Wireless Sensor Networking over White Spaces

Abusayeed Saifullah, Chenyang Lu
Washington University
St. Louis, MO
{saifullah,lu}@wustl.edu

Jie Liu, Ranveer Chandra, Sriram Sankar
Microsoft Research
Redmond, WA
{liuj,ranveer,sriram.sankar}@microsoft.com

ABSTRACT

White spaces represent unoccupied TV spectrum that could be used for wireless communication by unlicensed devices, just like Wi-Fi and ZigBee devices can use the ISM bands. This low-frequency spectrum enables high bandwidth communication at long transmission ranges. In this paper, we propose a system that leverages the benefits of white spaces for sensor networks, with a view to overcoming the limitations associated with coverage, bandwidth, time synchronization, and scalability of current wireless sensor networks (WSNs). However, the applications and requirements of WSNs pose interesting challenges for adopting white spaces. In this challenge paper, we address both the opportunities and the challenges of using white spaces for WSNs.

1. INTRODUCTION

Emerging large-scale wireless sensing and control systems will require many sensors connected over long distances. Existing wireless sensor network (WSN) technology such as IEEE 802.15.4 has short range and low bandwidth that pose a significant limitation in meeting this impending demand. In this paper, we propose a system to meet this demand by designing sensor networks to operate over the TV *white spaces*, which refer to the allocated, but unused TV channels.

In a historic ruling in 2008, the FCC allowed unlicensed devices, such as Wi-Fi devices, to operate in unoccupied TV channels [1]. To learn about unoccupied channels at a location a device needs to either (i) sense the medium before transmitting, or (ii) consult with a cloud-hosted geo-location database [2], either periodically or every time the device moves 100 meters. Similar regulations are being adopted in Canada, Singapore, UK, and many other countries worldwide.

Since the TV transmissions are in the lower frequencies – VHF, and lower UHF (470 to 698 MHz) – white space transmissions have excellent propagation characteristics over long distance. They can easily penetrate walls and other objects, and hence hold enormous potential for wireless applications that need substantial bandwidth and long transmission range.

This potential of white spaces is mostly being tapped into for wireless broadband access. Industry leaders, such as Microsoft and Google, are exploring ways to leverage its properties for broadband access in underserved communities [3, 4]. Various standard bodies are modifying existing standards and also developing hardware platforms to accommodate this paradigm shift – of opportunistically using unoccupied spectrum, and consulting with a cloud service before using the medium. For example, the recently standardized IEEE 802.11af [5, 12] details techniques to enable Wi-Fi transmissions in the TV white spaces. IEEE 802.22 [6] working group intends to provide wireless broadband access over white spaces. IEEE 802.19 [7] is aimed at enabling effective use of white spaces by the family of IEEE 802 standards. Recently, Neul, Inc. has proposed the Weightless standard to cater for the FCC requirements [8]. Their preliminary study has demonstrated the feasibility of using white spaces for machine to machine communication within several kilometers with a battery life of several years. In parallel, the research community has been investigating techniques to make spectrum sensing low-cost [27], and more accurate [17, 18].

In this paper, we propose a new application over this spectrum, called Sensor Network over the White Spaces (SNOW). Such networks can help overcome limitations of existing technologies associated with communication range, bandwidth, time synchronization, and scalability. In conjunction with existing sensor network technologies, SNOW can accomplish pervasive connected sensors over large geographic areas. Furthermore, white spaces will enable large-scale wireless control systems that rely on real-time communication over wireless sensor-actuator networks. However, realizing SNOW requires several challenges to be solved. In this paper, we will first identify the tremendous opportunities white spaces bring to WSNs, followed by a discussion of the key research challenges that need to be addressed in order to realize the vision of sensor networks over white spaces.

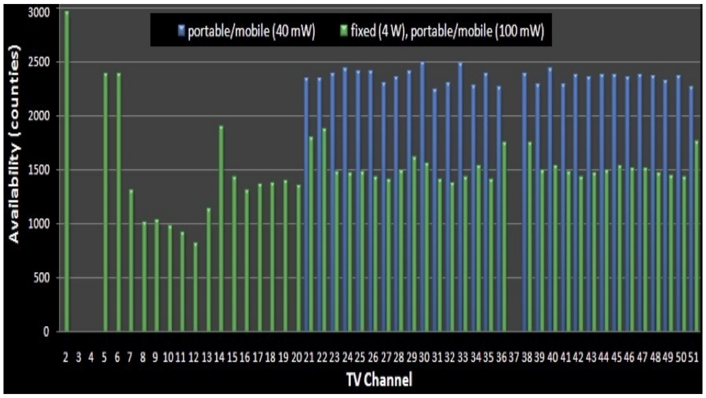

Figure 1: White space channel availability in US

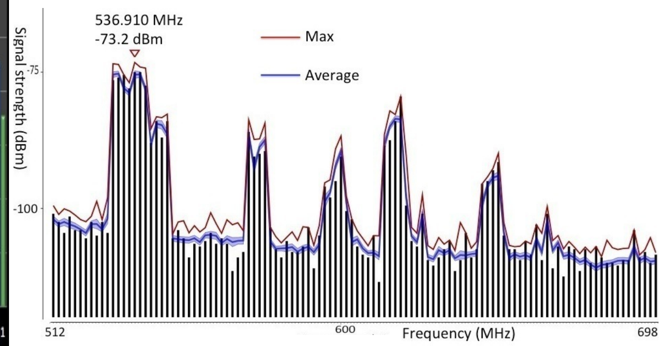

Figure 2: Spectrum (512-698MHz) inside a data center

2. WHITE SPACES CHARACTERISTICS FOR SENSOR NETWORKS

In this section, we discuss some physical characteristics of white space spectrum that can be exploited by a large class of WSN applications to overcome the traditional problems associated with capacity, range, and scalability.

2.1 Availability

Compared to IEEE 802.15.4 or Wi-Fi, white spaces offer a large number channels. For example, the spectrum between 512MHz and 698MHz has 30 channels (each TV channel is 6 MHz wide). The availability of white spaces depends on location. Rural (and suburban) areas, typically have more white spaces than urban areas due to lesser number of TV stations. Figure 1 shows county based availability of white space spectrum per channel from channel 2 to 51 (except 37 that is not allowed for usage) in USA, based on the statistics collected by Spectrum Bridge [9]. For a channel in x-axis, the corresponding value on y-axis indicates the number of counties among 3142 counties nationwide where the channel is available as white space. Many wireless sensor network deployments such as those for habitat monitoring [25], environmental monitoring [20], and volcano monitoring [26] are in rural areas, making them perfect users of white spaces.

White spaces also exist in many urban areas. For example, given the large area covered by today's data centers, it has become challenging to support real-time data center management (e.g., data center-wide power management) over traditional wireless sensor networks such as IEEE 802.15.4. To explore the availability of white space in data centers located in urban areas, we measured the spectrum availability in a data center (of dimension $50m \times 50m$) located in a metropolitan city. The Spectrum Bridge website showed that 13 channels were available at this location. We went a step further and performed measurements using a spectrum analyzer with a low noise floor. We represent the power density graph of spectrum between 512 MHz and 698 MHz when measured in the data center, in Figure 2. As we see in the Figure, large parts of the TV spectrum is left unused even in some urban areas.

2.2 Long Transmission Range

Due to the short communication range of the IEEE 802.15.4 devices (20-30m at 0 dBm), most WSNs form multi-hop mesh networks, making the protocol design and network deployment quite complicated. In contrast, due to lower frequency, white space radios have very long communication range. Previous experiments using software radio [12] and preliminary industrial study with the ongoing standardized devices [8] have shown the communication range to be of several kilometers. Long transmission range can overcome several inherent challenges faced by current WSN technologies such as time synchronization and real-time communication in wide-area deployments.

Reduced Time Synchronization Overhead. Time synchronization is a critical requirement in many WSN applications. However, accurate time synchronization in wireless mesh networks poses a serious overhead, especially in large-scale deployments. For example, FTSP [23], a widely used time synchronization protocol for WSNs, face scalability challenges due to its reliance on flooding. PulseSync [21] is another well-known protocol which does not use flooding but still incurs exponential worst-case error with hop count. The difficulty in time synchronization in multi-hop WSNs also motivated multiple out-of-band techniques that rely on separate radios or GPS, but these solutions require extra hardware. Replacing current WSNs with SNOWs will reduce many multi-hop mesh topologies to star topologies, or will reduce the number of hops significantly. This will reduce time synchronization overhead, and enhance scalability.

Reduced Latency. Reduced hop counts will in turn result in shorter communication latency making SNOW much more suitable for large real-time control applications than existing WSN technologies. Consider the following analytical results as an example. The lower bound latency for convergecast in a WSN of n nodes of linear topology is $(3n-3)$ time slots [15], whereas in a single hop SNOW the upper bound of this latency is n slots, thereby reducing the latency by a factor of 3 in SNOW. If n_k represents the maximum number of nodes in a subtree, then the lower bound latency for convergecast in tree networks is $\max(3n_k-3,n)$ slots [15], whereas in a single hop SNOW the upper bound of this latency is n slots.

2.3 Wall Penetration

One of the primary benefits of the white space is the improved propagation compared to 2.4 and 5 GHz band. Wavelengths in this band are quite large and are in the range between $43cm$ and $59cm$. They can therefore easily penetrate walls, unlike 802.15.4 or 802.11 transmissions, allowing better non-line-of-sight connectivity. SNOW will thus be suitable for many environments that are challenging for current WSN technologies, e.g., those requiring wall penetration for radio through office buildings. Obstacle penetration is a severe problem in industrial environment where moving objects often block wireless communication. WirelessHART networks [10] handles this through a high degree of redundancy where a packet is transmitted multiple times and through multiple paths hindering scalability. Transmissions over the white spaces can cover the entire network with negligible signal decay through walls and obstacles. The lower frequency signals are less susceptible to human presence, multi-path and fading, and exhibit exceptional indoor penetration, and vary less over time. Hence, they hold potential for better indoor applications and localization [14].

2.4 Higher Bandwidth

It is challenging to transmit and receive large amounts of data over current WSNs that typically have low bandwidth (i.e., 250Kbps maximum for IEEE 802.15.4). In many WSN application such as volcano monitoring [26], where data from seismometers and microphones must be recorded at relatively high rates (100Hz sampling rate) with adequate per-sample resolution (24 bits), it is infeasible to continuously transmit the full resolution signal through a 802.15.4 network. A white space channel can offer a throughput of up to 22Mbps, which can be even further increased by channel bonding [12]. Higher bandwidth in SNOW will enable data intensive applications such as volcano monitoring [26], structural health monitoring [19], and camera sensor network to deliver data in real-time, which is extremely challenging for existing WSNs.

3. BENEFITTING APPLICATIONS

SNOW will be a natural fit for wide-area sensing applications where current WSN technologies have faced significant challenge. More importantly, it will enable an emerging class of large-scale wireless control applications that require real-time communication over wireless sensor-actuator networks.

3.1 Wide-Area Applications

Given its long communication range, SNOW will greatly simplify wide-area sensing applications that must collect data from sensors spread over a large geographic area or distance. With current short-range wireless technology such as 802.15.4 and 802.11, wide-area sensor networks must form many-hop mesh networks and address significant engineering challenges in scalability. For example, the WSN deployment on Golden Gate Bridge forms a 46-hop network to cover 1280m main span of the bridge [19]. Due to this large number of hops and relatively large amounts of data, it requires approximately 10 hours to collect all sensor data from the bridge. Another example is ZebraNet deployed in Mpala Research Center for tracking zebras [16]. Because zebras are free to move in an area of approx. 200,000m^2, the network lacks continuous connectivity due to short radio communication range, and is managed through a delay-tolerant network which cannot deliver information in real time. All these WSNs can be covered in a single-hop white space network, thereby simplifying network design, deployment, and protocol design. SNOW may also require fewer nodes because it may not require any additional nodes to serve as relays.

3.2 Real-Time Applications

White spaces provide ample opportunities to enable large wireless cyber-physical systems involving feedback control loops that rely on real-time communication between sensors, controllers and actuators over wireless networks (e.g. process control [24] and civil structure control [22]). This new class of *real-time* applications represents a new paradigm of WSNs in sharp contrast to the traditional WSNs where best effort communication usually suffices.

For example, a current wireless technology for industrial process management is WirelessHART [10]. With the growing applications in industrial process management, these industrial networks now need to scale up to tens of thousands of nodes [11]. Based on existing guidelines that recommend the installation of a maximum of 80 nodes per gateway [10], a single WSN requires significant wiring to integrate numerous small networks in a large deployment, significantly reducing the benefit of wireless. A WirelessHART network also needs global time synchronization. The network is centralized, where a central network manager periodically monitors the health of all nodes and links. It then creates and distributes the schedules among individual nodes when there is a significant change in channel condition and link failure. The scalability and real-time performance of such networks is inherently limited by their multi-hop and time-varying topologies. In contrast, white spaces can cover an entire WirelessHART network in a single hop or a small hop count, simplifying management and protocol design.

Due to reduced number of hops in SNOW, real-time protocol design becomes simpler compared to that for multi-hop WSNs. For example, for a single-hop SNOW, the real-time schedulability condition can be developed by extending the traditional uniprocessor real-time scheduling analysis [13], and can be used for network planning and management. Namely, each periodic communication between the BS and a sensor node can be considered as a real-time task. When the Earliest Deadline First (EDF) policy is adopted for scheduling them, the task with the earliest absolute deadline has highest priority. If all tasks are periodic, preemptive, and have deadlines equal to their periods, preemptive EDF is able to achieve 100% network utilization [13]. But transmission scheduling has non-preemptive nature meaning that if a transmission starts, it cannot be stopped (preempted). Therefore, we can adopt a simple TDMA protocol using a non-preemptive EDF scheduling that approaches 100% network utilization. Consider ℓ_i as the packet length of task i with rate r_i, and ω as the data rate of the channel used. A set of n tasks for a single-hop SNOW under non-preemptive EDF [13] is schedulable (can meet their deadlines) if

$$\sum_{i=1}^{n} \frac{\ell_i}{\omega} r_i + b_j r_j \leq 1, \quad \forall 1 \leq j \leq n \qquad (1)$$

where b_i is the maximum blocking time of task i due to non-preemptivity of data transmission, and is given by

$$b_i = \max\{\frac{\ell_j}{\omega} | j \neq i\} \qquad (2)$$

4. CHALLENGES AND RESEARCH OPPORTUNITIES

Taking advantage of white spaces for sensor networking does not come for free. FCC poses strict rules on how the white spaces can be used, yet they are primarily designed for Internet connected, data access scenarios. Large communication ranges also bring forth challenges in energy efficiency, resource contention and network management. In this section, we outline some key challenges and future research directions to tackle them.

4.1 Energy Efficiency

Energy efficiency is a primary concern in sensor networks, especially for battery powered or energy harvesting deployments. Lessons learned from radios operating on ISM bands may not be immediately applicable to SNOW.

Packet Sizes. Given the long communication range, the energy spent on communicating a unit of data (e.g. measured in J/bit) over white spaces may be much lower than over multiple hops in ZigBee and WiFi. However, the network must be carefully designed to achieve high energy efficiency. For example, in IEEE 802.15.4, packet sizes are defined to be small, of maximum size of 128 bytes. This is because the bit error rate is high over low power wireless. Since communication over the TV white spaces are more reliable over shorter distances, it is likely beneficial to use longer packets when the devices are close by (high SNR). Using larger packet sizes will reduce the per-packet overhead, such as preambles and headers, and hence reduce the energy consumed per useful bit of data.

Media Access Control Protocols. To achieve high energy efficiency, many sensor networks try to reduce the idle listening time for wireless nodes, thus employing techniques, such as low power listening and receiver initiated MAC. However, both cases require one side of the link to send extremely long preambles that is at least half the length of its counterpart's sleeping time. With large communication range, a great number of nodes can be in the one-hop communication range. Blindly applying existing sensor network MAC designs will cause all those nodes to wake up unintentionally.

When time synchronization becomes cheap in SNOW over a single hop, it becomes natural to consider TDMA-like MAC protocols to achieve better energy efficiency by avoiding carrier sensing and collisions. Another direction is to exploit the potential asymmetry in energy sources. For example, if the base station is connected to a power source, then it can be more liberal in spending energy to manage the network, and keep the activity level of battery powered nodes low.

Antenna Design. On the physical form factor perspective, antennas that operate in white spaces are usually large, due to the longer wave length in TV bands. For sensor nodes that have size or weight constraints, these antennas become cumbersome. Smaller antenna are typically less sensitive, which can make the devices less energy efficient. In addition to seeking better materials and better geometry for antenna design, multi-antenna designs that can achieve high gain and directionality at the base station nodes, can greatly boost the energy efficiency at remote nodes.

4.2 SNOW Applicability

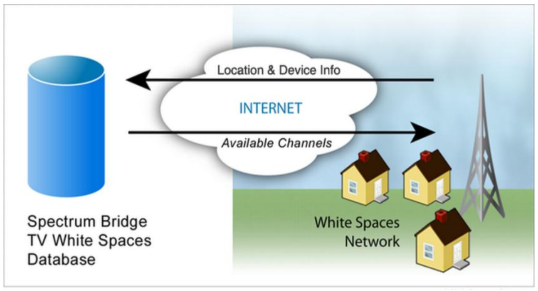

Figure 3: Accessing white space

Long communication range gives SNOW more flexibility for network topology control to achieve high efficiency and low interference, which was not practical for IEEE 802.15.4-based sensor networks. Long range can be an overkill for some applications where short range is enough, and is prone to interference from other networks that share the spectrum. When provision a SNOW, it is important to tradeoff between infrastructure nodes and commutation range/power. As a preliminary idea, we propose that two parameters need to be considered for choosing applications for SNOW: communication pattern and deployment density. Considering this, a number of applications may not be well-suited for white space. For example, systems that require local peer to peer communication and in-network processing; systems that employ tight local control loops and long-range interference is undesirable; and systems like body-area sensor network. In these applications, at least from the network perspective, using white spaces will cause unnecessary interference. Hybrid sensor networks architecture spanning over 2.4 GHz and the white spaces may achieve better overall performance.

Deployments for real-time control applications over SNOW is also of great interest. Although the single hop topology brings theoretical simplicity and better analytical results for these network, practical solutions may not be straightforward. When a system spans over large a geographical space and many nodes, it is important to differentiate the quality of services that the network can provide to the applications. One may borrow ideas from Control-Area Network (CAN) to support feedback loops with stringent deadline requirements.

4.3 Spectrum Management

Primary User Detection. Primary users can be detected either through white space database access (Figure 3) or using spectrum sensing. In case of the latter, the channel must be sensed before every packet is sent. The first option requires that the nodes know their physical locations, and they can reach, possibly through the base station, the central database. Knowing nodes' physical location introduces extra system provisioning overhead, especially when the nodes do not have GPS or, in many indoor environments, when GPS receivers do not work.

For untethered sensor networks that do not have constant access to the Internet (e.g. through periodic data mule), spectrum sensing becomes the only option. How often to perform spectrum sensing and how to share the duty among sensor nodes to sense spectrum availability to reach the consensus on available channels is an interesting research question.

Handling Temporal Variation. Primary sources such as wireless microphones can become active at any time without warning, and the white space devices (that employ sensing) using those channels must disconnect immediately upon detecting any microphone usage and then rapidly reconnect using a different available channel. Such disruption can have effects on real-time applications such as critical process monitoring or security monitoring. Therefore, new protocols may be developed for SNOW to ensure that the temporal disruptions do not affect the QoS of the underlying applications, or the nodes can proactively occupy more than one channels to ensure a smooth transition.

Node Joining. Another important challenge lies in handling a node joining process. When a node needs to join the network, it needs to send its location to the database server node (which is either the BS or another node). However, to avoid interference with primaries, the client cannot transmit its location to the database server node unless it knows the available white spaces at its own location. A possible solution can be to use periodic beacons by the database server node containing information of the available channels of all locations in its coverage. But this is a long and costly process, and improved methods must be developed.

Handling Mobility. Handling mobility in SNOW raises a number of challenges. According to the FCC regulations, fixed and portable devices have different requirements among which is the separation distance from digital and analog TV protected contours. In particular, fixed devices are not allowed to operate on first adjacent channels to a TV station. On the other hand, portable devices are allowed to operate on first adjacent channels subject to lower maximum transmit power constraints. This translates to different white space availability for fixed and portable devices (Figure 1). Spectrum availability also changes due to location change. There are multiple scenarios. Both a sensor node and its corresponding database server node can move. On the other hand, a sensor node can be mobile while its database server node does not move. In either case, the sensor node needs to periodically send its new location. The period is a function of their moving speed. In the second case, when only a sensor node is mobile, it needs to discover a new database server node based on a node joining protocol. In both cases, a safety margin on the communication range needs to be maintained based on their speed to avoid false positive channels due to mobility. Considering all of these issues, developing efficient WSN protocols to handle mobility opens a new area of research.

The above discussions assume that SNOW sticks with current rules that are primarily designed for broadband applications. Regulatory changes may potentially simplify this problem. For example, for sensing applications that only send data sporadically and in small quantities, their interference to primary users may be negligible. Is it possible to allow such transmissions at low power to co-exist with primary users? A deeper study of the interference between sensor network-like traffic and digital TV transmission can reveal this possibility.

5. CONCLUSIONS

Ubiquitous white spaces offer a new paradigm for wireless sensor networking to overcome the limitations on scalability and coverage. This paper has proposed Sensor Networking Over White space (SNOW) with the vision to supersede current wireless sensor networks and to unleash efficient WSN applications in cloud computing, cyber-physical systems, real-time applications, and remote monitoring. We present advantages, opportunities, and challenges in adopting SNOW with some potential proposals to overcome the challenges. We have opened a new door of research that will have tremendous impacts on the future of wireless sensing and control.

6. REFERENCES

[1] FCC, ET Docket No FCC 08-260, November 2008.

[2] FCC, Second Memorandum Opinion and Order, ET Docket No FCC 10-174, September 2010.

[3] `http://www.microsoft.com/africa/4afrika/`.

[4] `https://sites.google.com/site/tvwsafrica2013/`.

[5] `http://www.radio-electronics.com/info/wireless/wi-fi/ieee-802-11af-white-fi-tv-space.php`.

[6] `http://www.ieee802.org/22/`.

[7] `http://www.ieee802.org/19/`.

[8] `http://www.weightless.org/`.

[9] `http://spectrumbridge.com/Libraries/White_Papers/TV_WhiteSpaces_Usage_Availability_Analysis.sflb.ashx`.

[10] `http://www.hartcomm2.org`.

[11] WirelessHART system engineering guide. `http://www2.emersonprocess.com/siteadmincenter/PM%20Central%20Web%20Documents/EMR_WirelessHART_SysEngGuide.pdf`.

[12] P. Bahl, R. Chandra, T. Moscibroda, R. Murty, and M. Welsh. White space networking with wi-fi like connectivity. In *SIGCOMM '09*.

[13] G. C. Butazzo. *Hard Real-Time Computing Systems*. Springer, 2005. 2nd edition.

[14] Y. Chen, D. Lymberopoulos, J. Liu, and B. Priyantha. Fm-based indoor localization. In *MobiSys '12*.

[15] S. Gandham, Y. Zhang, and Q. Huang. Distributed time-optimal scheduling for convergecast in wireless sensor networks. *Comput. Netw.*, 52(3):610–629, 2008.

[16] P. Juang, H. Oki, Y. Wang, M. Martonosi, L. S. Peh, and D. Rubenstein. Energy-efficient computing for wildlife tracking: design tradeoffs and early experiences with zebranet. In *ASPLOS-X '02*.

[17] H. Kim and K. G. Shin. Fast discovery of spectrum opportunities in cognitive radio networks. In *DySpan '08*.

[18] H. Kim and K. G. Shin. In-band spectrum sensing in cognitive radio networks: Energy detection or feature detection? In *MobiCom '08*.

[19] S. Kim, S. Pakzad, D. Culler, J. Demmel, G. Fenves, S. Glaser, and M. Turon. Health monitoring of civil infrastructures using wireless sensor networks. In *IPSN '07*.

[20] K. Langendoen, A. Baggio, and O. Visser. Murphy loves potatoes: experiences from a pilot sensor network deployment in precision agriculture. In *IPDPS '06*.

[21] C. Lenzen, P. Sommer, and R. Wattenhofer. Optimal clock synchronization in networks. In *SenSys '09*.

[22] B. Li, Z. Sun, K. Mechitov, C. Lu, D. Dyke, G. Agha, and B. Spencer. Realistic case studies of wireless structural control. In *ICCPS'13*.

[23] M. Maróti, B. Kusy, G. Simon, and A. Lédeczi. The flooding time synchronization protocol. In *SenSys '04*.

[24] A. Saifullah, C. Wu, P. Tiwari, Y. Xu, Y. Fu, C. Lu, and Y. Chen. Near optimal rate selection for wireless control systems. In *RTAS '12*.

[25] R. Szewczyk, A. Mainwaring, J. Polastre, J. Anderson, and D. Culler. An analysis of a large scale habitat monitoring application. In *SenSys '04*.

[26] G. Werner-Allen, K. Lorincz, J. Johnson, J. Lees, and M. Welsh. Fidelity and yield in a volcano monitoring sensor network. In *OSDI '06*.

[27] X. Ying, J. Zhang, L. Yan, G. Zhang, M. Chen, and R. Chandra. Exploring indoor white spaces in metropolises. In *MobiCom '13*.