

Springbrook: Challenges in Developing a Long-Term, Rainforest Wireless Sensor Network

Tim Wark¹, Wen Hu¹, Peter Corke¹, Jonathan Hodge², Aila Keto³, Ben Mackey¹, Glenn Foley¹, Pavan Sikka¹, Michael Brünig¹

¹ Autonomous Systems Laboratory, CSIRO ICT Centre, Brisbane, QLD, Australia
<firstname.lastname>@csiro.au

² Environmental Protection Agency (EPA), Brisbane, QLD, Australia
jonathan.hodge@epa.qld.gov.au

³ Australian Rainforest Conservation Society (ARCS), Bardon, QLD, Australia
aila.keto@rainforest.org.au

Abstract

We describe the design, development and learnings from the first phase of a rainforest ecological sensor network at Springbrook — part of a World Heritage precinct in South East Queensland. This first phase is part of a major initiative to develop the capability to provide reliable, long-term monitoring of rainforest ecosystems. We focus in particular on our analysis around energy and communication challenges which need to be solved to allow for reliable, long-term deployments in these types of environments.

1. INTRODUCTION

A. Background

Wireless sensor networks (WSNs) continue to emerge as a technology that will transform the way we measure, understand and manage the natural environment. Whilst there has been considerable progress made over the past several years in bridging the gap between theoretical developments and real deployments, the field is still relatively immature and widespread use of environmental WSNs is not yet a reality. As such there is still a wide range of fundamental and applied research questions around dealing with the constraints placed on environmental WSNs which need to be addressed.

In this work, we describe the design, development and learnings from the initial phase of a rainforest ecological sensor network at Springbrook, Queensland — see Figure 1. The Springbrook site is part of a World Heritage precinct centred on the extinct Tweed shield volcano. Diverse and compressed environmental gradients present significant technical and ecological challenges and opportunities.

The focus for the Springbrook initiative centres on cost effective recovery of several hundred hectares of globally significant forests that were cleared during the last century. Traditional methods are prohibitively expensive and risk not achieving historical fidelity, viability, resilience or adaptive capacity for the forests and their ancient lineages of fauna

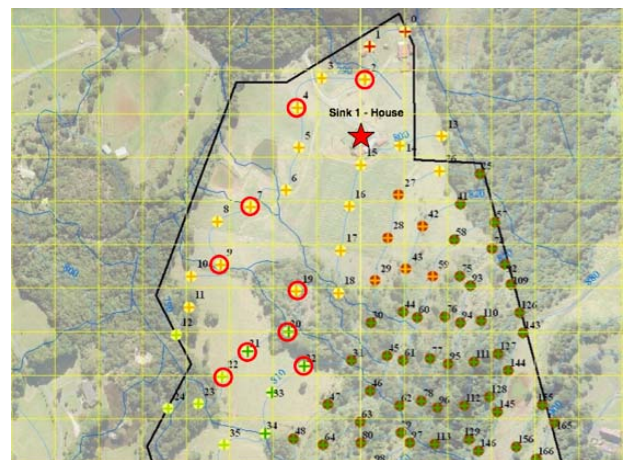


Fig. 1: Aerial photograph of the Springbrook site. Phase 1 node locations (shown in red) are overlaid on positions of future node positions for a phase 2 deployment. A tenth node "log-runner" is also located at an ecologically significant, elevated site just outside the limits of this map. (Aerial photography was provided to ARCS by Gold Coast City Council.)

and flora. To achieve the optimal balance between natural and assisted restoration, identification of biotic and abiotic barriers to natural regeneration at ecologically significant spatial and temporal scales will be critical. As such, no other technology holds the same potential as sensor networks for in situ, real time monitoring of multiple parameters for multivariate analysis of complex, interlinked systems of climate, vegetation, fauna, soils, hydrology and geomorphology involved in ecological restoration.

B. Related Work

A number of recent studies have increased focus on issues and challenges around power management and reliability in long-term, environmental sensor networks [1], [2]. These studies have recognised the wide gap that often exists between predicted performance of networks and observed performance in the field.

The potential of sensor networks as a means for answering important ecological questions is also well recognised [3], [4] in the sensor networks and ecological communities. A survey of 52 recent papers from the journal *Ecology*, indicated that nearly all experiments were conducted at either very low temporal-resolution or over small spatial areas [3].

Within the sensor networks community, a number of studies have focussed on relatively short deployments to help develop new tools and methodologies needed for environmental sensor networks for both plant and animal species observations. The *LUSTER* [5] WSN explored a number of the issues around delay-tolerance and reliability in networks measuring light levels under shrubs. The extension from static environmental monitoring to bird monitoring has been explored in a small study around the use of small RFID tags attached to birds [6]. Likewise tasks such as bio-acoustic monitoring is also an emerging area of research interest [7].

For all environmental sensor networks however, one of the most dominant challenges is finding a long-term source of energy. In order for these networks to be sustainable, nodes must be able to be powered without the need for human intervention for long periods. The source of energy that has been most actively explored to date is that of solar energy. A number of recent studies [8], [9] have quantified the amount of solar energy which can be harnessed under various environmental conditions. These studies have also investigated the power management designs which are most effective for recharging batteries.

C. Motivation

The ultimate goal of our broader work is in the development of sensor networks as a means for tracking restoration of biodiversity. Our primary motivation in this first phase has been to develop a better understanding around the challenges in deploying long-term, low-power, wireless sensor networks in environments such as rainforests. These environments are typically characterised by areas with very limited solar energy and adverse and dynamic radio environments. In order to develop the network and energy management protocols that are required to enable robust and reliable performance of long-term, rainforest networks, we must first fully quantify the performance of current networks under these types of conditions.

Our key contributions are:

- Development and deployment of a new multi-hop, microclimate rainforest sensor network.
- Full evaluation of energy characteristics through open and rainforest areas.
- Full evaluation of network performance under open and rainforest areas given a dynamic link-quality (LQ) routing protocol.

2. SYSTEM OVERVIEW

A. Deployment Design

The first phase of the Springbrook deployment consists of 10 microclimate nodes returning data to a single sink/gateway

point. Node locations, as shown in Figure 1, were chosen to maximise the amount of ecological information that could be derived from the node data as well as to ensure we covered the whole range of conditions from open areas to dense rainforest. A particular focus was placed on measuring the extremes in daily energy availability and wireless link quality in various areas of the rainforest region.

Each node was equipped with sensors for wind speed and direction, temperature, relative humidity, soil moisture and leaf wetness. Node electronics, batteries and radio antennas were designed to fit within a single plastic housing as shown in Figure 2. In order to mount nodes, a custom frame was designed and constructed from light-weight, hollow metal poles. The frame was deployed by inserting a star-picket into the ground and then attaching the frame to the star-picket using U-clamps. This design proved to be very sturdy as well as relatively easy to deploy by a team of two people.



Fig. 2: A microclimate node deployed in the Springbrook rainforest.

B. Hardware Platform

The hardware platform for all nodes is based around the FleckTM-3 wireless sensor platform [8]. The FleckTM-3 uses the Atmega128 micro-controller along with the Nordic NRF905 radio transceiver operating in the 915MHz band. The device also incorporates a real-time clock chip to reduce micro-controller overheads for timing operations. The architecture relies heavily on the SPI bus where the Atmega128 acts as the SPI master and can communicate with the radio, the flash memory and the real-time clock over the SPI interface.

Nodes are powered from a combination of rechargeable NiMH batteries and solar cells [8]. In our current design we use three 1.2V 2700mAh in series and mono-crystalline solar panels capable of supplying up to 300mA of current.

C. Software Platform

Within the WSN operating system community there are two broad approaches for embedded operating systems: event-based and thread-based systems — and there is healthy debate about the relative merits of different approaches [10]. We have developed our own operating system FleckOS (FOS) which falls into the latter category [11]. FOS provides a priority-based, non-preemptive (cooperative) threading environment

with separate stacks for each thread, which has the advantage of providing a simple concurrent programming model which does not require semaphores. All application software on the microclimate nodes ran on top of this operating system.

D. Network Protocols

The network topology of the phase 1 network was a multi-hop configuration, where data was directed downstream from each node to a single sink point. Previous outdoor sensor network deployments [12] show that the quality of outdoor wireless links are highly dynamic because of the frequent changes of weather conditions such as fog and rain. Therefore, an Expected Number of Transmissions (ETX) based routing protocol [13], which takes the link quality of both transmission directions into account, was used to increase end-to-end transmission delivery rates and reduce network energy consumption. For a path comprising h hops, the ETX was calculated as:

$$ETX = \sum_{i=1}^h \frac{1}{p_i q_i} \quad (1)$$

where p_i is the upstream link quality, and q_i is the downstream link quality of hop i .

A Low Power Listening (LPL) [14] Medium Access Control (MAC) protocol was used to duty cycle the NRF905 transceivers, which consume a significant portion of the energy in the nodes. The checking period of our LPL implementation was 3 ms and the length of one LPL cycle was 60 ms. Therefore, the lower bound of the transceiver duty cycle is 5%. Whilst there are a number of alternative MAC protocols we could have used, the purpose of this first phase of work was to evaluate the characteristics of network link qualities rather than implement the optimal MAC protocol.

Link layer retransmission is commonly used to increase end-to-end transmission reliability in wireless sensor networks. Previous research [13] shows that, for a radio transceiver that can support variable packet lengths, ETX explicit acknowledgment (eAck) transmits fewer bytes (and is therefore more energy efficient) than ETX implicit acknowledgment (iAck) when link quality is dynamic. In this deployment, the transmission time of a data packet is 60 ms and the transmission time of an acknowledgment packet is approximately 6 ms, (i.e. the cost of data packet transmission is approximately 10 times more than the acknowledgment packet), thus we have chosen to use explicit acknowledgement as the link layer retransmission mechanism.

E. Data Management and Visualisation

The network “gateway” connects the sensor network to the outside world. We have designed a custom gateway box that consists of a compact *Alix PC*¹ with a USB connection to a base FleckTM-3 node along with an ethernet connection to a 3G (Telstra NextG) modem. This hardware was set up indoors at the gateway location, shown by the star in Figure 1. A

¹<http://www.pceingines.ch/alix.htm>

Universal Power-Supply (UPS) was also installed to ensure the gateway would not be affected by any short power outages.

In order to enable fast and easy installation of the sensor network gateway, we have developed custom software that enables the gateway to auto-configure itself for streaming data over the backlink. Given the name and location of the remote database, the gateway box will automatically establish a secure data tunnel from the gateway to the database and start the client process which will stream sensor data from the network to the database. The current database is MySQL running on a Linux-based server however we are currently transitioning to a more powerful Oracle database.

Our system also features a web interface, designed to allow easy visualisation of data returned from each node as well as to flag problems in the network - e.g. nodes not heard from for a while or network backlink down. A screenshot from our newly developed web interface is shown in Figure 3 demonstrating the graphing capability of the interface.

3. EVALUATION

The network was deployed at the Springbrook site in May 2008 and has been running continuously since. For the phase 1 deployment, a sample interval of 5 minutes was chosen for each sensor. In addition to logging all environmental measurements, we also returned all relevant “engineering” data including solar and battery voltage and current, hop counts, parent node IDs and link qualities. We present and evaluate the performance of various aspects of the network in the following subsections.

A. Network Performance

As described in Section 2-D, data was returned over the network via a multi-hop, routing protocol where the best quality links were favoured as the means to pass data downstream towards the sink. Figure 4 shows the average number of hops that each node took to get data back to the sink. As expected there is a clear correlation between those nodes which are geographically close to the sink (e.g. 2, 4) and a low number of hops. Nodes in the rainforest (e.g. 19, 20) tend to have the highest number of hops due to the larger geographic distances. As we discuss later, the rainforest foliage has clear effects on link quality of longer hops, meaning shorter hops are generally chosen.




Fig. 4: Average hopcount from each node back to sink over a week period. “Log R” refers to the log-runner node - see footnote on next page.

Figure 5 shows the throughput (number of messages received at base compared with number that should have been

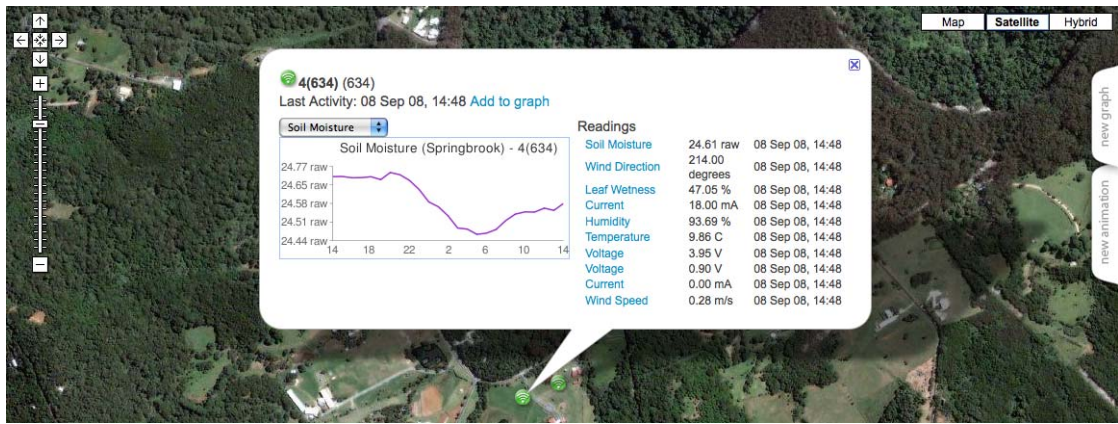


Fig. 3: Screenshot of the web interface used to visualise nodes at Springbrook rainforest.

theoretically sent) for each node. The top plot shows the performance when the network was first deployed. It can be seen that nodes in the open generally have throughputs of around 80% however throughput drops dramatically for rainforest nodes. As we discuss in Section 3-B this drop is not just a function of reduced radio performance but also due to the fact that there is not enough solar energy for rainforest nodes to run continuously once their battery stores are exhausted.

Due to the ecological importance of data coming from the “log runner”² node, during July 2008 an extra relay node was installed to assist with data coming back from this node. (The log runner node was much further from the sink than any other nodes.) At the same time an additional relay node was also placed on the roof of the house. As such, the bottom plot of Figure 5 shows the effect after these relays node had been installed. We can see that the throughput for all nodes has increased with the exception of node 19 which is at zero. This is due to virtually no solar energy at this point, so once the batteries were exhausted after initial deployment, no more data was sent back.

To help illustrate the variability of the radio link-quality, we can show how much the parent node varies over time. Figure 6 shows the proportion of parent nodes which were chosen by each node. As can be seen, most nodes have a dominant parent node and choose between three of four other neighbours from time to time.

B. Energy Characteristics

One of the most important variables to characterise is the amount of solar energy that is available in all parts of a rainforest network such as Springbrook. The impact of rainforest cover on incoming solar energy is significant as shown in the comparison of typical solar panel current in Figure 7.

We can further explore the impact of rainforest cover by looking at the relationship between solar panel voltage and output power of the panel. Figure 8 shows this relationship for an open and rainforest node for various operating points

²The *log runner* site is unique in that it is one of the few known locations of the rare log-runner bird species.




Fig. 5: Average percentage of messages (one week period) returned when nodes were initially deployed and when relay nodes were later added to the network. It can be observed that node 19 has since stopped sending any messages - due to no solar energy being received at this location.

throughout a typical day. In the case of the rainforest node the maximum power generated only just reaches 200mW whereas the node in the open reaches 1000mW.

Finally, we can calculate the average daily energy harvested by each node in the network by integrating the solar output power over time — see Figure 9. Whereas nodes in the open typically harvest over 10kJ/day, some rainforest nodes³ harvest as little as 100J on average — only 1% of the energy of nodes in the open!

C. Wireless Link Quality

Whilst lack of solar energy is clearly the biggest impact of rainforest cover on network nodes, it is also important to assess the impact on radio performance. As described in Section 2-D our routing protocol continually calculates the upstream (RX) and downstream (TX) link qualities (LQ) for each node.

³To keep important data coming back from the logrunner node, a temporary external solar panel has since been placed in a sunny location, attached by a long cable.




Fig. 6: Proportion of parent nodes utilised by each node within network over a week period.




Fig. 7: Plot of typical daily patterns of incoming solar panel current from nodes in open and rainforest locations. (Note the different axis scales)

Figure 10 shows the relationship between these metrics for a group of open nodes and a group of rainforest nodes over a several day period. Whilst the rainforest nodes have a slightly lower distribution of LQ scores, they still roughly match the distribution of nodes in the open.

We can gain more insight about the reduced performance of rainforest nodes by also looking at the relationship between node throughput and hop-count. Figure 11 is derived by calculating the node throughput during blocks of 6 hours over a one week period. During each block the average hop-count is also calculated. For nodes in the open, throughput stays very high with hop variations of between 1 and 4. For rainforest nodes, the hop-count is a lot higher - however this is partly due to the fact that most of these nodes are further from the sink.

The most interesting observation however is that there are two clear clusters of throughput for rainforest nodes. One cluster has very high throughput (like nodes in the open) however the other cluster is very low. Given that LQs, as per Figure 10, are very similar between open and rainforest nodes then our conclusion must be that these periods of low throughput are partly as a result of nodes having times of no




Fig. 8: Plot of relationship between solar panel current and solar panel output power for various operating points throughout a single day for node in open and node in rainforest location.




Fig. 9: Average daily energy harvested by each node.

or intermittent energy.

The final environmental factor we wish to analyse is the impact of rain. We can detect a rain event by a step increase in soil moisture. Figure 12 shows timeseries data for a period of time with two rain events for node 32. This node sits on the border of rainforest which blocks any direct radio communications back to the sink. It is clear that the link quality of the node drops dramatically after each rain event which also results in an increase in the number of hops for the node for a short period. It is also interesting to observe the relationship between relative humidity and link quality. It would appear that during periods of very dynamic humidity, the node link quality also drops significantly. Ongoing observation of network performance will be needed to clarify this observation — this will be a focus of future work.

4. DISCUSSION AND FURTHER WORK

We have described and evaluated the performance of our phase 1 rainforest network which has been running since May 2008. For nodes in the open we observe that there is more than sufficient solar energy to sustain continuous operation where data throughput is nearly 100%. As expected, results from the network show clear challenges for node locations which are in thick parts of the rainforest. These nodes receive as little as 1% of the daily solar energy that nodes in open areas receive. Given the current load of each node this is not enough energy to sustain continuous operation.

An interesting observation is that link qualities for rainforest nodes are generally the same as nodes in open locations. It is also clear that these link qualities drop dramatically after rain events. As such our conclusion would be that under normal conditions the effect of the rainforest on radio links is limited, however when there is significant moisture around the radio quality drops dramatically.




Fig. 10: Distribution of network routing link qualities (RX and TX) for all open and forest nodes over a four day period.




Fig. 11: Illustration of differences in hop count and throughput between open and rainforest nodes.

As we move toward the second phase of the network (where we expect to increase the network to over 200 nodes), our work is focussing on finding new ways that nodes in these types of locations can significantly reduce their workload to make it possible to sustain long-term monitoring of biodiversity. This can include forcing these nodes to be leaf (non-routing) nodes only, as well as only sending back the necessary information to summarise measurements rather than every data sample.

Our ultimate aim is for nodes to adaptively reduce their workload to maintain continuous operation rather than exhausting all energy resources and remaining offline for long-periods. In particular we will focus on new methods to allow independent adaptation of radio and sensing duty cycles given that samples do not necessarily have to be sent with every transmit cycle in order to effectively describe the phenomena being measured.

ACKNOWLEDGMENTS

The authors wish to thank Markus Burger, Les Overs, Stephen Brosnan and John Witham (CSIRO) for undertaking software development and hardware engineering in the development of the network. We also acknowledge the expert advice and support of Dr Keith Scott and Peter Johnson (ARCS) in the design and deployment phases.

REFERENCES

- [1] K. Langendoen, A. Baggio, and O. Visser, "Murphy loves potatoes: Experiences from a pilot sensor network deployment in precision agriculture," in *Parallel and Distributed Processing Symposium (IPDPS)*, 2006.
- [2] G. Barrenetxea, F. Ingelrest, G. Schaefer, M. Vetterli, O. Couach, and M. Parlange, "Sensorscope: Out-of-the-box environmental monitoring," in *Information Processing in Sensor Networks (IPSN)*, pp. 332–343, 2008.

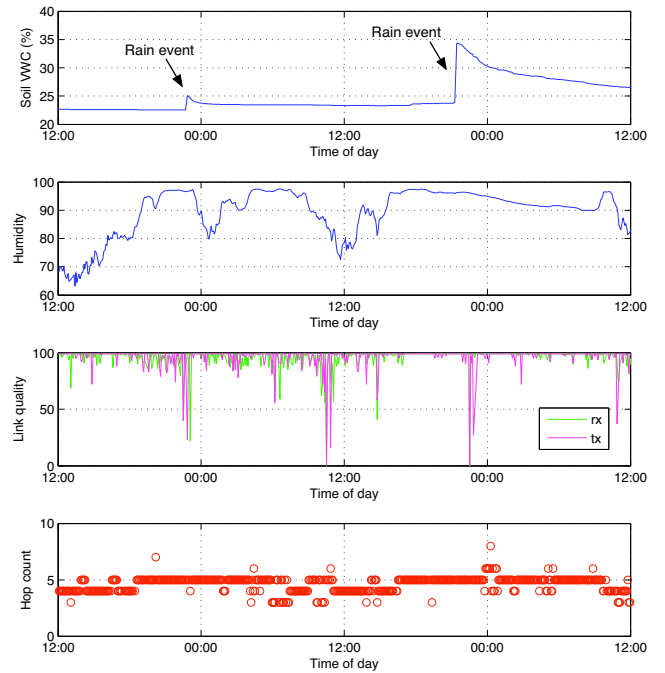


Fig. 12: Example of the effect of rain events and changes in relative humidity on wireless link quality for node 32.

- [3] J. Porter, P. Arzberger, H.-W. Braun, P. Bryant, S. Gage, T. Hansen, P. Hanson, C.-C. Lin, F.-P. Lin, T. Kratz, W. Michener, S. Shapiro, and T. Williams, "Wireless sensor networks for ecology," *BioScience*, vol. 55, no. 7, pp. 561–572, 2005.
- [4] J. Polastre, R. Szewczyk, A. Mainwaring, D. Culler, and J. Anderson, *Wireless Sensor Networks*, ch. Analysis of Wireless Sensor Networks for Habitat Monitoring, pp. 399–423. Springer, 2007.
- [5] L. Selavo, A. Wood, Q. Cao, T. Sookoor, H. Liu, A. Srinivasan, Y. Wu, and W. Kang, "Luster: wireless sensor network for environmental research," in *Proceedings of the 5th international conference on Embedded networked sensor systems (SenSys)*, pp. 103 – 116, 2007.
- [6] T. Naumowicz, R. Freeman, A. Heil, M. Calsyn, E. Hellmich, A. Brändle, T. Guilford, and J. Schiller, "Autonomous monitoring of vulnerable habitats using a wireless sensor network," in *RealWSN*, 2008.
- [7] M. Allen, L. Girod, R. Newton, S. Madden, D. T. Blumstein, and D. Estrin, "Voxnet: An interactive, rapid-deployable acoustic monitoring platform," in *Information Processing in Sensor Networks (IPSN) SPOTS*, 2008.
- [8] P. Corke, P. Valencia, P. Sikka, T. Wark, and L. Overs, "Long-duration solar-powered wireless sensor networks," in *EmNets*, pp. 33–37, 2007.
- [9] J. Taneja, J. Jeong, and D. Culler, "Design, modeling, and capacity planning for micro-solar power sensor networks," in *Information Processing in Sensor Networks (IPSN)*, pp. 407–418, 2008.
- [10] R. von Behren, J. Condit, and E. Brewer, "Why events are a bad idea (for high-concurrency servers)," *Proceedings of the 2003 HotOS Workshop*, May, 2003.
- [11] P. Corke, "FOS — a new operating system for sensor networks," in *European Conference on Wireless Sensor Networks (EWSN)*, (Poster), 2008.
- [12] T. L. Dinh, W. Hu, P. Sikka, P. Corke, L. Overs, and S. Brosnan, "Design and deployment of a remote robust sensor network: Experiences from an outdoor water quality monitoring network," in *Second IEEE Workshop on Practical Issues in Building Sensor Network Applications (SenseApp 2007)*, (Dublin, Ireland), Oct. 2007.
- [13] S. Rothery, W. Hu, and P. Cork, "An empirical study of data collection protocols for wireless sensor networks," in *ACM Workshop on Real-World Wireless Sensor Networks, (REALWSN'08)*, (Glasgow, Scotland), April, 2008.
- [14] J. Hill and D. Culler, "Mica: a wireless platform for deeply embedded networks," *IEEE Micro*, vol. 22, pp. 12–24, Nov. 2002.