

# IEEE 802.15.4 Channel Diversity in an Outdoor Environment

Atis Elsts  
SICS Swedish ICT\*  
atis.elsts@sics.se

Hjalmar Wennerström, Christian Rohner  
Department of Information Technology  
Uppsala University, Sweden  
{name.surname}@it.uu.se

## ABSTRACT

Low-power wireless link quality is known to be frequency-dependent because of multipath fading and other factors. We present a performance study of IEEE 802.15.4 radio links that quantifies and analyzes this frequency-specific performance in a clear-field outdoor environment. Using data from 16 channels on 240 links, we show that effect from channel selection on the average link is up to 4.89 dB, comparable with the effect from 38.7 °C change in temperature. These results provide a performance baseline for other environments, as the diversity can be expected to further increase in environments with more obstacles and external interference.

## 1. INTRODUCTION

There is a growing realization within the low-power wireless research community that making use of frequency diversity through channel switching and selection techniques can significantly improve radio link reliability, in this way leading to more dependable & energy efficient sensing systems.

However, it is commonly believed that these techniques are primarily suited for environments that are either interference-rich or obstacle- and reflection-rich. This paper contributes to the state of art by presenting the results of an IEEE 802.15.4 radio link performance study in an environment that fits in neither of these two categories, and therefore provides a performance baseline in conditions not specifically rewarding for frequency diversity techniques. Our results show that even in our test network, channel selection still matters; therefore we add one more datapoint in support of using multichannel techniques.

We start by describing the experimental setup of the data-collection experiments. The results are then analyzed to quantify both the network-wide and per-link effects from channel selection. We then compare these results with temperature effects on link performance, speculate about the causes of this diversity, and present initial results of trying to extrapolate link performance between channels.

\*This work has been done while at Uppsala University.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from [permissions@acm.org](mailto:permissions@acm.org).  
Copyright 20XX ACM X-XXXXX-XX-X/XX/XX ...\$15.00.




Figure 1: **The test site.** The figure shows the nodes mounted on the poles, as well as the relatively small number of stationary obstacles in the site, such as the mast of the weather station.

## 2. EXPERIMENTAL SETUP

We used a 16-node test network (Fig. 1) permanently deployed in a field in Marsta (near Uppsala, Sweden). During June and July of 2015 we collected link quality data as well as temperature and humidity measurements. The physical setup was the same as described in [8].

The test application broadcasted out bursts of 100 packets (each 66 bytes) from each node in the network in a sequential order; all non-sending nodes listened and logged the number of packets received as well as RSSI and LQI. After all nodes had served in the sender role, the active IEEE 802.15.4 2.4 GHz channel was switched using a pre-computed pseudorandom schedule. After all 16 channels had been tested, a new schedule was computed and the whole cycle restarted.

During June we used Tmote Sky nodes; for the experiments in July we replaced every other Sky node with Zolertia Z1; both platforms have the TI CC2420 radio chip. We selected data<sup>1</sup> from 7 consecutive days in June and 6 consecutive days in July for further analysis in this paper, as during these periods the experiment was running without interruptions. In the first period, around 57 million packets were sent; in the second, around 59 million. The effective frequency was approximately one 100-packet test each 5 minutes on each of the  $16 \times 240 = 3840$  different channel-links in the network.

## 3. ANALYSIS OF THE RESULTS

**Overall performance.** Some channels are clearly better than others: the average PDR varies from 48.9% on channel 11 to 57.2% on channel 22 (Fig. 2), and the number of

<sup>1</sup>The trace files are available at <http://www.it.uu.se/research/group/coregroup/software/marsta>

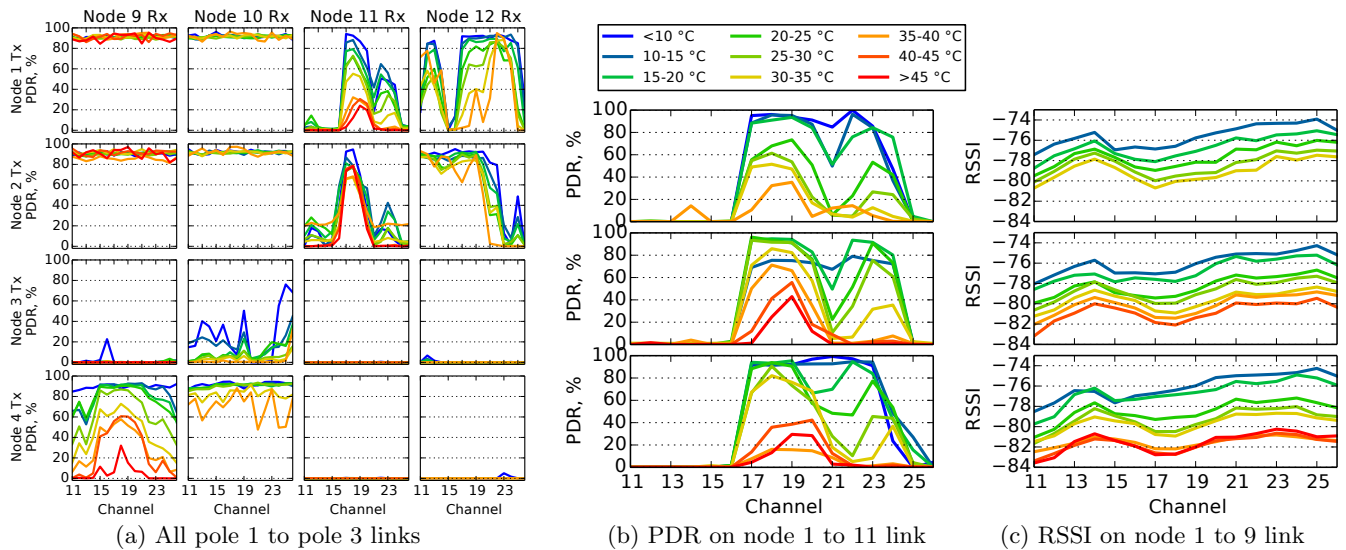


Figure 4: **Weekly (June 6-12) average temperature-dependent performance of all pole 1 to pole 3 links (left) and the daily-average performance of two links in June 10-12 (center and right).**  $T^{\circ}$  measurements from receiver nodes are used for the  $T^{\circ}$  categories. (a) Three classes of links are clearly visible: (1) in the four figures at the upper left corner, upper-upper node links have only small random packet loss; (2) the upper right and lower left corners display intermediate links (upper-lower & lower-upper nodes) with highly varied PDR; (3) links at the lower right corner (lower-lower nodes) receive only occasional packets. The channel selection and temperature effects are only visible on intermediate-quality links. As nodes 10 and 12 are shaded by the pole, they have no data about the high extremes of the temperature. (b) PDR of a link with RSSI close to the noise floor. (c) RSSI of a different link (with stable PDR).

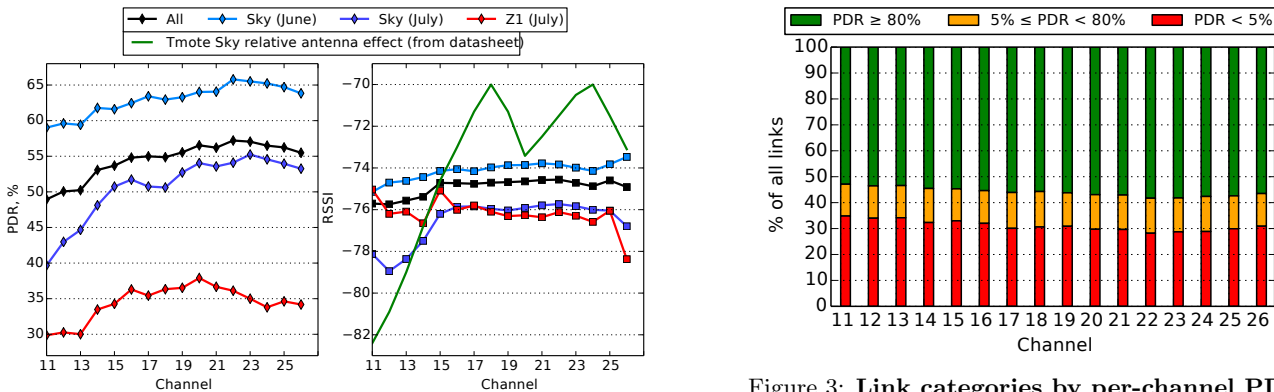


Figure 2: **All link-average PDR and RSSI values for the extended measurement periods.** Channels 11-13 show reduced performance. The Tmote Sky antenna parameter curve [3] does not match the observed RSSI performance.

intermediate and good links (i.e., with  $PDR > 5\%$ ) varies from 65% on channel 11 to 72% on channel 22 (Fig. 3). In particular, channels 11-13 have noticeably lower PDR and RSSI. We are tempted to explain this as a CC2420 radio chip-specific behavior: the PDR is smaller there even when the RSSI remains at the same level, especially for Zolertia nodes. If some unknown sources of external interference were the cause, the RSSI would be higher.

**Link-specific performance.** The effects are far stronger if considered per-link basis. We want to especially stress that for links with RSSI close to the noise floor, PDR is critically dependant on both temperature effects and channel selection (Fig. 4b). Many links are in this class (Fig. 4a). Figure 3 shows that while the proportion of intermediate links (with  $5\% < PDR < 80\%$ ) per channel is only 12-13%, around 20% of links have at least one intermediate channel.

Figure 3: **Link categories by per-channel PDR**

**Comparison of & interaction with temperature.** The impact of temperature on low-power wireless communication is well known in the research literature [1] [8], and is highly visible in our test environment (Fig. 4).

In our data, the average effect of  $10^{\circ}\text{C}$  increase in temperature on both transmitter and receiver nodes leads to 1.26 dB decrease in RSSI. This value is calculated by averaging the linear regression coefficients of Tx and Rx temperature effects on each link-channel with at least 3% PDR; RSSI values that are either outliers ( $\geq 3 SD$  from the mean) or close to CC2420 nonlinearities (see [1]) are ignored. As the average per-link *RSSI spread* (defined here as the difference between the average RSSI in the best and in the worst channel) in our data is 4.89 dB (with large  $SD=2.54$  between links), choosing a good channel can compensate up to  $38.7^{\circ}\text{C}$  increase in temperature on the average link.

Our results show that the dependence of RSSI on channel selection is described by average link-specific Pearson's correlation 0.40, which is smaller but comparable with RSSI dependence on temperature (negative correlation 0.58).

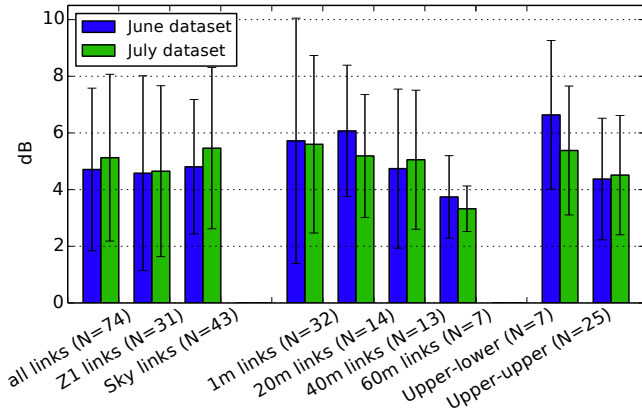


Figure 5: **The average RSSI spreads** of different groups of links. The bar “Z1 links” for June shows the data from the Tmote Sky nodes that were later replaced by Z1. Upper-lower and upper-upper links do *not* include links between nodes on the same pole.

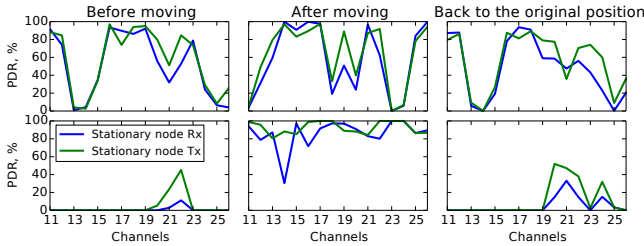


Figure 6: **The change in the measured PDR levels after some nodes are displaced.** The two rows show data from two different pairs of nodes. In each of these pairs, one node is displaced and the other one remains stationary. The columns show: (1) the initial, baseline performance; (2) performance after the mobile node is placed 20 cm away from its initial position; (3) performance after the mobile node is returned to its initial position.

We were interested in, but failed to detect any statistically significant *interactions* between channel diversity and temperature. It seems that the relative channel performance remains approximately the same throughout a large temperature range (from 5 to 50 °C; Fig. 4b). If so, this makes reasoning about these two performance factors simpler, as they can be treated as independent variables.

**The possible causes of channel diversity.** We measure the amount of channel diversity on a link by its RSSI spread, as defined above. We calculate the average RSSI spread of a link by averaging its RSSI spread values in all measured points of time. However, naïvely averaging all data would systematically underestimate this spread, especially for links close to the noise floor. The problem is that a radio chip can only detect RSSI of packets above its sensitivity threshold. To mitigate the effect of the unseen packets, we first exclude tests in which some channels have received no packets; if a link has more than 10% of tests excluded in this way, we ignore that link completely.

The main link-specific cause of the channel diversity is likely to be multipath effects. This suggests that links with more obstacles nearby should have correspondingly larger spread. To quantify this, we separate the links in several classes (Fig. 5). Unexpectedly, the graph suggests that longer links have lower RSSI spreads (despite more obstacles in path, on average!). In contrast, upper-upper node links (the

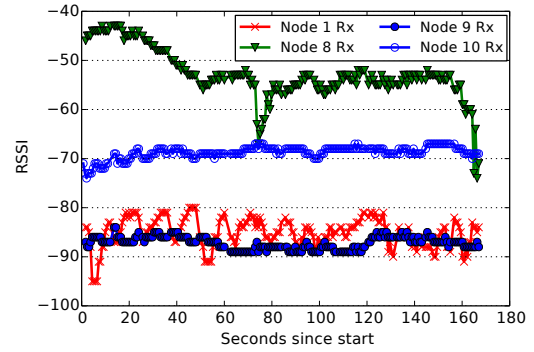


Figure 7: **Measured RSSI levels during the second experiment with movement.** Data from four different stationary nodes; the transmitter node is continuously and slowly moved in a straight line from its original position.

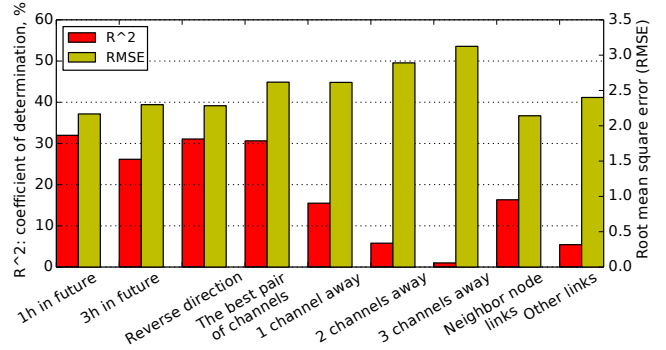


Figure 8: **The accuracy of channel-specific performance predictions based on various types of data.** The quantity predicted is the relative RSSI (positive or negative) of each channel, compared to the 16-channel average RSSI of the link at each point in time. The two prediction metrics are  $R^2$  (coefficient of determination, higher is better) and RMSE (root mean square error, lower is better).

Fresnel zone completely free) have lower RSSI spread. It is therefore more likely that the channel diversity is caused by reflections from the ground rather than from the small obstacles. In contrast, we detect no significant differences between Sky and Z1 performance in this aspect; the performance differences between the months dwarf the differences between the HW platforms.

We perform two additional tests to confirm that multipath fading is a likely cause of this diversity. In the first test, four nodes are moved for 20 cm from their initial position, and the PDR on other nodes is measured before moving, after moving, and after placing them back. The results show clear change in per-channel performance during the displacement in 11 cases (two of them are shown in Fig. 6), while just in 2 cases gives contradictory results (performance more different before and after, rather than during the displacement).

In the second test, we measure the effect on RSSI when a transmitter node is slowly moved along a continuous line (Fig. 7). The receivers show large & different RSSI variances in this test, irrespective of their absolute RSSI values. These variances have 0.41 correlation with link RSSI spreads (Pearson; not statistically significant due to the small number of receivers taken into account, i.e., with  $\geq 90\%$  PDR).

**Can channel diversity be predicted?** Before attempting to answer the question, consider that the outcome would be informative in either way. If channel-specific performance

varies independently on each link, then this variance can be exploited to achieve higher reliability of network protocols very easily, e.g., by simple random channel hopping. If, in contrast, some channels are closely correlated, this fact can be exploited to predict PDR and RSSI on these channels by measuring just one of them, thus leading to reduced energy consumption for link estimation. On the other hand, it is likely to reduce the effectiveness of random channel hopping; smarter strategies must be used to avoid fixating on a set of all-bad channels.

The research literature suggests that even neighboring channels and neighboring nodes should show independent performance in most cases. Watteyne *et al.* [7] names 5 MHz as the empirical coherence bandwidth and 5.5 cm as the coherence displacement (for >4m long links), both sufficient to move an IEEE 802.15.4 link out of a deep fade zone.

To establish a set of baselines, we first compute the predictive power of data on the same channel-link combination, but taken either 1 h or 3 h in the past, and also in the reverse direction of the same link (Fig. 8). Both past data and reverse-direction data are commonly used for wireless link quality estimation.

We compare the baselines with the predictive performance of data from neighboring channels. The results show that even the directly neighboring channels have predictive power that is on average probably too small for practical applications ( $R^2=16\%$ ). However, occasionally some pairs of channels match each other very well. This can be exploited to find such a good pair on almost every link; extrapolating from one channel to the other in this pair is on the average comparable with predicting just 1h in the future ( $R^2=31\%$  vs.  $R^2=33\%$  respectively).

Data from neighboring nodes<sup>2</sup> have practically no predictive power on relative PDR, but, interestingly, have some non-negligible correlation ( $R^2=17\%$ ) with relative RSSI, suggesting that spatially shared environment plays a role in determining the per-frequency diversity. However, the RSSI values are slightly similar even between *non-neighboring* nodes; this suggests that some of the relative RSSI change is explained by the hardware factors common to all nodes.

## 4. RELATED WORK

Initially, many low-power wireless protocols made the assumption that performance is the same on all channels [5]. Since then, the conventional wisdom in the research community has changed and now includes the fact that external interference and multipath fading vary across channels [7] leading to performance differences. However, based on the belief that “although multipath propagation is a characteristic of almost all wireless links, it is particularly problematic in indoor environments such as offices” [7] the existing experimental work has focused on measuring channel diversity in either interference-rich or reflection-rich environments, e.g., offices [7] [5] [6], a factory [2], an oil refinery [6], or underground mines [4]. With the exception of [6], none of them provides diversity results for an outdoor field environment.

The only exception in this list [6] shows a plot of RSSI values (Fig. 4a in [6]) in a “clean outdoor” environment. In contrast to our work, this study has a very limited scope (data from just a single link). Nevertheless, the per-channel differences in RSSI in their results are similar to ours, in-

<sup>2</sup>Neighbors always have the same HW platform; we do not include the July data in the analysis for this subsection.

cluding the fact that channels 11-14 have worse RSSI; this serves as evidence that our observations are not accidental.

The conclusions from the other environments suggests that performance in the frequency band of interest varies for up to 7-8 dB in mines (Fig. 7. in [4]), that a single link can have “significantly different stability values for different channels” in a factory [2], and that “geographic proximity is not sufficient for [link performance] similarity” [2]. All of these results match ours.

## 5. CONCLUSIONS

We collected and analyzed a dataset about the relative performance of IEEE 802.15.4 channels in the 2.4 GHz frequency band on 240 links. We show that even though neither external interference nor obstacle-heavy environment is present, channel selection still matters, both network-wide and especially per individual links. The average maximal effect from channel selection on a single link is 4.89 dB, comparable with the effect from a 38.7°C change in temperature. Both our data analysis and two additional experiments suggest that multipath fading is one of the main factors responsible for this diversity.

The results are similar between both of our hardware platforms (sharing the same CC2420 radio chip), but are likely to be chip dependent, as suggested by the performance dip in channels 11-13. We did not detect a clear interaction between channels and temperature, even though  $T^\circ$  changes in the wide range from 5 to 50°C in our data.

Finally, we discovered that on the average performance of an individual link-channel varies independently both from neighboring channels and from neighboring links, even though some highly correlated link-specific channel pairs do exist. This means good news both for random channel hopping protocols and for cross-channel link estimation techniques.

## 6. REFERENCES

- [1] C. A. Boano, H. Wennerström, et al. Hot packets: A systematic evaluation of the effect of temperature on low power wireless transceivers. In *ExtremeCom*, 2013.
- [2] L. Doherty, W. Lindsay, and J. Simon. Channel-specific wireless sensor network path data. In *IEEE ICCCN*, pages 89–94, 2007.
- [3] Moteiv Corporation. Tmote Sky: Low Power Wireless Sensor Module, 2006.
- [4] C. Nerguizian, C. L. Despains, S. Affès, and M. Djadel. Radio-channel characterization of an underground mine at 2.4 ghz. *IEEE Transactions on Wireless Communications*, 4(5):2441–2453, 2005.
- [5] K. Srinivasan, P. Dutta, A. Tavakoli, and P. Levis. An empirical study of low-power wireless. *ACM Transactions on Sensor Networks*, 6(2):16, 2010.
- [6] T.-D. Tran, R. Silva, D. Nunes, and J. S. Silva. Characteristics of channels of IEEE 802.15.4 compliant sensor networks. *Wireless Personal Communications*, 67(3):541–556, 2012.
- [7] T. Watteyne, S. Lanzisera, A. Mehta, and K. S. Pister. Mitigating multipath fading through channel hopping in wireless sensor networks. In *IEEE ICC*, 2010.
- [8] H. Wennerström, F. Hermans, O. Rensfelt, C. Rohner, and L.-Å. Norden. A long-term study of correlations between meteorological conditions and 802.15.4 link performance. In *SECON*, pages 221–229, 2013.