Understanding the Link-level Behaviors of A Large Scale Urban Sensor Network

Jiliang Wang* and Wei Dong†
*School of Software and TNLIST, Tsinghua University
†College of Computer Science, Zhejiang University
jiliangwang@tsinghua.edu.cn, dongw@zju.edu.cn

***Abstract*—We present the first comprehensive link-level measurements in an operational large-scale urban sensor network. By carefully analyzing the performance metrics, we seek to answer several fundamental questions: what are the characteristics of links in a real large-scale network, and what causes link performance degradation? The key findings of this study are that (1) the performance of intermediate links is the most unpredictable and some links exhibit highly periodic patterns, (2) the width of the reception "transitional region" is much larger than those reported in previous experiments, indicating that an outdoor environment might have a larger impact on the link performance and current protocol parameters should be carefully designed, and (3) different from previously reported results, link performance degradation has a relatively weak correlation with the corresponding RSSI (Received Signal Strength Indicator) values fluctuating near the noise floor.**

I. Introduction

As an emerging technology that bridges cyber systems and the physical world, wireless sensor networks (WSNs) are envisioned to support numerous applications that are unthinkable before [1], [2], [3]. Today, technological advances already make it possible to apply WSNs in many areas. Many real-world deployments are witnessed: VigilNet [1] includes 200 nodes to support military surveillance; ExScal [4] consists of over 1000 sensor nodes and 200 backbone nodes; SensorScope [2], a real-world WSN system on rock glacier, consists of nearly 100 sensor nodes.

We are still facing severe challenges, however, in designing scalable, long-lived, and high-performance WSN systems. Some of the difficulties come from the fact the current understanding of WSN system has some limitations. Theoretical models can hardly be applied in practical systems with complex interactions among different components [5], [6]. Therefore, it is necessary to conduct empirical measurements in real-world WSN systems, so that we can better under the behavior of a large scale WSN and facilitate effective and efficient design of such systems.

Indeed, there are many relevant measurement studies in the past years. Aguayo et al. [7] present a link-level measurement study on an 802.11b mesh network. Zhao et al. [8] report a measurement study on packet delivery performance using 60 Mica nodes. Srinivasan et al. [9] measure packet delivery performance of the Telos and Micaz platforms. Natarajan et al. [10] measure and analyze the link layer behavior of a body area network at 2.4GHz. Dawson-Haggerty et al. [11] analyze the effect of link churn on wireless routing. Clearly, the above measurement studies target at a certain subset of networking metrics and do not provide a systematic study of operational WSNs, and they all suffer from short term and relatively small scale.

It is well-known that the measurement study on the Internet is an important topic since its birth. Numerous efforts are paid for characterizing a network with respect to different metrics, testing or validating certain network protocols and software, monitoring or evaluating network performance, and diagnosing network anomalies [12], [13], [14], [15], [16], [17].

Measurements in WSN systems, on the other hand, are very different from those in Internet, with different intrinsic difficulties. On one hand, the Internet has evolved to a relatively mature system for development and measurement, while there is few, if not none, deployed working WSN system for measurement. On the other hand, Internet today already consists of both computers and users, while WSNs mainly consist of sensor nodes, interacting with outside objects like unpredictable environment dynamics.

To address these challenges, we have deployed large-scale WSN systems to explore the intrinsic properties. Previously, we have reported the measurements on a WSN with 330 nodes, and identified three important factors that may interfere with scalable WSNs [18], namely bottlenecked nodes, inherent concurrency of network operations, as well as environmental dynamics. Although the previous study gives us important understanding for designing scalable WSNs, it does not deeply investigate the link-level behaviors and thus fails to give detailed guidelines to protocol designs. Motivated by this need, we conduct a link-level measurement in a large-scale urban sensor network system which consists of up to 1,200 sensor nodes, providing us unprecedented opportunities to observe system behaviors at scale. Our goal is to better understand the link-level behaviors which form the foundation of packet delivery performance in practical large scale network. We are interested in the following questions: What are the temporal-spatial characteristics of links, and what causes link performance degradation?

Understanding the above issues would have broad implications for many aspects of sensor networking, such as carrier sensing, link estimation, wireless routing, and etc. The major findings of this study are as follows:

1) The performance of intermediate links is most unpre-

dictable and some links exhibit highly periodic patterns.
2) The width of the reception "transitional region" is much larger than those previously reported in indoor experiments [11], [9], indicating that the environment might have a large impact on the link performance and protocol parameters should be carefully designed.
3) Differing from previously reported results, link performance degradation has a weak correlation with the corresponding RSSI (Received Signal Strength Indicator) values fluctuating near the noise floor.

The rest of this paper is structured as follows. Section II discusses related work. Section III introduces the measurement platform and data sources. Section IV presents the link-level performance, and temporal and spatial characteristics of all visible links. Section V explores the root causes of link loss. Finally, we conclude this study in Section VI.

II. Related Work

An operational networking system can be viewed as the integration of two parts: the inherent network infrastructure (gateway, router, switch, cable, protocol stack, etc.) and miscellaneous network traffic (data packets and control signals) running on it. While the network infrastructure is usually prebuilt as a relatively fixed component, the network traffic is often out of full control of network administrators, due to the distributed and autonomous nature of network entities. Network measurement is therefore a significant issue since the birth of networking systems and carries many functions, such as characterizing a network with respect to different metrics, testing or validating certain network protocols and software, monitoring or evaluating network performance, and diagnosing network anomalies [12], [13], [14], [15].

Measurement of the Internet is one of the main flows in network measurement studies. As the Internet is popular worldwide nowadays, the research focus of Internet measurements ranges from the microscopic experimental study on lower level of the protocol stack to macroscopic understanding of the networks in various aspects.

Turner et al. [12] present a methodology for understanding the causes and impact of link failures. They opportunistically mine data sources that are already available in modern network environments and analyze failure events of over five years in a large regional network. Wang et al. [13] conduct a measurement study on the impact of routing events to the end-to-end path performance. They show that end-to-end Internet path performance degradation is correlated with routing dynamics and analyze the root cause of the correlation between routing dynamics and such performance degradation. Mahimkar et al. [14] study the impact of upgrades on network key performance indicators in a large operational network. Gill et al. [15] study network failures in data centers. They find that data centers are mostly reliable and network redundancy is only 40% effective in reducing the impact of failure.

The openness of the Internet presents sufficient work space for measurements. Researchers deploy powerful agents to sniffer the network, utilize technically-mature tools to conduct targeted measurements, and even collect huge volumes of data traces for fine-grained offline analysis. In contrast, the existing studies on WSN measurements are by nature restricted due to the limited scale of real deployments.

A WSN for habitat monitoring [19] is deployed at Great Duck Island in 2002. Tolle et al. [20] report a sensor network consisting of 33 nodes to monitor the microclimate of a redwood tree, covering an area of about 50 square meters. Werner-Allen et al. [21] deploy a WSN of 16 nodes to monitor an active volcano. Experiences from those systems provide basic findings that act as the early guidance to design the building blocks of modern WSNs. Due to the limited system scales and deployment periods, those efforts fail to capture the fundamental challenges in large-scale WSN construction, nor can they experience the real stumbling blocks towards a sustainable WSN.

VigilNet [1] employs 200 sensor nodes to support military surveillance, covering an area of 100×100 square meters. ExScal [4], [22] used to attempt to deploy a WSN with over 4,000 nodes but the system does not run for a long time. SensorScope [2] is a real-world WSN system, in which the largest deployment consists of 97 nodes. Those WSN deployments are originally designed for a specific application, without many efforts on measurement studies. In words, there is still a huge gap between the disclosed facts and the need of understanding practical WSNs.

The work in [18] for the first time reports the measurements on a 330-node outdoor sensor network, GreenOrbs. The study identifies three important factors that may interfere with scalable WSNs, namely bottlenecked nodes, inherent concurrency of network operations, and environmental dynamics.

In order to get more insights, there have been many existing works with particular interest on a certain aspect of WSNs. Aguayo et al. [7] present a link-level measurement study on an 802.11b mesh network. Subramanian et al. [23] present a measurement study on channel and interface heterogeneity in wireless mesh networks. Zhao et al. [8] carry out measurements on packet delivery performance using 60 Mica nodes. Maheshwasi et al. [24] carry out measurements on interference modeling and scheduling with two 20-node TelosB testbeds. Srinivasan et al. [9] present measurements of packet delivery performance of the Telos and Micaz platforms. Natarajan et al. [10] measure and analyze the link layer behavior of a body area network at 2.4GHz. Dawson-Haggerty et al. [11] analyze the effect of link churn on wireless routing. They validate the hairy edge hypothesis under practical routing protocols, i.e. whether the most important links for routing are the most difficult to predict.

III. Background and Data Sources

A. Application Background

"The world has just ten years to bring greenhouse gas emissions under control before the damage they cause becomes irreversible." Global climate change causes a series of damage to the earth environment, such as global warming, glacier melting, sea level rise, and extreme weather events. It poses

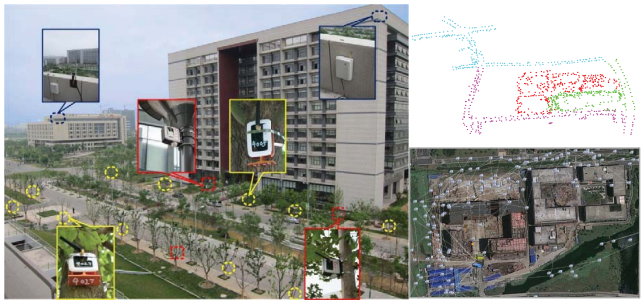

Fig. 1: The network deployment

increasingly severe threats to the subsistence of human beings and the sustainable development of human society.

People have taken various efforts to combat global climate change. Environmental research reveals that the increasing emission of greenhouse gases, of which carbon dioxide is a major component, is the main culprit of global climate change. According to the statistical results, forest sequestrates 77% of the total carbon stocks in the terrestrial ecosystem, while CO_2 emissions from cities account for 75% of the total emissions of human activities. Correspondingly, people realize that there are two aspects of countermeasures. One is to reduce the emission of carbon dioxide by advocating low-carbon economy, and the other is to enhance the sequestration of carbon dioxide by developing forest and urban greening.

The goal of our project, CitySee, is to deploy thousands of wireless sensor nodes in an urban area of Wuxi City, China, such that multi-dimensional data including CO_2, temperature, humidity, illumination, location, and etc., could be collected in a real-time manner for further analysis. The need of the long-term, large-scale, continuous, and synchronized surveillance of urban areas poses new challenges to the design of a WSN system [25] [26] [27]. Thus we exercise a series of innovative designs with our system.

Figure 1 plots the current deployment. The left subfigure shows a real deployment scene, where we can see all the three types of nodes deployed at different positions. The upper right subfigure shows the overall node distribution, where every dot represents a node and different subnets are differentiated in different colors. The lower right subfigure shows the first subnet of CitySee on the real satellite map, where a snapshot of routing paths in the subnet is included as well.

B. Data Sources

The measurement is carried out mainly using the data from normal nodes. Four types of data are collected, defined as packet types C1, C2, C3, and C4, respectively.

Type C1. It includes two kinds of information: (1) sensor readings, such as CO_2, temperature, humidity, light, and battery voltage; (2) routing metrics: path-ETX [28] from the original node to the sink node and node IDs along a routing path (the maximum length is set to 10).

Type C2. It contains the routing table with a pre-configured maximum size of 10. Each routing entry includes the neighbor node ID, the RSSI value from the neighbor, the link-ETX, and the path-ETX. We denote them as nb[i], RSSI[i], linkETX[i], and pathETX[i] ($0 \leq i < 10$).

Type C3. It contains different counters: (1) the CPU counter records the accumulated task execution time; (2) the radio counter records the accumulated radio-on-time in millisecond; (3) the transmission counter records the accumulated number of transmitted packets; (4) the reception counter records the accumulated number of received packets; (5) the drop_no_ack counter records the accumulated times of packet drops because of exceeding the retransmission threshold (set at 30 in CTP); (6) the drop_overflow counter records the accumulated times of packet drops due to queue overflow; (7) the loop counter records the accumulated number of detected routing loops; (8) the MACI counter records the accumulated number of the MAC initial backoffs; (9) the MACC counter records the accumulated number of MAC congestion backoffs.

Type C4. It contains other networking information about the current packet: (1) a sequence of node IDs along the routing path; (2) queue length at each forwarding node upon the packet's arrival; (3) the retransmission number at each forwarding node.

The four types of packets share a common routing header, which is formatted to include the following fields: (1) origin - the packet origin; (2) SEQ - sequence number of CTP [29] forwarding; (3) THL - the hop count of the incoming packet; (4) SOURCE_TIME - the packet transmission time in the local clock of the original node; (5) SINK_TIME - the packet reception time in the local clock of the sink.

Each sensor node transmits four kinds of packets every 10 min. We conduct our experiments using all data collected in four months.

IV. Link Characteristics

In this section, we analyze the packet trace to explore the link performance, and its temporal and spatial characteristics.

A. Methodology

We mainly utilize type C2 packets to extract the link statistics. A C2 packet contains information of up to ten links. linkETX[i] ($0 \leq i < 10$) denotes the ETX of the link origin $\rightarrow$ nb[i]. RSSI[i] denotes the RSSI from nb[i] to origin upon a beacon's reception.

Our current protocol uses the 4 bit link estimation method [30] to estimate the link quality. The 4 bit link estimation combines both control plane information (i.e. using beacons) and data plane information (i.e. using data packets). For non-routing links (i.e. nb[i] is not the parent), since there is no data traffic and there is no exchange of the link quality in the current implementation, we use linkETX as the incoming link quality from nb to origin. For routing links (i.e. nb[i] is the current parent), we use linkETX as the outgoing link quality from origin to nb[i].

We extract link statistics from the packet trace. We then obtain time series of all visible links. Previous wireless measurements explore the set of all possible links while in this

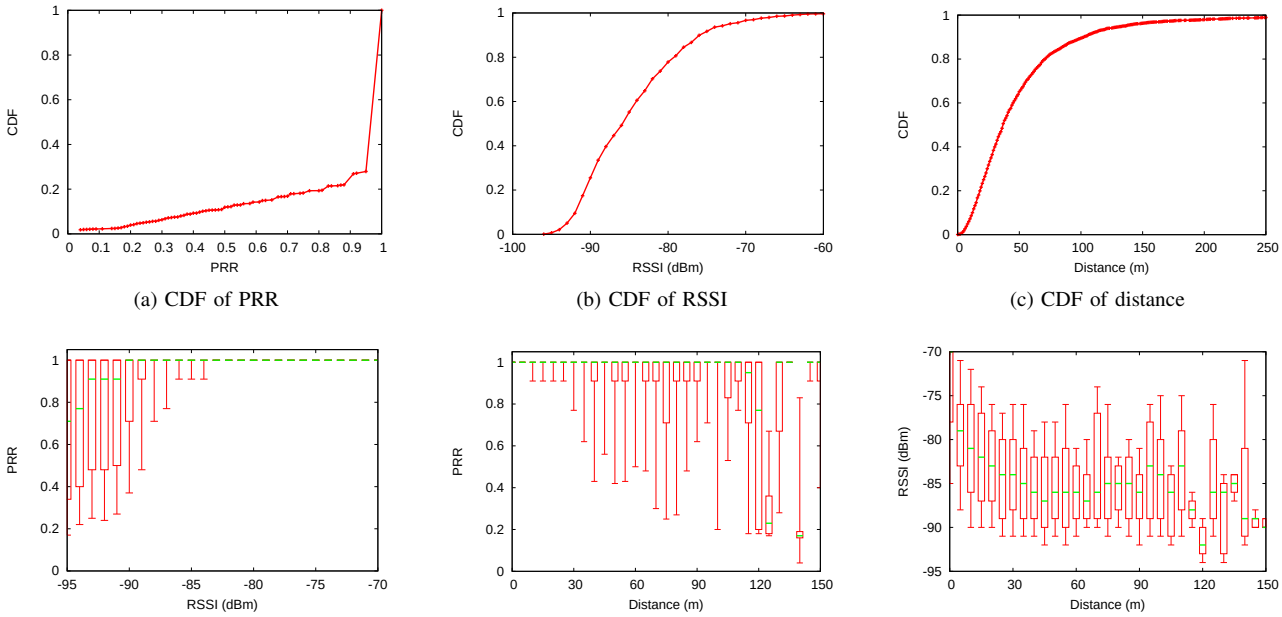

(a) CDF of PRR (b) CDF of RSSI (c) CDF of distance

(d) RSSI vs PRR (e) Distance vs PRR (f) Distance vs RSSI

Fig. 2: Link performance

study we consider all visible links. We reduce the complexity of recording all possible links while preserving the set of most important links (since they are candidates for routing).

In particular, a visible link has a time series about its PRR (1/linkETX), i.e., link a→b: PRR(t_1), ..., PRR(t_n), where t_i ($1 \leq i \leq n$) denotes the timestamp. It also has a time series about its RSSI, i.e. link a→b: RSSI(t'_1), ..., RSSI(t'_m), where t'_j ($1 \leq j \leq m$) denotes the timestamp.

In order to find the correlation between RSSI and PRR, we use a link's PRR time series and RSSI time series. We match the PRR value to the nearest RSSI value with respect to time. If the nearest time difference exceeds a threshold (10 min), we do not match the corresponding PRR and RSSI.

We also find that there are circumstances that the values of three fields in C2 packets (i.e. RSSI, linkETX, ETX) remain unchanged for a long time period. This is because the statistics of the corresponding link are not updated during the period. The link record remains in the neighbor table since the current CTP does not have an effective aging mechanism to rule it out. We discard these data in order not to bias our analysis.

In the deployment process, we use GPS cameras to record node locations. Hence, the link distance can be calculated using the node locations. In particular, the link distance d can be calculated by the haversine formula:

$$d = 2r \arcsin \sqrt{\sin^2 \frac{\phi_2 - \phi_1}{2} + \cos\phi_1 \cos\phi_2 \sin^2 \frac{\psi_2 - \psi_1}{2}} \quad (1)$$

where ϕ_1 and ϕ_2 are the longitudes of two nodes, ψ_1 and ψ_2 are latitudes, and r denotes the radius of the earth.

When necessary, we use the SING dataset [9] for comparison. This dataset contains data for 100 telosB nodes in the Mirage indoor testbed.

B. Link Performance

In this section, we analyze the link performance in terms of the observed PRR and RSSI. We also examine the relationships among PRR, RSSI, and distance.

PRR. Figure 2(a) gives the CDF of PRRs of all visible links. The PRR of a visible link is the median of its PRR time series. We can see that the majority of links are good, i.e. 75% link PRRs are higher than 0.9. We also see that the PRRs have an almost uniform distribution in the range of 0–0.9. This is because the observed PRRs are moving averages. Thus it is possible that a short time duration of poor channel condition can cause modest link performance in a long time duration.

RSSI. Figure 2(b) shows the CDF of RSSIs of all visible links. We can see that the average RSSI is relatively low under the current settings with the maximum transmission power (i.e. 0dBm) and the distribution has a wide range from -95dBm to -60dBm.

Distance. Figure 2(c) shows the CDF of distance of all visible links. The average link distance is 60m in our deployment.

Researchers spent many efforts to explain the link performance. There are many relevant works which examine the relationships among PRR, RSSI, and distance.

Aguayo et al. [7] points out that PRR has a weak correlation with RSSI for 802.11 radios while Srinivasan et al. [9] concludes that PRR has a strong correlation with RSSI for 802.15.4 radios.

There is also a consensus that there exists a transitional region where the PRR performance is highly unpredictable with respect to distance.

Cerpa et al. [31] conduct measurements in indoor (Office) and outdoor (Habitat) environments using Mica 1 and Mica

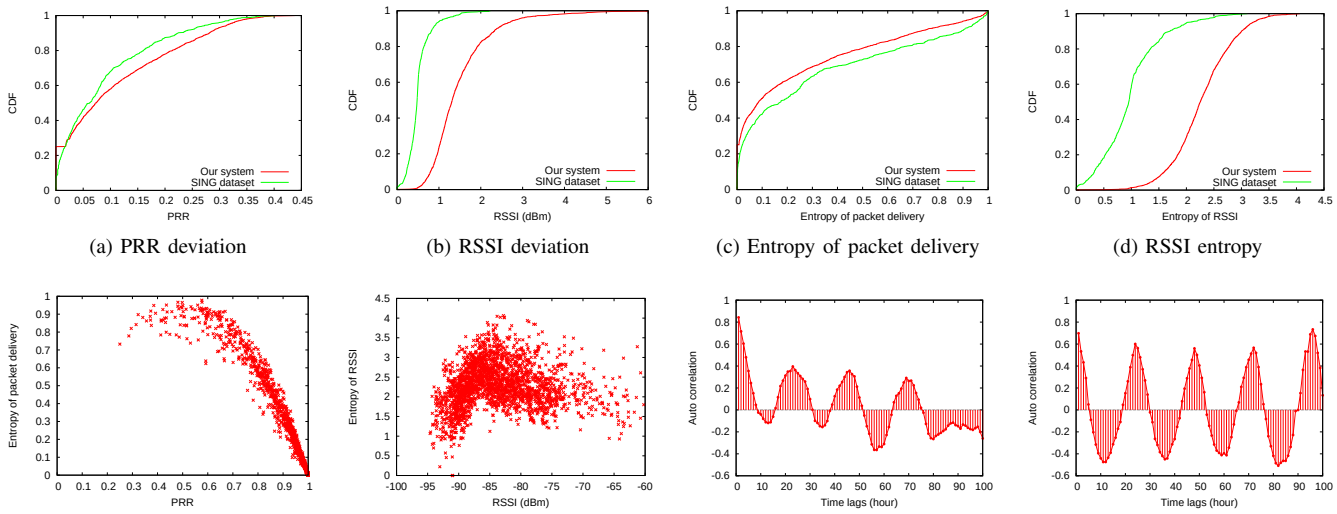

(a) PRR deviation (b) RSSI deviation (c) Entropy of packet delivery (d) RSSI entropy

(e) PRR vs entropy of packet delivery (f) RSSI vs RSSI entropy (g) PRR autocorrelation (h) RSSI autocorrelation

Fig. 3: Link temporal characteristics

2 platforms with different power levels, namely -10 dBm, -6 dBm and 1 dBm. They show that the width of the transitional region is significant, ranging from 50% up to 80% of the transmission range. On the other hand, Zhao et al. [8] show that the transitional region width is smaller, almost one-fifth to one-third of the transmission range. Using CC2420 radios, Dawson-Haggerty et al. [11] find out that the transitional region is smaller than what was previously observed with CC1000 radios: it is only 25% of the size of the good region.

We revisit these issues in our practical deployments since the quantitative results can have a significant implication on system design and network deployment. For example, if we know where the transitional region locates, we can avoid the region in order to improve link predictability.

RSSI vs PRR. Figure 2(d) shows how PRR varies with RSSI. At each RSSI value, we show the 10% percentile, 25% percentile, median, 75% percentile, and the 90% percentile of PRRs. At a glance, the RSSI-PRR curve is consistent with previous studies: PRR increases with the increase of RSSI and there is a grey region in which the PRRs are highly unpredictable. Looking deeper, there are observations worth noting.

First, the RSSI threshold, beyond which the PRR is almost 100%, is -86dBm. This threshold is much lower than what is used in many existing protocols. For example, the current CCA (Clear Channel Assessment) value in TinyOS is set to -77dBm. That means the channel will be considered as idle when the signal power is lower than -77dBm. However, from this figure, we see that the transmission would be successful w.h.p. when the signal power is higher than -86dBm. This observation implies that the current CCA value may cause heavy interference since the MAC will mistakenly consider a busy channel (e.g. with signal power between -86dBm and -77dBm) to be idle.

Second, the width of the grey region is almost 11dBm and it is much greater than 6dB reported by [9]. This may be due to several reasons. First, the RSSI is measured upon a beacon's arrival rather than on a per-packet basis. Besides, the PRR value is a moving average. These facts cause the PRR-RSSI correlation to be weaker. Second, both hardware diversity (e.g. the radio receive sensitivity is different for different nodes) and environmental factors can also cause a wider grey region.

Distance vs PRR. Figure 2(e) shows how PRR varies with distance. We see that there is a general trend that the PRR decreases as the link distance increases. We define the start of the transitional region to be the distance at which 90% links are above 90% PRR, and the end to be the distance at which 90% links are below 10% PRR. We observe that the transitional region is almost 80m, twice the width of the good region. This result deviates significantly from previous studies. It means that the link performance might have a huge variance due to environmental factors.

Distance vs RSSI. Figure 2(f) shows how RSSI varies with distance. RSSI decreases as the distance increases. We also see that even under the same distance, the RSSI variation is huge. This raises practical challenges for RSSI-based localization [32]. Besides, we observe that RSSI values seem insensitive when the distance is relatively long.

The overall link characteristics give us some practical guidelines to real-world WSN deployments. First, as a rule of thumb, the link will almost be reliable when RSSI>-86dBm. This threshold should be carefully measured in practical protocol design. Increasing the transmit power (e.g. using external antennas) is obviously effective in improving the link performance. However, special care should be taken to avoid potential interference. Second, current CCA value in TinyOS CSMA MAC is likely to cause unnecessary collisions since it might mistakenly consider a busy channel to be idle.

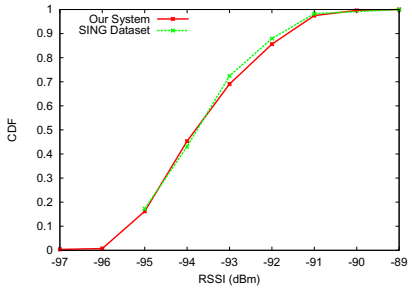

Fig. 4: Noise floor

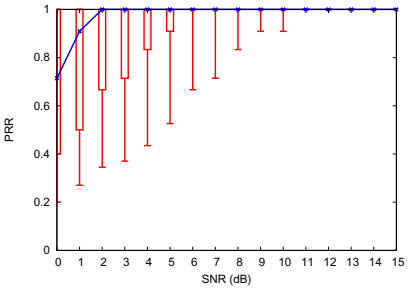

Fig. 5: SNR vs PRR

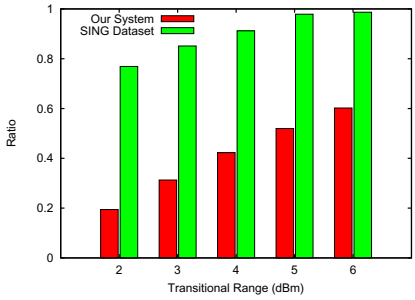

Fig. 6: Correlation ratio

C. Temporal Characteristics

We now look at the temporal characteristics by calculating the standard deviation and entropy.

Standard deviation. Figure 3(a) plots the CDF of standard deviation of PRR and Figure 3(b) plots the CDF of standard deviation of RSSI. We compare our results to those using the SING dataset [9]. The average PRR deviation is 0.107 for our dataset, comparing to 0.087 for the SING dataset. The average RSSI deviation is 1.5dBm for our dataset, comparing to 0.5dBm for the SING dataset. While PRR dynamics are comparable, our system has higher RSSI dynamics. The instability of our system is probably due to the environmental impact.

Entropy. We examine a more fundamental notion of unpredictability by using the concept of entropy. We define the entropy $H(x)$ by

$$H(X) = -\sum_{i=1}^{n} p(x_i)\log_2 p(x_i) \tag{2}$$

We measure the entropy of packet delivery and observed RSSI. For packet delivery, X indicates a binary random variable of packet delivery and x_i is a value of the random variable, which indicates either a success or a failure. For RSSI, X indicates a discrete random variable of RSSI and x_i indicates possible RSSI values.

Figure 3(c) shows the CDF of the entropy of packet delivery and Figure 3(d) shows the CDF of the RSSI entropy. We can see similar results as shown in Figure 3(a) and Figure 3(b), i.e. the predictability of PRRs is comparable while the predictability of RSSI is lower in our system.

We proceed to look at the correlation among entropy, RSSI and PRR.

PRR vs packet delivery entropy. Figure 3(e) shows the scatter plot of PRR and the entropy of packet delivery. We see that high quality links (and low quality links) have a more predictable performance while the intermediate links are less predictable. This result is consistent with previous results.

RSSI vs RSSI entropy. Figure 3(f) shows the scatter plot of RSSI and the RSSI entropy. Similar to Figure 3(e), intermediate links are less predictable. Different from Figure 3(e), the variation is higher for the intermediate links, i.e. while some intermediate links are more unpredictable, some others are as predictable as good links.

Autocorrelation. Being aware that the link performance can have a high variation, we are interested in whether or not the changes have some intrinsic patterns. To this end, we measure the autocorrelation of PRR and RSSI. We use the following formula to calculate the autocorrelation:

$$\hat{R}(k) = \frac{1}{(n-k)\delta^2}\sum_{t=1}^{n-k}(X_t-\mu)(X_{t+k}-\mu) \tag{3}$$

where k denotes the time lag (hours); δ denotes the variance of total samples; and μ denotes the mean of total sample. The value of 1 indicates the strongest positive correlation while the value of -1 indicates the strongest negative correlation.

Figure 3(g) shows the PRR autocorrelation of a representative link 1186–1153. Figure 3(h) shows the RSSI autocorrelation of a representative link 1021–1363. We indeed see that some links exhibit periodic patterns. The period is approximately 24 hours, indicating that environmental differences in day and night may cause link quality changes.

V. Causes of Link Degradation

In this section, we seek to answer the following important questions:

1) Can RSSI or SNR (Signal to Noise Ratio) predict the link performance in practice?
2) To what extent RSSI/SNR can predict the link performance?

The work in [9] concludes that link degradation correlates closely with RSSI fluctuating near the noise floor while the work of [7] finds a weak correlation between link performance and RSSI. We desire to quantify the correlation between RSSI/SNR and link degradation.

As we can see in Figure 2(d) that the RSSI grey region is about 11dB which is wider than the sensitivity of 3dB reported in the CC2420 datasheet. One reason is due to the hardware diversity, i.e. the RSSI noise floors (the RSSI in absence of packet transmissions) of different nodes are different. Thus SNR is a better indicator of signal reception. SNR measures the signal power relative to the noise floor. It can isolate the impact of hardware diversity and is considered to have a better correlation with link performance [24].

In order to calculate SNR, we first need to calculate the noise floors of all nodes. The detailed calculation process can be referred to [9]. The basic idea is to find out the minimum

valid RSSI value. It is worth noting that the approximated RSSI noise floor may be larger than the real value since the node collects the RSSI upon a beacon's arrival. However, the error should be small because (1) a node can hear weak signals from far away nodes by chance, (2) the RSSI time series usually contains sufficient data including low RSSI values.

Noise floor. Figure 4 shows the approximated noise floors in our system , which are in general similar to previous results [9].

After obtaining the noise floor, we revisit the SNR-PRR correlations. It is supposed to have a better correlation with PRR and the grey region will be smaller than we see in Figure 2(d).

SNR-PRR correlation. Figure 5 shows the SNR-PRR correlation in CitySee. We clearly see that the width of the grey region is approximately 6dB, smaller than 11dB as observed in Figure 2(d), indicating that SNR-PRR correlation is better than the RSSI-PRR correlation.

In order to quantify whether the PRR degradation correlates closely with RSSI fluctuating near the noise floor, we make the following definitions:

1) PRR degradation event. For time instant t_i, if PRR(t_i)<0.85, it is considered as a PRR degradation event for the given link.
2) Weak signal event. For time instant t_i, if SNR(t_i)=RSSI(t_i)-noise<w, it is considered as a weak signal event.
3) For a given link, if the PRR degradation event and weak signal event have been observed close in time, i.e. for a given link, PRR(t_i)<0.85, SNR(t_j)<w, and $|t_i - t_j|$<10min, we consider that the PRR degradation event is correlated with the weak signal event. It is highly probable that the PRR degradation is caused by the weak signal power.

It is worth explaining some of the parameters mentioned above. We denote w as the sensitivity of the radio. It is reported as 3dB in the CC2420 datasheet [33], and is measured as 6dB in both our study (cf. Figure 5) and the previous study [9]. We consider it to be in the range of 3dB–6dB.

We then calculate the ratio of the number of correlated events to the number of all PRR degradation events. SNR can better explain the degradation if the ratio is closer to 1.

Figure 6 shows the ratio with respect to w. We can see that the percentage of the correlated events increases with the increase of w. This is because the number of correlated events will be larger with less stringent criteria. Our system has a lower ratio than that of [9]. With w=6dB, the ratio of SING dataset is 98% while the ratio is 60% for our system, indicating that PRR does not correlate closely with RSSI fluctuating near the noise floor, and there might be other factors impacting link degradation, e.g. interference, multipath, etc.

VI. Conclusion

In this paper, we present a measurement study on the link-level behaviors in a large-scale urban sensor network. By carefully analyzing the performance metrics, we investigate the system performance, characterize link behaviors, and investigate the causes of link loss and the impacts of link performance on wireless routing.

The contributions of this work are follows: (1) We design and deploy a large-scale urban sensor network, providing us the measurement platform for observing system behaviors at scale. (2) We conduct a comprehensive measurement study to investigate the system performance, characterize link behaviors, explore the causes of link degradation, and examine the effect of link performance on wireless routing. (3) We present observations in our network. The empirical findings have many implications to WSN deployments and network protocol designs.

Future work leads to two directions. (1) We would like to investigate other important aspects of sensor networking, such as routing dynamics, traffic distributions, and etc. (2) We would like to improve our system design based on the implications.

Acknowledgement

This work is in part supported by NSFC under grant 61572277, NSFC Joint Research Fund for Overseas Chinese Scholars and Scholars in Hong Kong and Macao under grant 61529202, NSFC key program 61532012 and Fundamental Research Funds for the Central Universities (No. 2016FZA5010).

References

[1] T. He, P. Vicaire, T. Yan, Q. Cao, G. Zhou, L. Gu, L. Luo, R. Stoleru, J. A. Stankovic, and T. F. Abdelzaher, "Achieving Long-Term Surveillance in VigilNet," *ACM Trans. on Sensor Networks*, vol. 5, no. 1, pp. 1–39, 2009.

[2] G. Barrenetxea, F. Ingelrest, G. Schaefer, M. Vetterli, O. Couach, and M. Parlange, "SensorScope: Out-of-the-Box Environmental Monitoring," in *Proc. of ACM/IEEE IPSN*, 2008.

[3] T. Yan, D. Ganesan, and R. Manmatha, "Distributed image search in camera sensor networks," in *Proc. of ACM SenSys*, 2008.

[4] A. Arora, R. Ramnath, E. Ertin, and et al., "ExScal: Elements of an Extreme Scale Wireless Sensor Network," in *Proc. of IEEE RTCSA*, 2005.

[5] K. Tan, J. Fang, Y. Zhang, S. Chen, L. Shi, J. Zhang, and Y. Zhang, "Fine-grained Channel Access in Wireless Lan," in *Proc. of ACM SIGCOMM*, 2010.

[6] S. Hong and S. Katti, "DOF: A Local Wireless Information Plane," in *Proc. of ACM SIGCOMM*, 2011.

[7] D. Aguayo, J. Bicket, S. Biswas, G. Judd, and R. Morris, "Link-level Measurements from an 802.11b Mesh Network," in *Proc. of ACM SIGCOMM*, 2004.

[8] J. Zhao and R. Govindan, "Understanding Packet Delivery Performance In Dense Wireless Sensor Networks," in *Proc. of ACM SenSys*, 2003.

[9] K. Srinivasan, P. Dutta, A. Tavakoli, and P. Levis, "Understanding the Causes of Packet Delivery Success and Failure in Dense Wireless Sensor Networks (Technical report SING-06-00)," Stanford University, Tech. Rep., 2006.

[10] A. Natarajan, B. de Silva, K. kiong Yap, and M. Motani, "Link Layer Behavior of Body Area Networks at 2.4GHz," in *Proc. of ACM MobiCom*, 2010.

[11] S. Dawson-Haggerty, J. Ortiz, X. Jiang, and D. Culler, "The Effect of Link Churn on Wireless Routing," University of California, Berkeley, Tech. Rep., 2008.

[12] D. Turner, K. Levchenko, A. C. Snoeren, and S. Savage, "California Fault Lines: Understanding the Causes and Impact of Network Failures," in *Proc. of ACM SIGCOMM*, 2010.

[13] F. Wang, Z. M. Mao, J. Wang, L. Gao, and R. Bush, "A Measurement Study on the Impact of Routing Events on End-to-End Internet Path Performance," in *Proc. of ACM SIGCOMM*, 2006.

[14] A. Mahimkar, H. H. Song, Z. Ge, A. Shaikh, J. Wang, J. Yates, Y. Zhang, and J. Emmons, "Detecting the Performance Impact of Upgrades in Large Operational Networks," in *Proc. of ACM SIGCOMM*, 2010.

[15] P. Gill, N. Jain, and N. Nagappan, "Understanding Network Failures in Data Centers: Measurement, Analysis, and Implications," in *Proc. of ACM SIGCOMM*, 2011.

[16] W. Wu and J. Lui, "Exploring the Optimal Replication Strategy in P2P-VoD Systems: Characterization and Evaluation," in *Proc. of IEEE INFOCOM*, 2011.

[17] C. Wu, B. Li, and S. Zhao, "Characterizing Peer-to-Peer Streaming Flows," *IEEE Journal on Selected Areas in Communications*, vol. 25, no. 9, pp. 1612–1626, 2007.

[18] Y. Liu, Y. He, M. Li, J. Wang, K. Liu, L. Mo, W. Dong, Z. Yang, M. Xi, and J. Zhao, "Does Wireless Sensor Network Scale? A Measurement Study on GreenOrbs," in *Proc. of IEEE INFOCOM*, 2011.

[19] R. Szewczyk, J. Polastre, A. Mainwaring, J. Anderson, and D. Culler, "Analysis of a large scale habitat monitoring application," in *Proc. of ACM SenSys*, 2004.

[20] G. Tolle, J. Polastre, R. Szewczyk, D. Culler, N. Turner, K. Tu, S. Burgess, T. Dawson, P. Buonadonna, D. Gay, and W. Hong, "A Macroscope in the Redwoods," in *Proc. of ACM SenSys*, 2005.

[21] G. Werner-Allen, K. Lorincz, J. Johnson, J. Lees, and M. Welsh, "Fidelity and Yield in a Volcano Monitoring Sensor Networks," in *Proc. of USENIX OSDI*, 2006.

[22] S. Bapat, V. Kulathumani, and A. Arora, "Analyzing the Yield of ExScal, a Large-Scale Wireless Sensor Network Experiment," in *Proc. of IEEE ICNP*, 2005.

[23] A. P. Subramanian, J. Cao, C. Sung, and S. R. Das, "Understanding Channel and Interface Heterogeneity in Multi-channel Multi-radio Wireless Mesh Networks," in *Proc. of the 10th International Conference on Passive and Active Network Measurement*, 2009.

[24] R. Maheshwari, S. Jain, and S. R. Das, "A Measurement Study of Interference Modeling and Scheduling in Low-Power Wireless Networks," in *Proc. of ACM SenSys*, 2008.

[25] Z. Cao, Y. He, Q. Ma, and Y. Liu, "L2: Lazy forwarding in low duty cycle wireless sensor networks," *IEEE/ACM Transactions on Networking*, vol. 23, no. 3, pp. 922–930, 2015.

[26] Q. Ma, K. Liu, X. Miao, and Y. Liu, "Sherlock is around: Detecting network failures with local evidence fusion," *IEEE Transactions on Parallel and Distributed Systems*, vol. 26, no. 5, pp. 1430–1440, 2015.

[27] W. Du, J. C. Liando, H. Zhang, and M. Li, "When pipelines meet fountain: Fast data dissemination in wireless sensor networks," in *Proceedings of ACM SenSys*, 2015.

[28] D. S. J. D. Couto, D. Aguayo, J. Bicket, and R. Morris, "A High-throughput Path Metric for Multi-hop Wireless Routing," in *Proc. of ACM MobiCom*, 2003.

[29] O. Gnawali, R. Fonseca, K. Jamieson, D. Moss, and P. Levis, "Collection Tree Protocol," in *Proc. of ACM SenSys*, 2009.

[30] R. Fonseca, O. Gnawali, K. Jamieson, and P. Levis, "Four-Bit Wireless Link Estimation," in *Proc. of the 6th Workshop on Hot Topics in Networks (HotNets VI)*, 2007.

[31] A. Cerpa, N. Busek, and D. Estrin, "SCALE: A Tool for Simple Connectivity Assessment in Lossy Environments," University of California, Los Angeles (UCLA), Tech. Rep., 2003.

[32] S. Sen, B. Radunovic, R. Roy Choudhury, and T. Minka, "Precise indoor localization using PHY information," in *Proc. of ACM MobiSys*, 2011.

[33] Texas Instruments Inc., "CC2420 datasheet." [Online]. Available: http://www.ti.com/product/cc2420