreyBoulton_OpenAIREworkshopUMinho.pdf (accessed on 5 June 2019).
34. Shelley, W.; Lawley, R.; Robinson, D.A. Technology: Crowd-sourced soil data for Europe. *Nature* **2013**, *496*, 300. [CrossRef] [PubMed]

