

MobiCom 2009 Poster: Why Are Long-Term Large-Scale Wireless Sensor Networks Difficult? Early Experience with GreenOrbs

Yuan He[†]

heyuan@cse.ust.hk

Lufeng Mo[§]

mo.lufeng@stu.xjtu.edu.cn

Yunhao Liu[†]

liu@cse.ust.hk

[†]Hong Kong University of Science and Technology

[§]Xi'an Jiao Tong University

Wireless sensor networks in the past decade have achieved remarkable progress, while the real applications are still far from being long-term or large-scale. This paper presents GreenOrbs [1], the latest effort to explore the fundamental challenges of long-term large-scale wireless sensor networks. GreenOrbs supports a series of forestry applications. Based on the early experience with GreenOrbs, this paper further discusses the future research directions.

I. Introduction

Countless projects have been launched for study on wireless sensor networks (WSNs). In the past decade, there have been a number of well-known WSN systems and deployments, but none of them is indeed long-term large-scale. ExScal [2] included more than 1000+ nodes, but it was kept in operation for only a few days. A SensorScope [3] deployments is the longest among all the previous deployments, lasting for around 6 months, but the system scale is restricted to no more than 97 nodes.

Why are long-term large-scale WSNs difficult?

What are the fundamental challenges?

Bearing the above questions in mind, we have launched GreenOrbs, a collaborative research project to study long-term large-scale WSNs in the forest. The early experience in GreenOrbs delivers many lessons on deploying a long-term large-scale WSN in the wild environments. This paper presents our recent advances with GreenOrbs and discusses the future research directions.

II. The GreenOrbs Project

II.A. Forestry Applications

GreenOrbs is motivated by the need of long-term large-scale sensing for continuous forest surveillance, precise forestry measurements, and forestry research. The first application is canopy closure estimates [4]. Previous approaches of canopy closure estimates have either poor accuracy or prohibitive cost. Using WSN as a technique of quantitative measurement, GreenOrbs realizes accurate and economical estimates of vast forest. The second application is on carbon sequestration. The capacity of carbon sequestration of different tree species needs to be accurately measured, as can be realized with carbon dioxide sensors in the three-dimensional forest space.

GreenOrbs plans to support fire risk evaluation and study on biodiversity.

II.B. Implementation and Deployments

We use TeloB [5] motes with MSP430 F1611 processor and CC2420 radio. The software on the sensor nodes is based on TinyOS 2.1. CTP [6] is adopted for multi-hop data collection and modified to save communication cost. Data disseminations from the sink are enabled to control the nodes' operational parameters, such as the transmission power, sampling frequency, and duty cycle.

We have carried out a number of GreenOrbs deployments, including the mountain deployments and the prototype deployments. As of May 2004, the most lasting GreenOrbs deployment has been in continuous operation for 12 months. We've collected more than 5,240,000 data packets, which account for over 600M bytes.

The prototype system is deployed on the campus woodland, as shown in Figure 1. The deployment area is around 40,000m². The deployment started in May 2009 and included 50 nodes. In November 2009 it was expanded to include 330 nodes. The network diameter is 12 hops. The mountain deployment includes 200 nodes and has been in continuous operation since August 2009. The deployment area is around 200,000m². The network diameter is 20 hops. The duty cycle of GreenOrbs nodes is usually 5%.

III. Experience and Future Directions

III.A. Diagnosis

Deployed in the wild environments, sensor nodes tend to be unreliable and error-prone in operations. GreenOrbs has experienced various problems in the deployments, such as node failure, reading errors, packet loss, and software bugs, etc. Those problems lead to performance degradation of the network.



Figure 1: The bird's-eye picture of GreenOrbs prototype




Figure 2: Daily number of nodes that successfully return data

Figure 2 plots the daily number of nodes that successfully deliver their data to the sink during a 57-day 110-node deployment. We observe decreases of number for multiple times, which indicate latent problems in the network. Since by far there is no universal diagnosis tool, we can only conjecture the causes of those problems. For example, a node that successfully forwards data but fails to deliver its own data must undergo a node failure or packet loss. We used to take three attempts to resolve problems, e.g. by resetting the possibly problematic nodes. But the overall trend shows the number of nodes clearly keeps decreasing. Without comprehensive diagnosis, we are unable to locate, identify, and resolve the problems completely. To make it even more challenging, it is in nature difficult to collect all the required information from a resource-constrained WSN via unreliable wireless communications.

III.B. Localization

Localization is a fundamental issue that has been extensively studied in the literature. The real-world experience from GreenOrbs reveals that localization in the wild environments remains very challenging, in spite of the substantive efforts existing in the literature.

The non-uniform deployment of sensor nodes inevitably causes anisotropic problem. Nevertheless, the received signal strength indications (RSSI) used for ranging are highly irregular, dynamic, and asymmetric between pairs of nodes. To make it even worse, the complex terrain and obstacles in the forest easily affect RSSI-based range measurements, thus




Figure 3: Errors of range measurements on the nodes

incurring undesired but ubiquitous errors. Figure 3 plots the mean ranging error of 100 GreenOrbs nodes. We can see that only 9% nodes have large ranging errors ($>5\text{m}$) and only 18% nodes have small ranging errors ($<1\text{m}$). The errors of the rest nodes (73%) are between 1m and 5m. Such errors cannot be easily detected but seriously degrade the overall localization accuracy.

III.C. Routing

WSNs mostly rely on multi-hop transmissions to deliver data packets. Figures 4 and 5 plot the distributions of traffic among the nodes and the time that a node is under high traffic load. 5% nodes carry over 80% of the total traffic, while the traffic load on nearly 90% nodes is very low.

Another interesting finding is about the causes of packet loss. Figure 6 shows `Transmit_Timeout` (the number of retransmissions exceeds the limit so that the packet is dropped by the sender) accounts for 61.08%, which are evenly distributed among many nodes. Meanwhile, `Receive_Pool_Overflow` (the receiving pool on a forwarding node is full so that an upcoming packet is dropped) accounts for the rest 38.92%, which takes places on only 5% nodes.

Based on the above observation, future research on routing aims to explore the answers to the following questions: Are there any critical nodes and links in a WSN with regard to routing? Do the existing routing mechanisms fully utilize the network bandwidth? How to measure the quality of routing paths, so as to provide a comprehensive routing metric?




Figure 4: Traffic distribution among the nodes




Figure 5: CDF of the time that a node is under high traffic load




Figure 6: Cumulative distribution on packet drops

III.D. Miscellaneous

Wireless reprogramming. The programs on the sensor nodes often need to be upgraded or replaced after deployment due to various reasons, such as correcting bugs and upgrading program functions. The GreenOrbs nodes are deployed in the wild environments where it is extremely difficult to collect back the deployed nodes. Enabling sensor nodes to be reprogrammable over the air is a crucial technique to support long-term continuous deployments of WSNs.

Wireless reprogramming of WSNs mainly faces the following challenges. First, the code size greatly affects the reprogramming efficiency. Most sensor nodes have very limited memory and cannot simultaneously accommodate two program images. A large code also results in a long transmission time and much energy consumption during dissemination. Second, the loading cost for executing the new code must be appropriate. Specifically, we should try to avoid hardware rebooting, which incurs higher energy consumption and loss of sensor data. Third, the relatively long dissemination process increases the vulnerability of reprogramming to unexpected failures and packets losses. It is hard to keep the sensor nodes controllable and consistent with each other throughout the reprogramming process.

Outlier detection. As we observe with GreenOrbs, outliers are frequently present in sensor data. We would like to emphasize that many WSNs applications, e.g. the forestry applications, demand highly complete data set. The existing schemes of outlier detection are either centralized or inefficient. Centralized detection incurs excessive processing delay, which fails to meet the needs by applications with stringent timing requirements. Considering that the general sensor nodes have very limited program memory and network bandwidth, distributed schemes for outlier detection should be not only efficient in computation, but also light-weight with respect to communication and memory cost.

The future research directions of GreenOrbs also include novel sensing devices for forestry sensing and the interdisciplinary study between WSN and forestry ecosystem research.

IV. Acknowledgement

Being established in China, GreenOrbs is a collaborative research project among more than 10 domestic and overseas universities [1].

We would like to thank the support from the co-investigators including Dr. Jizhong Zhao (XJTU), Dr. Guomo Zhou (ZJFU), Dr. Xiangyang Li (IIT), Dr. Guojun Dai (HDU), Dr. Ming Gu (THU), Dr. Huadong Ma (BUPT), and Dr. Mo Li (HKUST).

We would also thank the project team members who have contributed a lot to the system development and deployments of GreenOrbs, including Jiliang Wang, Xin Miao, and Zhichao Cao from HKUST, Kebin Liu from SJTU, Wei Dong from ZJU, Min Xi, Wei Xi, Shuo Lian, and Rui Li from XJTU, Xingfa Shen from HDU, Danning Chen and Guowei Zou from THU, Chao Wang from BUPT, Xufei Mao and Shaojie Tang from IIT, Shuo Liu from ZJFU, Tao Chen from NUDT, and Zhengge Guo from YSU of China.

References

- [1] GreenOrbs, <http://www.greenorbs.org/>
- [2] A. Arora, R. Ramnath, E. Ertin, *et al.*, "ExScal: Elements of an Extreme Scale Wireless Sensor Network," in *IEEE RTCSA*, 2005.
- [3] G. Barrenetxea, F. Ingelrest, G. Schaefer, and M. Vetterli, "The Hitchhiker's Guide to Successful Wireless Sensor Network Deployments," in *ACM SenSys*, 2008.
- [4] L. Mo, Y. He, Y. Liu, J. Zhao, S. Tang, X.-Y. Li, and G. Dai, "Canopy Closure Estimates with GreenOrbs: Sustainable Sensing in the Forest," in *ACM SenSys*, 2009.
- [5] J. Polastre, R. Szewczyk, and D. Culler, "Telos: Enabling Ultra-low Power Wireless Research," in *ACM/IEEE IPSN*, 2005.
- [6] O. Gnawali, R. Fonseca, K. Jamieson, D. Moss, and P. Levis, "Collection Tree Protocol," in *ACM SenSys*, 2009.