Mitigating Inter-network Interference in LoRa Networks

Thiemo Voigt
Uppsala University, Sweden
SICS Swedish ICT
thiemo@sics.se

Martin Bor, Utz Roedig
Lancaster University, UK

m.bor@lancaster.ac.uk
u.roedig@lancaster.ac.uk

Juan Alonso
Univ. Nac. de Cuyo and Univ.
Nac. de San Luis, Argentina
jmalonso@uncu.edu.ar

Abstract

Long Range (LoRa) is a popular technology used to construct Low-Power Wide-Area Networks (LPWAN). Given the popularity of LoRa it is likely that multiple independent LoRa networks are deployed in close proximity. In this situation, neighbouring networks interfere and methods have to be found to combat this interference. In this paper we investigate the use of directional antennae and the use of multiple base stations as methods of dealing with inter-network interference. Directional antennae increase signal strength at receivers without increasing transmission energy cost. Thus, the probability of successfully decoding the message in an interference situation is improved. Multiple base stations can alternatively be used to improve the probability of receiving a message in a noisy environment. We compare the effectiveness of these two approaches via simulation. Our findings show that both methods are able to improve LoRa network performance in interference settings. However, the results show that the use of multiple base stations clearly outperforms the use of directional antennae. For example, in a setting where data is collected from 600 nodes which are interfered by four networks with 600 nodes each, using three base stations improves the Data Extraction Rate (DER) from 0.24 to 0.56 while the use of directional antennae provides an increase to only 0.32.

Categories and Subject Descriptors

C.2.1 [**Network Architecture and Design**]: [wireless communication]; C.4 [**Performance of Systems**]: [reliability, availability, and serviceability]

General Terms

Experimentation, Measurement

Keywords

LoRa, Low-Power Wide-Area Network, Interference

1 Introduction

Long Range (LoRa) Low-Power Wide-Area Network (LPWAN) devices communicate directly with base stations which removes the need of constructing and maintaining a complex multi-hop network. Multiple LoRa networks may be deployed in the same physical space which leads to inter-network interference. For example, multiple smart city applications based on LoRa may be deployed in the same area and interference between these networks will occur. To ensure acceptable network performance this inter-network interference must be managed appropriately.

LoRa transceivers can use orthogonal transmission settings (such as frequency, spreading factor, bandwidth) which in principle can be used to prevent inter-network interference. However, there are drawbacks which make this approach less viable in practical settings. First, transmitter settings have an impact on transmission properties such as range, reliability and energy consumption which prevents nodes to select parameters freely. Second, dynamically choosing parameters requires a complex protocol and network cooperation which current LoRa systems do not support. Hence, current LoRa deployments typically use a default static setting which leads to inter-network interference. For these reasons it is desirable to find additional mechanisms for dealing with this interference.

In this paper we focus on two practical alternative methods to deal with interference. Both methods aim at improving the chance of decoding a message in presence of interference. First, we consider directional antennae to improve signal strength at the receiver without increasing transmission energy cost. Second, we consider the use of multiple base stations to improve the probability of decoding a message at at least one receiver. While both methods have advantages and disadvantages in terms of practicality (such as cost, method of deployment, maintainability) it is the question which of the approaches is more effective.

In this paper we answer this question by analysing the effectiveness of both approaches via comprehensive simulation. Our purpose built simulation environment is calibrated using LoRa testbed experiments to ensure simulation results match as close as possible practical setups. Our findings show that both methods are able to improve LoRa network performance in interference settings as it would be expected. However, the results demonstrate that the use of multiple base stations clearly outperforms the use of directional

antennae. For example, in a setting where data is collected from 600 nodes which are interfered by four other LoRa networks with 600 nodes each, the use of three base stations improves the Data Extraction Rate (DER) from 0.24 to 0.56 while the use of directional antennae increases it to 0.32.

The main contributions of this paper are:

- We evaluate the impact of inter-network interference on LoRa networks showing that such interference can drastically reduce the performance of a LoRa network.

- We quantify network performance gains by introducing directional antennae and multiple base stations.

- We show that adding more base stations rather than equipping nodes with directional antennae is more efficient when mitigating LoRa network interference.

In the next section, we present essential background on LoRa. Section 3 describes briefly our previous work on LoRa [2; 3] and the resulting simulation environment used for our evaluation. Section 4 describes the evaluation of performance gains by introducing directional antennae and multiple base stations. Before concluding, we discuss related work in Section 5.

2 Long Range (LoRa)

Long Range (LoRa) is a proprietary spread spectrum modulation technique by Semtech, derived from Chirp Spread Spectrum (CSS). Instead of modulating the message on a pseudorandom binary sequence, as is done in the well known Direct-Sequence Spread Spectrum (DSSS), LoRa uses a sweep tone that increases (upchirp) or decreases (downchirp) in frequency over time to encode the message. Spreading the signal over a wide bandwidth makes it less susceptible to noise and interference. CSS in particular is resistant to Doppler effects (common in mobile applications) and multipath fading. A LoRa receiver can decode transmissions 20 dB below the noise floor, making very long communication distances possible, while operating at a very low power. LoRa transceivers available today can operate between 137 MHz to 1020 MHz, and therefore can also operate in licensed bands. However, they are often deployed in ISM bands (EU: 868 MHz and 433 MHz, USA: 915 MHz and 433 MHz). The LoRa physical layer may be used with any MAC layer; however, Long Range Wide Area Network (LoRaWAN) is the currently proposed MAC. LoRaWAN operates in a simple star topology.

A LoRa transceiver has five runtime-adjustable transmission parameters: Transmission Power (TP), Carrier Frequency (CF), Spreading Factor (SF), Bandwidth (BW), and Coding Rate (CR). These parameters have an influence on the transmission duration, energy consumption, robustness and range.

Transmission Power (TP). TP on a LoRa receiver can be adjusted between −4 dBm and 20 dBm in 1 dB steps. Because of regulatory and hardware limitations, however, this is often limited between 2 dBm and 14 dBm. TP has a direct influence on energy consumption and the range of the signal.

Carrier Frequency (CF). CF is the centre frequency, which can be programmed in steps of 61 Hz between 137 MHz to 1020 MHz.

Spreading Factor (SF). SF determines how many chips are encoded in each symbol as a power of two, and can be set between 6 and 12. A higher spreading factor increases the Signal to Noise Ratio (SNR) and therefore receiver sensitivity and range of the signal. However, it lowers the transmission rate and thus increases the transmission duration and energy consumption. The SFs in LoRa are orthogonal. Consequently, concurrent transmissions with different SF do not interfere with each other, and can be successfully decoded (assuming a receiver with multiple receive paths).

Bandwidth (BW). BW can be set from (a fairly narrow) 7.8 kHz up to 500 kHz. In a typical LoRa deployment, only 125 kHz, 250 kHz and 500 kHz are considered. A wider bandwidth means a more spread-out and therefore more interference-resilient link. In addition, it increases the data rate, as the chips are sent out at a rate equivalent to the bandwidth. The downside of a higher bandwidth is a less sensitive reception, caused by the integration of additional noise.

Coding Rate (CR). CR is the amount of Forward Error Correction (FEC) that is applied to the message to protect it against burst interference. Higher CR makes the message longer and therefore increases the time on air. LoRa transceivers with different CR, and operating in 'explicit header mode', can still communicate with each other, as the CR is encoded in the header.

3 LoRa Simulation Environment

In our previous work [2] we investigated the general scalability of LoRa networks. For this study we carried out testbed experiments to characterise LoRa link behaviour. We then used the results of this study to develop the simulation tool LoRaSim[1]. LoRaSim models (i) achievable communication range in dependence of communication settings TP, SF and BW and (ii) capture effect behaviour of LoRa transmissions depending on transmission timings and power. We extend LoRaSim for the experiments in this paper with (iii) the ability to simulate directional transmissions. Correct representation of these three effects is important as they determine if interfering transmissions can be decoded by a receiver. How effect (i) and (ii) are represented by LoRaSim is described in detail in our previous publication [2]; we include here a brief summary.

LoRaSim. LoRaSim is a custom-build discrete-event simulator implement with SimPy [1]. LoRaSim allows us to place N LoRa nodes and M LoRa base stations in a 2-dimensional space. The communication characteristics of a LoRa node are defined by the transmission parameters TP, CF, SF, BW and CR. Furthermore, a node's transmission behaviour is described by the average packet transmission rate λ and the size of the packet payload B.

LoRaSim emulates LoRa base station chips such as the Semtech SX1301. This chip can receive up to eight concurrent signals as long as these signals are orthogonal, that is, they use different SF.

Communication Range. A transmission is successfully received if the received signal power P_{rx} lies above the sensitivity threshold S_{rx} of the receiver. The received signal power P_{rx} depends on the transmit power P_{tx} and all gains and losses

---

[1] Available at http://www.lancaster.ac.uk/scc/sites/lora/

along the communication path. We use the well known log-distance path loss model [11] which is commonly used to model deployments in built-up and densely populated areas.

On the transmitter side, range can only be changed by changing the transmit power. The range can also be influenced by the use of a directional antenna (described later). Other parameters like SF, BW and CR do not influence the radiated power, or any other gains and losses. On the receiver side, the range is limited by the sensitivity threshold S_{rx}, which is influenced by the LoRa parameters SF and BW.

To determine P_{rx}, the path loss model must be configured and the communication distance d must be known. In our simulations we configure the path loss to reflect a built up environment. S_{rx} depends on the selected BW and SF. We use the measured sensitivity from calibration experiments based on the Semtech SX1272 LoRa transceiver to determine sensitivity in dependence of BW and SF.

Collision Behaviour. When two LoRa transmissions overlap at the receiver, there are several conditions which determine whether the receiver can decode one or two packets, or nothing at all. These conditions are Carrier Frequency (CF), Spreading Factor (SF), power and timing. As LoRa is a form of frequency modulation, it exhibits the *capture effect*. The capture effect occurs when two signals are present at the receiver and the weaker signal is suppressed by the stronger signal. The difference in received signal strength can therefore be relatively small. An increase of signal strength as present when using a directional antenna has significant impact on this behaviour.

Collision behaviour including capture effect is modelled in LoRaSim to match a Semtech SX1272.

Directional Antenna. We extend LoRaSim with directional transmissions. We model our transmissions according to the SPIDA antenna [8], an electronically switchable directional (ESD) antenna designed for low-power wireless sensor networks. SPIDA has six parasitic elements that can be individually grounded or isolated via a software control at negligible energy cost. If all the parasitic elements except one are grounded, the direction of the maximum antenna gain is towards the isolated element. In the experiments we let the direction of maximum gain point towards the receiving base station. As a result, the received signal power of transmissions at the intended receiver increases while it might increase or decrease at other receivers depending on their location. This increase or decrease is based on our previous measurements with SPIDA [13; 14]. We emulate an antenna that behaves approximately as SPIDA with a gain of 4 dBi in the main direction, i.e., when the parasitic element pointing towards the base station is isolated. When the two neighbouring elements are isolated, the same gain is achieved. If the parasitic element opposed the base station is isolated, the gain is decreased with 3 dBi, while for the other two elements the gain is decreased with 4 dBi. We also emulate an improved theoretical antenna where we double these values. For example, in the direction towards the base station, the gain with this antenna is 8 dBi.
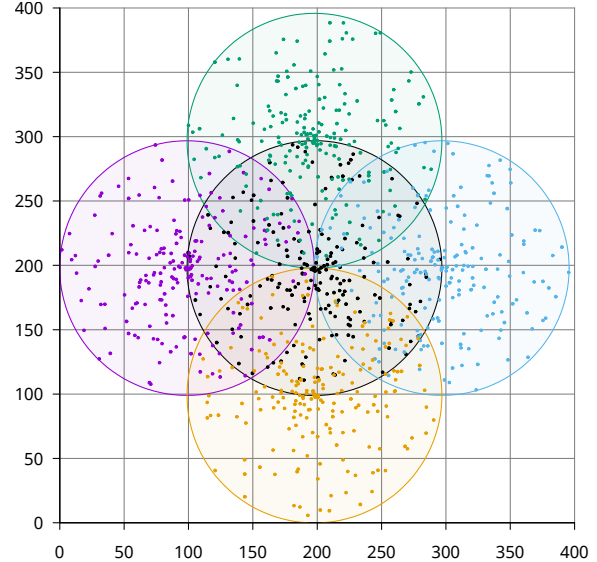

Figure 1. Example configuration used in our simulations. The black network in the centre is interfered by four other networks.

4 Evaluation

We use DER as metric for evaluating network performance. DER is defined as the ratio of received messages to transmitted messages over a period of time. Note that a message is regarded as received correctly if at least one LoRa base station of the corresponding network receives it. DER does not capture individual node performance but looks at the network deployment as a whole. When all transmitted messages arrive successfully at one of the base stations, then $DER = 1$.

All experiments use the same node configuration set, $SN = \{TP, CF, SF, BW, CR, \lambda, B\}$, whereby TP, CF, SF, BW and CR are the transmission parameters as previously defined, λ the average packet transmission rate and B the packet payload. In particular, we study a set we call SN^1 where $TP = 14$ dBm, $CF = 868$ MHz, $SF = 12$, $BW = 125$ kHz, $CR = 4/8$, $\lambda = 16.7$ min and $B = 20$ B. SN^1 corresponds to the most robust LoRa transmitter settings. SN^1 transmissions have the longest possible airtime: 1712.13 ms. Due to space constraints, we do not present results for other node settings. We have, however, verified in our simulations that other settings and in particular a setting called SN^3 [2] shows the same trends. SN^3 is similar to SN^1 except for a lower coding rate which reduces the time on air and leads to fewer collisions. Current LoRa deployments use static configurations such as SN^1 or SN^3; for example, LoRaWAN based deployments use SN^3.

In our experiments we create networks by placing N nodes randomly within a circle of radius R around a base station. The distance between nodes and base stations is such that all nodes can reach the base station with the given transmitter settings. If no interference occurs transmissions of nodes reach their base station without loss. In the experiments we use a radius of $R = 99$ m which represents a realistic range for built-up environments [3]. In the experiments
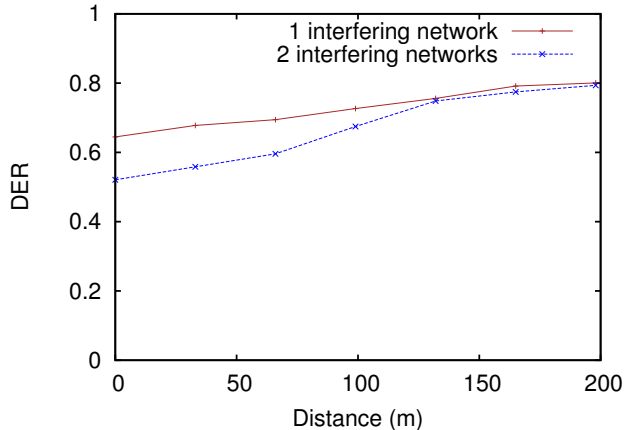
Figure 2. When the distance to the interfering base stations increases, the DER of a deployed LoRa network increases as expected.

Figure 3. With more interfering networks, DER decreases significantly in particular when the number of nodes is high.

we deploy a variable number of interfering networks around the network of interest (called the interfered network); interfering networks are deployed in the same way as our main network. Figure 1 shows an example configuration; the black network is the network of interest; the other 4 networks are interfering systems. In all experiments, we assume a 20 Byte packet is sent by each node every 16.7 min representing a realistic application; the main network and interfering networks use this transmission pattern.

4.1 Impact of Interfering Networks

We use a setup as depicted in Figure 1 and described previously. The network of interest, i.e., the interfered network, is the one in the centre (see Figure 1). The interfered network's transmissions to the base station might interfere with packets from other networks (or with packets from the interfered network itself) depending on the position of the base stations. Our goal is to evaluate performance of this network in form of DER. In the first experiment, we assume that each network has $N = 200$ nodes. We vary the distance between the base stations of the interfering networks to the base station of the interfered network. The purpose of this first experiment is to show how inter-network interference impacts on DER. In the second experiment, we vary the number of nodes per network and the number of interfering networks. The distance between the base stations in this second experiment is 99 m which is also the network radius.

The results of the first experiment are shown in Figure 2. When all base stations are placed at the same location, i.e., the distance is zero, the interference is the highest since the transmissions of all nodes in the interfering network can interfere with the transmissions of the interfered network. With an increasing distance, less nodes of the interfering networks interfere with the transmissions of the interfered networks which leads to a higher DER. When the distance between the base stations is 200 m, no interference between the base stations is possible. DER is here around 0.8 which is the maximum achievable performance due to interference from within the own network.
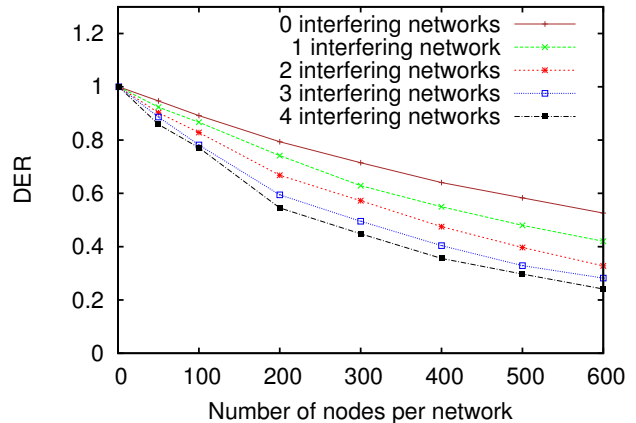
Figure 3 depicts the results of the second experiment. As expected, when the number of interfering networks increases, DER decreases significantly, in particular when the number of nodes is high. For example, with 500 nodes per network, the DER of the interfered network decreases from 0.58 without interference to ca. 0.3 when there is interference from four networks.

The experiments show that the deployment of non-cooperating LoRa networks in the same space has a significant impact on network performance. Hence, it is desirable to address and mitigate this issue.
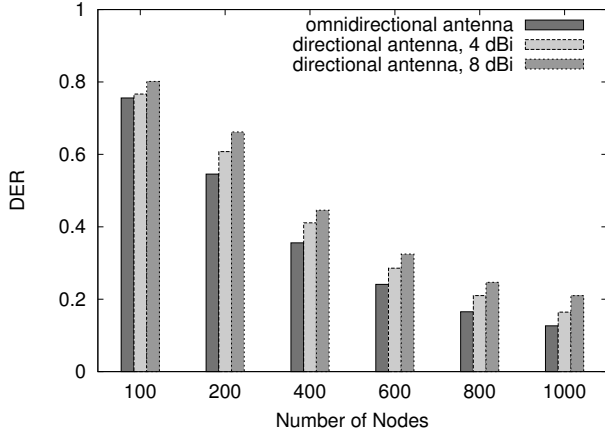
4.2 Using Directional Antennae

In the experiments in this section we evaluate to which extend directional antennae can improve the DER of an interfered LoRa network. We expect that this is possible as the directional antennae can radiate more energy towards the intended base station thereby increasing the signal strength at the base station. Directional antennae also reduce the interference at other base stations which is likely to increase the overall performance of all networks. This is, however, not the focus of this study. The experimental setup is equivalent to the previously used setups; however, nodes in the network of interest (the centre network in Figure 1) are equipped with directional antennae.
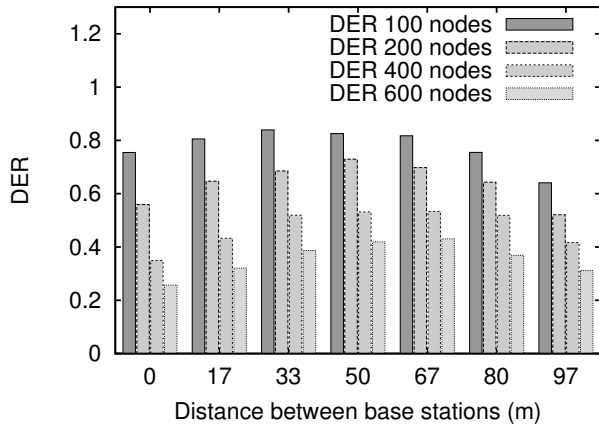
Figure 4 depicts the results. The figure shows that as expected, directional transmissions improve DER, in particular when the number of nodes is high. This is the case as the signal strength of nodes of the interfered network increases. As consequence, it is more likely that due to the capture effect these transmissions succeed even when there are collisions. For most of the setups, DER increases by about 0.04 when we equip the nodes in the interfered network with directional antennae. Using even better directional antennae (8dBi gain), the DER increases by another 0.04.

4.3 Using Additional Base Stations

Our previous work has shown that one way to make LoRa networks scale is to increase the number of base stations [2]. In the experiments in this section, we evaluate whether this

Figure 4. Comparison with omnidirectional antenna, directional antenna and better directional antenna (8 dBi) for four interfering networks. Directional antennae increase the DER.
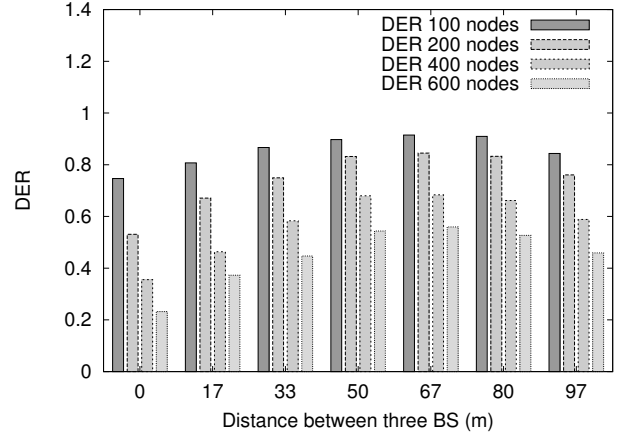

Figure 6. Impact of distance between three base stations on DER for different number of nodes.

Figure 5. Impact of distance between two base stations on DER for different number of nodes.

Figure 7. Summary of results: Deploying multiple base stations is more efficient than using directional antennae.
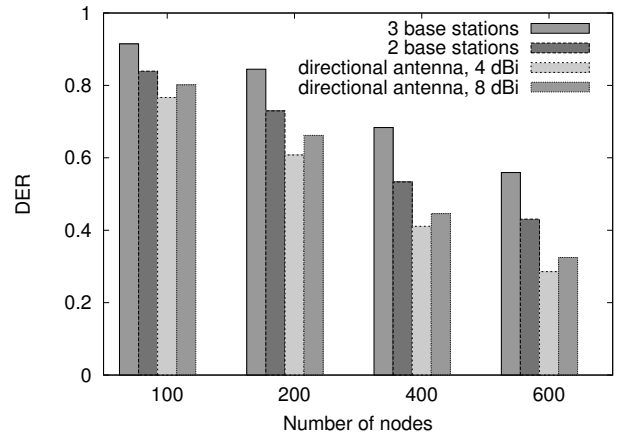
is also true in interference settings.

We replace the base station in the centre of the setup shown in Figure 1 by two and three base stations respectively. We place these additional base stations at a distance d from the original base station. For two stations we move the original base station d to the right (leaving its vertical position as it is) and add an additional base station d to the left of the original location. When replacing the original base station with three base stations, we move one base station upwards by d and the other two $45°$ down and to the left and right respectively, so that the distance is also d from the original location of the base station. The placement of the sensor nodes is unchanged, i.e., they are placed within the radius r around the original location of the base station. A packet transmission is counted as successful if either of the base stations receives it. All four interfering networks are active.

Figure 5 shows the results when the original base station is replaced with two base stations. The figure depicts that

with a distance of $0\,$m the DER is quite low. There is no improvement compared to the results with the omnidirectional antennae in Figure 4: placing two base stations at the same place does not change anything as they will receive exactly the same packets. For larger distances like $97\,$m, the DER is even lower since some nodes might not even reach the base station. The best distances are between $33\,$m and $67\,$m for all setups. For 600 nodes, the DER increases from 0.35 to 0.53 for distances between the base stations of $50\,$m and $67\,$m.

Figure 6 depicts the results when the original base station is replaced with three base stations. In general, while the trends are similar to those in Figure 5, the DER is higher than with two base stations. In particular, the results with larger distance, e.g., $97\,$m are much better. The reason is the distribution of the base station that ensures that all nodes are in reach of a base station which was not the case for two base stations. Also, the overall results are higher since the chance that a transmission finds a base station where the capture effect comes into play increases.

4.4 Discussion

Using the experimental results, we can now answer the question if it is better to equip sensor nodes with directional antennae or to deploy additional base stations to achieve a high DER under interference. The result in Figure 7 shows that to achieve a high DER under interference, deploying multiple base stations is more efficient than using directional antennae. Moreover, Figure 3 shows that with multiple base stations, LoRa can achieve a DER that is higher than for one base station without inter-network interference. Even when there is no inter-network interference, the transmissions of the nodes from the own network can cause collisions. In our previous study without interference, we have already seen that using multiple base stations is an efficient way to scale LoRa networks [2]. Note that from a practical point of view, it is also easier to deploy multiple base stations than to equip sensor nodes with directional antennae, in particular for a sub-1 GHz frequency where antennae are larger in size than antennae for higher frequency bands. Combining both methods (additional base stations and directional antennae) is theoretically possible but seems an impractical choice.

5 Related Work

There is limited published work discussing interference issues and scalability of LoRa. Closest to this paper is the work by Petäjäjärvi et al. who present an evaluation of LoRa link behaviour in open spaces [9]. In another paper [10], the same authors evaluate the coverage and reliability of a LoRa node operating close to a human in an indoor area. The authors also analyse the capacity and scalability of LoRa in a more general approach [6] but in contrast to our work [2] this seems to be based mostly on the theoretical data rather than real-world calibrated simulations as we do [3]. Georgiou and Raza [5] show that the performance drops exponentially as the number of end-devices grows, similar to what we have seen in our previous work [2]. None of these previous efforts considers the case of interference from co-located, non-cooperating LoRa networks.

Saifullah et al. present SNOW [12], a long-range sensor network that operates on white spaces and in many aspects is similar to LoRa. They present a distributed implementation of OFDM that allows them to decode a large number of concurrent transmissions. In contrast, we assume a base station that has more constraints which limits scalability.

Directional antennae have been widely used in cellular and other wireless networks [4]. In the context of wireless sensor networks, Mottola et al. show how only minor modifications to an existing protocol originally designed for omnidirectional antennae can bring performance improvements when using directional antennae [7]. Varshney et al. use them to improve the performance of bulk transfers [13].

6 Conclusions

In this paper we have evaluated the impact of inter-network interference on LoRa networks. Through simulations based on real experimental data, we have shown that interference can drastically reduce the performance of a LoRa network. Our results demonstrate that directional antennae and using multiple base stations can improve performance under interference. Our simulations show that deploying multiple base stations outperforms the use of directional antennae.

8 References

[1] SimPy – Event discrete simulation for Python. `https://simpy.readthedocs.io`. Accessed: 24-05-2016.

[2] M. Bor, U. Roedig, T. Voigt, and J. Alonso. Do LoRa Low-Power Wide-Area Networks Scale? In *Proceedings of the Conference on Modeling, Analysis and Simulation of Wireless and Mobile Systems (ACM MSWiM)*, 2016.

[3] M. Bor, J. Vidler, and U. Roedig. LoRa for the Internet of Things. In *Proceedings of the 2016 International Conference on Embedded Wireless Systems and Networks*, EWSN '16, pages 361–366, USA, 2016. Junction Publishing. ISBN 978-0-9949886-0-7.

[4] H. Dai, K.-W. Ng, M. Li, and M.-Y. Wu. An overview of using directional antennas in wireless networks. *International Journal of Communication Systems*, 26(4):413–448, 2013.

[5] O. Georgiou and U. Raza. Low power wide area network analysis: Can LoRa scale? *arXiv preprint arXiv:1610.04793*, 2016.

[6] K. Mikhaylov, J. Petäjäjärvi, and T. Haenninen. Analysis of capacity and scalability of the LoRa low power wide area network technology. In *22th European Wireless Conference*, pages 1–6, 2016.

[7] L. Mottola, T. Voigt, and G. Picco. Electronically-switched directional antennas for wireless sensor networks: A full-stack evaluation. In *The IEEE Communications Society Conference on Sensor, Mesh and Ad Hoc Communications and Networks (IEEE SECON)*, 2013.

[8] M. Nilsson. Spida: A direction-finding antenna for wireless sensor networks. In *Proceedings of the Workshop on Real-World Wireless Sensor Networks (REALWSN)*, pages 138–145. Springer, 2010.

[9] J. Petäjäjärvi, K. Mikhaylov, A. Roivainen, T. Hanninen, and M. Pettissalo. On the coverage of LPWANs: range evaluation and channel attenuation model for LoRa technology. In *ITS Telecommunications (ITST), 2015 14th International Conference on*, pages 55–59, Dec 2015. doi: 10.1109/ITST.2015.7377400.

[10] J. Petäjäjärvi, K. Mikhaylov, M. Hämäläinen, and J. Iinatti. Evaluation of LoRa LPWAN technology for remote health and wellbeing monitoring. In *2016 10th International Symposium on Medical Information and Communication Technology (ISMICT)*, pages 1–5. IEEE, 2016.

[11] T. S. Rappaport et al. *Wireless communications: principles and practice*, volume 2. Prentice Hall PTR New Jersey, 1996.

[12] A. Saifullah, M. Rahman, D. Ismail, C. Lu, R. Chandra, and J. Liu. SNOW: Sensor network over white spaces. In *Proceedings of the International Conference on Embedded Networked Sensor Systems (ACM SenSys)*, 2016.

[13] A. Varshney, L. Mottola, M. Carlsson, and T. Voigt. Directional transmissions and receptions for high-throughput bulk forwarding in wireless sensor networks. In *Proceedings of the 13th ACM Conference on Embedded Networked Sensor Systems*, 2015.

[14] T. Voigt, L. Mottola, and K. Hewage. Understanding link dynamics in wireless sensor networks with dynamically steerable directional antennas. In *European Conference on Wireless Sensor Networks*, pages 115–130. Springer, 2013.