Improving LP-WAN performance in Dense Environments with Practical Directional Clients

ABSTRACT

Directional antennas are a promising solution for improving the range of client devices in wireless networks. Unfortunately, in LP-WAN systems directional antennas tend to be both large and expensive due to operating at sub-GHz frequencies. However, if a client device is willing to forgo improvements in antenna gain, it is possible to realize compact and low-cost antennas that provide spatial diversity control. In this paper, we show that, by increasing spatial diversity in LP-WAN clients with limited (or no) client gain, we can dramatically increase overall network capacity and improve client battery life by avoiding re-transmissions. This type of directional control can also be used for hot-spot management by more effectively load balancing clients across gateways.

We performed a sensitivity analysis in simulation to explore the impact of various switchable antenna geometries on network capacity under a variety of deployment configurations. We then designed and evaluated three prototype multi-sector array clients: (1) a switchable patch antenna configuration, (2) a digital phase-shift nulling configuration, and (3) a low-cost switched PCB element phase-shift system. We see that our real antenna beam patterns, captured in an anechoic chamber, perform in a similar manner to our simulated predicted models with up to 28% improvement in capacity from interference isolation alone and up to 95% when offloading hot spots between four gateways. We also perform a small measurement study of how often our final design changes its configuration when deployed over multiple days on a campus testbed.

KEYWORDS

Low-Power Wide-Area Network (LPWAN), Cloud Computing, Interference Mitigation, Co-existence

ACM Reference Format:

. 2022. Improving LP-WAN performance in Dense Environments with Practical Directional Clients. In *Proceedings of The 20th ACM Conference on Embedded Networked Sensor Systems (ACM Sensys'22).* ACM, New York, NY, USA, 13 pages. https://doi.org/XXXXXXX.XXXXXXX

1 INTRODUCTION

Low-Power Wide-Area Networks (LP-WANs) are a promising solution for low data-rate Internet of Things (IoT) applications where base stations cover several kilometers supporting clients with multiple-year battery lives. As these systems are adopted in utility sensing, traffic monitoring, and other urban infrastructure applications, they will face significant challenges in terms of client density and the total number of nodes. One common approach for increasing network capacity is using directional antennas to improve spatial reuse. Directional antennas focus radio frequency (RF) signals, allowing clients to transmit farther, at lower power, and with more stable links. Unfortunately, at the sub-GHz frequencies used in most LP-WAN technologies, for client devices to adopt directionality, it would require large (tens of cm) and expensive antenna arrays.

One often overlooked attribute of directional antennas is that, by focusing the RF energy, the client is naturally reducing interference with other nearby nodes and base stations. In this paper, we explore the impact of increased spatial diversity on client devices in LP-WANs in terms of overall network capacity as opposed to just client gains. Given the nature of LP-WAN networks, where there are thousands of clients per gateway and long transmission distances, we see that reducing interference among neighbors not only has a significant impact on overall system capacity but also reduces message retries that are costly in terms of battery life. Our key insight is that it is possible to create low-cost and compact directional designs at the expense of gain. Normally, this is a counterintuitive operating point in antenna design for a client, since it requires adding complexity with no increase in range or transmit power. However, in an altruistic LP-WAN context, this capability has a significant impact on overall network performance. We demonstrate that it is possible to create a variety of directional antenna designs that have similar coverage to an omnidirectional antenna while remaining simple and compact. By reducing off-axis interference alone, we can improve overall LP-WAN capacity by as much as 28%. We also highlight that at overloaded (i.e., hotspot) base stations, directional antennas can be used to more effectively shed load to alleviate congestion.

We explore the potential of client-side spatial diversity in LP-WAN systems by first simulating a variety of city-scale scenarios with different antenna geometries. Our goal is to understand the trade-off between the number of sectors and antenna gain on overall network capacity. As part of this effort, we design and evaluate three generations of antenna beam steering hardware with a software scheme that is protocol compatible with LoRaWAN networks: (1) a rather expensive switchable patch antenna array, (2) a digitally controllable phased array for dual antenna nulling, and (3) a compact and low-cost switchable PCB trace delay phase array. For our switchable patch array, we used a miniature (undersized) ground plane to reduce size. Normally, the ground plane would need to be larger than 10 cm to operate efficiently, so by using smaller ground planes we trade-off size for antenna gain. We passed the measured antenna pattern from an anechoic chamber into our simulator to validate the expected real-world performance. While quite faithful to simulation, patch antennas are relatively expensive and mechanically awkward. We assume clients do not move as is the case with many more traditional WSN applications. Based on our simulation results, we find that there are benefits from having as few as two sectors, but that in terms of size, weight, power, and

ACM Sensys'22, November 07–09, 2022, Boston, United States

ACM ISBN 978-1-4503-XXXX-X/18/06…$15.00
https://doi.org/XXXXXXX.XXXXXXX

cost (SWAP-c), we see that beyond a four or five switchable sector configuration yields diminishing returns. To reduce cost and complexity, we then explore using a nulling controlled interference from a programmable phase difference between two (or more) offset Isotropic antennas. Nulling is a well-established technique for reducing electromagnetic interference but is more difficult to design and often does not increase gain as significantly as other directional techniques. Steerable nulling systems require the ability to change the phase of a signal feeding a set of offset antennas. This can be done with phase-shifting circuits or simply by switching the RF path through longer or shorter PCB traces (discussed next). We design and evaluate a digitally controllable phase offset front-end as a proof-of-concept, called NulLoRa v1, that provides fine-grained control of phase. Unfortunately, digital phase-shifting hardware is expensive, adds insertion loss, and is power-hungry. Our nulling-based system did not match our simulated beam pattern predictions as well as the sector design, but in some cases, the overall network capacity was larger than expected due to deeper nulls. Even with irregular and more difficult to predict beam patterns, when a node selected the strongest primary signal to a gateway, the deeper nulls dramatically reduced neighbor interference. Finally, we brought all of our design lessons together and built a discrete (as opposed to digitally controllable) phase delay system with RF switches and a PCB delay path called NulLoRa v2. We again employ the same anechoic chamber antenna trace-driven simulation of our nulling and diversity antenna design and compare it against the simulated predictions. In our simulations, a switchable nulling antenna with just four modes outperformed our switchable patch antenna design at a lower cost and with less mechanical complexity.

One of the benefits of our final design is that, with multiple antennas, it is also possible to leverage pure antenna diversity by entirely selecting a single Isotropic antenna (albeit at the cost of neighbor interference). Through various RF mapping studies with LoRa radios on campus, we found that across half a wavelength in distance (about 8cm) nodes experience on average 6 dB of difference due to fading and shadowing. As with most other wireless systems, simply selecting the best antenna often provides significant gains. In practice, this means that a client node that has a weak connection to a gateway has the option to simply use antenna diversity without additional phase delay to boost reception or active nulling.

Another natural benefit of steerable antennas is the ability to reduce the load on hotspot gateways that have significantly more client load compared to neighboring base stations. LP-WAN networks can already associate nearby nodes with other base stations, but the interference still dramatically reduces throughput. We extend our simulation work to show the potential of alleviating hotspots with and without steerable antennas. In the most simple case with two gateways, it was possible to improve capacity by over 95%.

It is worth mentioning that all of these techniques can be applied to gateways as well as clients. We focus on clients since the majority of the traffic is upstream. Gateways can also afford more complex RF front-ends like MIMO arrays. Though our work is largely trace-driven, we also perform a small measurement study of how often our final design changes its antenna selection when deployed over multiple days on a campus testbed. We did see that the best antenna choice changed over time on the order of a few times per day. This indicates that electronic steering is useful compared to just static placement.

In summary, our main contributions are as follows:

(1) We provide open-source extensions to an NS3 LoRaWAN simulator to support high-fidelity antenna simulations. We use the simulator to provide a sensitivity analysis of the impact of antenna geometry (i.e. number of switchable sectors) on overall network capacity;
(2) We design and evaluate a LoRa client with a compact switchable sector antenna that has the gain of an omnidirectional antenna by the ability to dramatically reduce neighbor interference;
(3) We design and evaluate a potentially more practical and lower-cost phase-based adaptive steerable nulling RF front-end that outperforms our switchable sector design;
(4) We provide an initial proof-of-concept that directional antennas can aid in hotspot offloading even without increased gain.

2 BENEFITS OF DIRECTIONALITY IN LORAWAN

Different LP-WAN implementations have unique architectural trade-offs that impact hardware and deployment costs. The cellular industry uses a licensed spectrum for LP-WAN, thus minimizing interference and ensuring QoS. Other standards such as LoRaWAN and SigFox, operate in an unlicensed ISM spectrum that is often easier to deploy, but its shared nature can lead to performance concerns, especially in densely populated areas. This is already a problem in systems using other shared ISM bands such as WiFi, but the lower frequency and scale of LP-WAN exacerbate these challenges. Throughout this work, we use LoRaWAN because it is available as an open-source platform that works on the unlicensed spectrum. But the same techniques can be applied independently or in combination with other device-centric networks on the market (i.e. Sigfox, Weightless [7, 27]). A more detailed overview of LoRa and LoRaWAN can be found here [15].

LoRaWAN networks consist of a set of well-provisioned gateways that are usually powered endpoints on a broadband network. These gateways often listen on multiple channels simultaneously and can benefit from high-performance (potentially directional) antennas. Each gateway can support thousands of low-cost battery-operated clients. Clients are low-cost, resource-constrained devices that transmit asynchronously to a selected gateway. A negotiation protocol at startup is used to establish the minimum transmit power and spreading factor for the target gateway, but each message can be overheard by other nearby gateways to improve reliability. The Spreading Factor controls how much coding is applied to a packet allowing clients to trade off airtime for range and reliability (high spreading factors transmit farther but take significantly longer). Clients are assumed to use Isotropic antennas with the combination of power control and spreading factor to reduce interference on neighboring gateways. Upstream and downstream traffic uses different channels, so client upstream traffic should not interfere with client downstream traffic.

This architecture leads to the interesting trade-off where reception by multiple gateways increases the reliability of a single packet

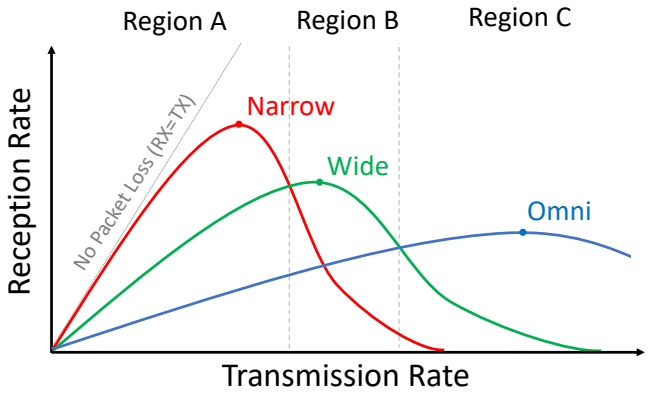

Figure 1: Illustration of idealized network performance given load with different spatial diversity capabilities.

at the cost of increasing interference globally. In a local sense, transmitting a packet with higher transmit power should increase reliability. In a global sense, high-powered transmissions increase interference with other gateways, which decreases network capacity. This is similar to the principle behind CSMA network capacity [31] except now in the context of multiple receiver gateways. This scenario becomes more complex when one considers the energy cost of retransmitting lost packets in LoRaWAN. In practice, we see scenarios where the network capacity could be quite high, but nodes are transmitting multiple times which leads to poor energy efficiency. Control over spatial diversity allows clients to go beyond power control and spreading factors to direct RF energy toward a gateway of their choice. In theory, it is possible to leverage spatial diversity to allow for higher power transmissions at lower spreading factors without interfering with neighbors.

2.1 Spatial Diversity and Nulling

In this paper, we present two mechanisms for controlling spatial diversity. The first is switchable sector antennas that consist of an RF switch that can cycle through several patch antennas tiled in different directions. The patch antennas slightly overlap in terms of coverage with their adjacent antennas to avoid holes in coverage. In our designs, we assume that as the number of sectors increases the width of each sector decreases to increase pointing resolution. The second mechanism we explore is the use of phase-offsets in an array of antennas to programmatically create nulls in the antenna radiation pattern. In this case, the antenna can be configured to adjust where maximal power output should go with deep nulls in off-axis directions. Unlike the switchable patch antennas, nulling antenna configurations tend to be less regular. In Section 4, we discuss how we design and evaluate these different configurations.

Figure 1 shows an ideal illustration of how we would expect spatial diversity to impact overall network performance. On the x-axis, we see transmission rate as a measure of load on the network. The y-axis shows the reception rate to indicate how many unique packets are received by the overall network across all gateways (i.e., multiple copies of the same packet only count as a single packet received). In an ideal environment, the transmission rate would match the reception rate shown by the "No Packet Loss" line. In practice, contention and collisions result in packet loss. We show three different antenna configurations with their respective curves that define network performance. The peak of each curve can be thought of as the maximum network capacity (peak goodput) for any configuration. The first line (shown in red) defines what we might expect from a network where all the client devices have fine-grained *Narrow* controllable antenna beams. This could be achieved in any number of ways including MIMO beamforming [20], switchable sectors, etc. It is no surprise that a network with highly directional ideal clients has the highest network capacity. If we look at the second *Wide* line (shown in green), we see the performance of the network if the clients had wider beam patterns. Keep in mind that the amplitude and position of these curves are highly dependent on several factors such as antenna gain and side-lobe interference. Generally, we might expect that all things equal the network capacity would be lower compared to the Narrow system. The last line (in blue) represents how the network would respond with omnidirectional antennas on each client.

One interesting aspect to note is that peak capacity might not tell the entire story about what is important for clients. As we can see, there are three regions labeled A through C that indicate different levels of increasing transmission rate. If the transmission rate is increasing, but the peak reception rate is lower, that indicates that packets are getting lost due to contention and require re-transmissions. This implies that each client is expending more energy per unique packet. In this example, the Narrow clients have both a higher peak reception rate and the peak occurring in Region A compared to the Wide clients that have a lower peak that occurs in Region B. It would be possible however for Wide to have a potentially equal or higher peak in Region B compared to the Narrow clients, but that would indicate this is coming at a penalty in terms of client energy (due to retries). Region C shows an example where the omnidirectional clients are still able to get data through compared to the Narrow and Wide clients that have already reached an overload state where most devices are constantly colliding and retrying. Normally this would be resolved with better MAC tuning like in p-persistent CSMA [31], but this is difficult in practice in LoRaWAN networks where each client is device-centric making mostly local decisions without a constant central coordinator.

Through the rest of this paper, we will explore how different mechanisms change the network capacity of various client and gateway configurations. We will highlight the trade-offs in terms of network capacity and average client energy to help inform the influence of various techniques under different scenarios.

3 RELATED WORK

There has been a significant body of work around improving the performance of LP-WANs, in particular LoRaWAN at multiple layers of the network stack [10, 13, 19, 25, 26]. At the physical layer, Charm [9] exploits the observation that, while signals from certain clients may attenuate significantly, they are still likely to be received by multiple gateways in a dense network. Charm introduces a hardware and software design at the gateways that identify and transports weak received signals to the cloud. This is achieved

through a joint decoding system in the cloud that coherently combines weak signals received across multiple city gateways to decode the underlying data. As a result, Charm both expands the decoding range of the LP-WAN network and improves battery life for nodes already in range – allowing client devices to spend less energy per transmitted bit. OPR [4] uses a software-only solution to recover packets by pooling information from multiple gateways at the link layer, which would otherwise be lost due to short-lived interference to LoRa. OPR works by disabling the packet rejection that normally takes place at LoRa gateways when packets are subject to Cyclic Redundancy Checks (CRC) and/or Forward Error Correction (FEC) errors. Instead of being discarded, these corrupt packets are collected by a network service that groups them based on geographic proximity and reception time. A scheme of XOR-based combination and majority voting is used to estimate the most likely bit error locations. Then, Cloud-based processing is used to find the most likely correct bit combination.

Unlike the previous approaches, which are focused on combining multiple signal versions to improve the SNR, in [18], a scheme is presented for Deep Neural Network (DNN) assisted demodulation, effectively lowering the SNR threshold required to correctly decode a given SF. The scheme employs a dual-DNN model, consisting of a noise filter and a chirp symbol decoder, for noise reduction and chirp symbol decoding, respectively. Most of these approaches at the physical layer can be used in parallel with spatial diversity mechanisms.

Other LoRaWAN optimization works exploit SF orthogonality to increase network capacity by performing adequate SF assignment. In [32], the authors propose an SF assignment method based on the Lagrange multipliers technique, which improves the probability of data delivery, at the cost of a minor increase in the energy consumption of the devices. In [11], SF optimization schemes based on gradient ascent and game theory are proposed to maximize throughput in multi-operator networks. In [24], the problem of SF and transmit power optimization is modeled as an integer linear programming problem in two phases: firstly, the SF of each node is optimized, which can be done base on one of two objective functions, either minimizing the collisions in the most overloaded SF or balancing the load in each SF for all gateways; secondly, the transmit power of each device is optimized. This results in a higher delivery ratio when compared with other techniques, namely the assignment of the minimum SF and airtime balancing between SFs.

During the last decade, the potential advantages of using directional antennas in Wireless Sensor Networks (WSN) have been acknowledged in the literature, mostly in the context of multihop WSNs [12, 14]. The identified advantages encompass security, localization, increased range and link quality, energy efficiency, and interference mitigation. In this context, some smart antenna prototypes of the switched beam type, suitable for operation in the 2 GHz frequency band, were developed and tested. In [6], a switched beam antenna with omnidirectional capability was developed, comprising eight micro-strip antennas composed of a rectangular two-element patch antenna array, and a vertical half-wavelength dipole antenna. Such design was capable of switching between nine antenna patterns, being eight directional (half-power beamwidth of nearly 60°) and one omnidirectional. The same team has later proposed another antenna design comprising four identical antennas, each composed of an array of two L-shaped quarter-wavelength slot antenna elements [5]. Based on this, the switching circuit can select between eight different radiation patterns, with a half-power beamwidth of 45°. In [8], a design is proposed with eight Yagi-Uda antennas, capable of switching in steps of 45°. The main lobe presents a half-power beamwidth of 60°, with a gain of 5.5 dBi, and sidelobes gains 11 dB smaller. In [23], the author proposes the SICS Parasitic Interference Directional Antenna (SPIDA). This Electronically Switched Parasitic Element (ESPE) antenna design consists of a central active element surrounded by "parasitic" elements. The central element is a conventional quarter-wavelength whip antenna. The parasitic elements consist of six legs made of 1-mm copper wire. This allows switching between six different directions, resulting in a smooth radiation pattern with an antenna gain of approximately 4 dBi and a front-to-back ratio of 11 dB. The SPIDA design was recovered in [29], where the measured gain was approximately 6.8 dBi.

The use of directional antennas in WSNs raises important issues, such as head-of-line blocking and directional hidden terminal problems in multihop communication. Specific MAC protocols have been proposed to mitigate these problems, such as in [17]. These problems do not usually apply in LPWANs, since, by default, they are a single hop and rely on ALOHA-like MAC protocols, as in LoRaWAN and SigFox [3, 34]. On the other hand, since, in theory, LPWANs rely on a simple hardware design, relying on omnidirectional broadcasting to reach different gateways for sake of diversity, the use of directional antennas, being an apparent contradiction, deserves a careful evaluation. In [33], SPIDA type directional antennas mounted on the end devices were considered as a way to improve signal strength at the receiving gateway, enabling the capture effect. The study concludes that increasing the number of gateways leads to a more significant performance improvement than using directional antennas. However, no exhaustive evaluation is made of the impact of directional antennas, namely for interference mitigation. In fact, in the presented simulations, only the devices connected to the gateway of interest are equipped with directional antennas, not those belonging to the interfering networks.

By avoiding signal propagation towards null directions, directional antennas can also be used to mitigate interference and increase spatial reuse, which has been researched in the context of infrastructure WLANs. Proposed schemes employ a central controller, which decides on scheduling and AP assignment to user devices, being able to perform load balancing between APs. In [16], the authors propose a Multi-Channel Multi-Sector Directional Antenna Wireless LAN (MCMSDA WLAN). This scheme integrates multi-sector switched beam antennas in user terminals, a TDMA MAC, and a centralized scheduling algorithm to provide load balancing in infrastructure WLANs. The scheduling mixed-integer optimization problem targets the minimization of overall transmission time. In [30], a scheme is proposed, where indoor space is covered by several APs equipped with directional antennas. Based on station location data, a central controller selects the most suitable AP, which provides the highest throughput. Handover in overlapping zones is done based on the direction of movement of the mobile client.

This paper aims to provide an extensive analysis, of both the theoretical and practical use of directional antennas in LoRaWAN end devices. It proposes and evaluates practical directional antenna designs, having in mind the application in LPWAN end devices.

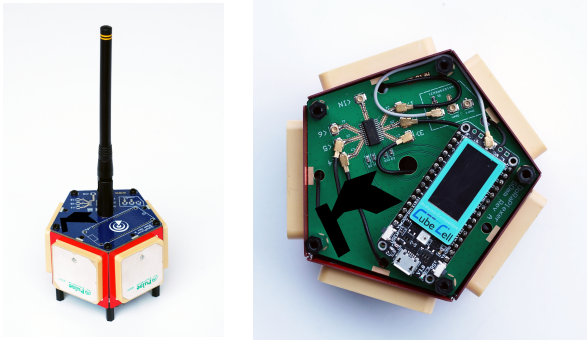

Figure 2: DoRa client patch array prototype is composed of five PulseLARSEN W3215 Ceramic Patch antenna and a Linx ANT-916-PML whip antenna. Antenna selection is performed by the Analog Devices HMC252A SP6T RF switch using GPIO from a CubeCell HTCC-AB02 LoRaWAN end-device (shown on bottom).

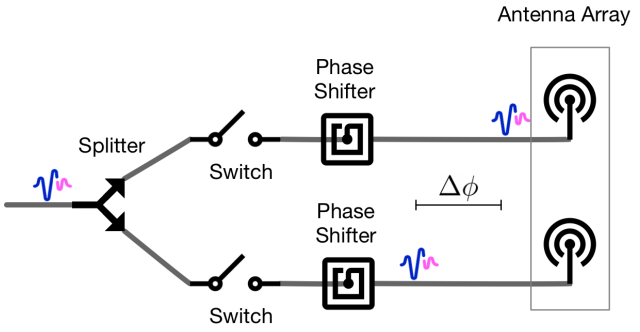

Figure 3: Block diagram for adaptive steerable nulling.

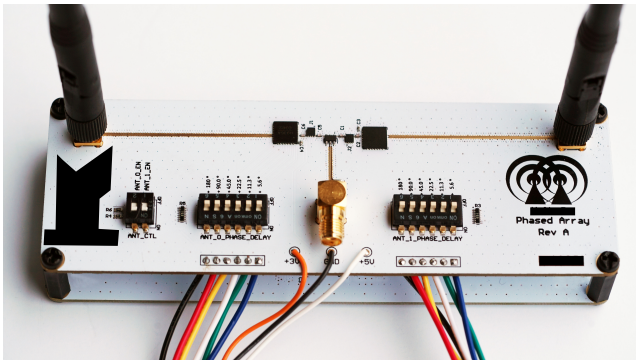

Figure 4: NulLoRa programmable nulling front-end prototype is composed of a Mini-Circuits BP2C+ Power Splitter/Combiner, two Skyworks SKY13347-360LF RF Switches (diversity control), two Analog Devices HMC936ALP6E 6-bit digital phase shifters and two Linx ANT-916-CW-HWR-SMA whip antenna.

The performance of these designs is evaluated based on prototype implementations. Computer simulation is used to evaluate the performance and benefits of these antennas in larger-scale deployments, including the performance gains of gateway hotspot offloading. To the best of our knowledge, this is the first study specifically focused on this topic.

4 SYSTEM DESIGN

In this section, we propose two distinct ways of reducing interference and increasing the overall network capacity of a LoRa network. We validate and optimize our solutions using the NS3 [28] discrete-event network simulator and build hardware prototypes of a directional antenna array (DoRa) and a nulling array front-end (NulLoRa)[1]. We also fine-tune our initial simulation models with real measurements captured in an anechoic chamber and compare their performance at scale under several different gateway and client configurations.

4.1 DoRa: Directional Patch Array

Traditional directional antennas come with many benefits (increased gain and directivity), but also have drawbacks, such as the size, especially when operating on sub-GHz bands, where antenna geometries are a function of wavelength. The use of Ceramic as a substrate in a ceramic patch antenna enables it to act as electrically larger when compared to its operating wavelength. Compared to an omnidirectional monopole antenna, the peak gain of a ceramic patch can be as high as 4.5 dBi [2] at sub-GHz bands, but to achieve this theoretical gain, large ground planes are required.

To keep our design compact, we opted to sacrifice the maximum achievable gain while maintaining beamwidth performance. Before designing our prototype, we evaluated the trade-off between the number of sectors and beamwidth, and its impact on network capacity (Section 5.1). Figures 5a and 5b show the estimated radiation pattern of our idealized 0 dB gain 5 patch array and the anechoic measurements from our prototype (shown in Figure 2).

[1]We will be releasing both open-source designs with Gerber files on GitHub

The prototype also includes an omnidirectional monopole that can be switched on as a fallback in cases when a deployed protocol doesn't support antenna selection, and it is also used as the baseline for our current systems.

We make use of an Analog Devices HMC252A SP6T RF switch to select the active patch antenna. When active, the switch adds 11 mW, which is negligible compared to the RF transmission but does add 0.8 dB of insertion loss. For this prototype, we make use of Pulse LARSEN W3215 ceramic patch antennas. The total Bill-of-Materials cost for the antennas and the switch is about $80 US (clearly cost-prohibitive for production systems). As seen in Figure 5b, the peak gain of each sector varies from 0.3 dB and -1.5 dB in the direction of interest.

4.2 Nulling Prototype

As an alternative to ceramic patch antennas, we explore the idea of using nulling based on a two Isotropic antenna array with controllable phase offset between them. A relatively large number of patterns can be achieved by varying the phase and distance between the two omnidirectional antennas. In order to reduce the number of antenna elements and keep the design compact, there

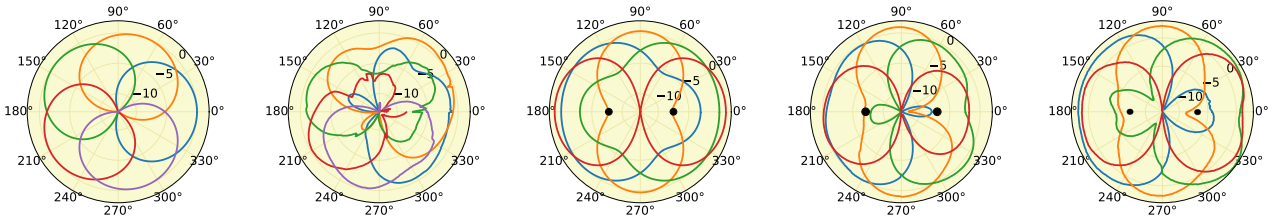

(a) Ideal five sector directional patch antenna with a beamwidth of 110 deg and 0 dB gain.

(b) Measurements from our five sector directional patch antenna prototype.

(c) Select patterns used in our two antenna nulling-array simulations.

(d) Measurements from our two antenna nulling-array prototype.

(e) Measurements from our two antenna passive nulling-array.

Figure 5: Antenna gain (in dBi) radiation patterns used in our simulations, derived from mathematical models (a, c) and compared to the measurements taken in an anechoic chamber (c, d, e). Black dots show monopole locations.

is a limitation in directivity that such a system can achieve. Three (or more) antenna elements could have been used to improve gain and directivity but this comes at a significant increase in cost and size. We opted for two Isotropic antennas since this would also enable more traditional diversity through non-phase-shifted antenna selection.

Figure 3 shows a block diagram and the functional parts of our nulling Prototype. It works by splitting/combining the signals between our transceiver and the two antennas while passing through programmable digital phase shifters. Figure 4 shows our Nulling prototype front-end that can accept the signal from any LoRa client via an SMA connector, which then transmits a nulling pattern controlled by some GPIO pins. The prototype board allows us to generate phase-offsets between the two antennas with up to 6° resolution and a range from −354° to 354°. We used the Skyworks SKY13347-360LF as an RF switch along with HMC936ALP6E chips to allow for fine-grained digital control of the phase going into each antenna. Though not ideal in a production system, this front-end allowed us to easily experiment with different nulling patterns. Each digital control pair costs $85 US, consumes 3mW and adds 5.0dB of insertion loss.

Even though we can generate a larger number of radiation patterns, in practice a limited number of patterns that can replicate the effective cumulative gain of Isotropic antenna equivalent is sufficient (discussed in Section 6). We used the MATLAB Antenna Toolbox [1] to derive the separation distance between antennas that minimizes the number of patterns, as well as maximizes their nulling regions. Figure 5c shows the derived four nulling patterns by separating the two antennas by $0.37\lambda \simeq 12$ cm at −90°, 0°, 90° and 180° of phase-offsets. Our hardware prototype was also configured based on this configuration, and the radiation pattern of each phase offset is shown in Figure 5d.

Our NulLoRa nulling front-end also includes a pair of RF switches that can be used to disable one of the antennas. This allows to selectively choose the best antenna in a particular RF environment, which is also known as antenna Diversity. We compare the performance gains of antenna diversity with nulling in Section 5.2.

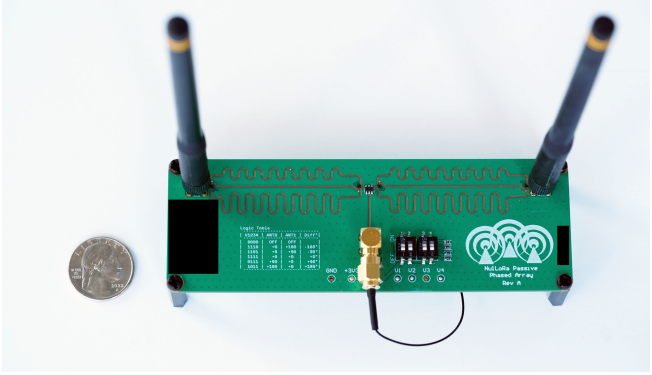

Figure 6: Our final Passive version of the NulLoRa programmable nulling front-end is composed of a Mini-Circuits BP2C+ Power Splitter/Combiner and four Infineon BGS13S4N9 SP3T RF switches. It uses two Linx ANT-916-CW-HWR-SMA whip antenna.

4.3 NulLoRa: Passive Nulling Antenna Array

We believe that cost plays a major factor in Internet of Things (IoT) hardware adoption. Thus, we have developed a passive version of the Nulling Antenna Array. As shown in Figure 6, we have replaced the expensive phase shifters used in the prototype with discrete trace delays and SP3T RF switches. NulLoRa adds 1.1 mW of energy consumption when active, and adds 3.5 dB of insertion loss (which gets offset by the gain in the pattern beam forming). The passive Nulling Antenna Array design can be added to any commercial LoRaWAN design for as little as $2.96 in parts.

4.4 Anechoic Chamber Experiments

We performed back-validation of our prototype designs against our simulations and commercially available hardware in an anechoic chamber. As shown in Figure 7, the device-under-test (DUT) is placed on a computerized rotary table that measures the received signal strength under different incident angles (0° to 359° azimuth and −45° to 45° elevation at 1° increments) and across the US915 ISM band. For the simplicity of data interpretation in this paper, simulation antenna models and anechoic measurements use the

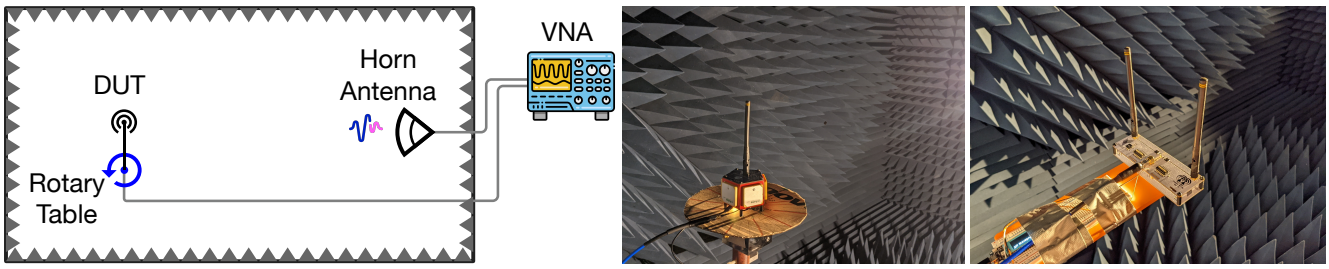

Figure 7: Anechoic chamber experimental setup (left), DoRa five-sector board (middle) and Nulling front-end (right)

values at the interception of the horizontal plane (0° elevation) and consider the average gain across the US915 ISM band (not the peak Gain).

When averaged across the US915 band, and due to hardware imperfections, we can observe the radiation patterns of our patch array (Figure 5b) and nulling array (Figure 5d) have the same effective gain as an omnidirectional antenna with 0dB gain (we use a Linx ANT-916-CW-HWR-SMA Whip antenna as our benchmark).

4.5 Client Antenna Selection Procedure

Client devices periodically scan the status of the surrounding gateways, by cycling through a ping message on each antenna. We extend the current LoRaWAN association protocol used for Adaptive Data Rate (ADR) to include cycling through multiple antennas (similar to mark and sweep algorithms). When a client is cycling through antennas, each gateway keeps track of the mean packet reception ratio (PRR) and RSSI of a sequence of messages for each client. The closest gateway coordinates with other nearby gateways and returns the average PRR and RSSI values to the client in an appended downstream message. This coordination assumes an interconnect between nearby gateways which is not uncommon given LoRaWAN's MAC in the cloud architecture. Gateways that have low PRR or RSSI values do not send their information to clients. The device can then use the set of gateways, PRR, and RSSI values to switch to the antenna where the highest PRR was achieved that was above a certain minimum RSSI threshold. The scan operation can also be triggered if the device stops receiving any response in the currently selected antenna, or if the RSSI reported by the gateway drops below the minimum RSSI threshold for a number of sequential packets. The final association decision can be made by the client or pushed down by the gateway depending on the architecture. Standard LoRaWAN systems would allow the client to decide while systems that employ hotspot offloading might leverage more information from the network. This protocol can easily be implemented as an application layer scheme in the context of LoRaWAN.

4.6 Hotspot Offloading Algorithm

Besides interference isolation, beam selection can also be used to perform hotspot offloading, directing transmissions away from saturated gateways, thus improving load balancing. The challenge here is to make sure that the underlying signaling is scalable and converges in a manner where nodes don't oscillate from one gateway to another unnecessarily. In order to validate the potential of hotspot management, we adopt a simple centralized algorithm, which is formally described in Algorithm 1. The algorithm receives as input parameters a set of N gateway descriptors, $\Gamma \triangleq \{\gamma_n\}_{n=1}^{N}$, a set of M end device descriptors, $\Delta \triangleq \{\delta_m\}_{m=1}^{M}$, and a set of K antenna pattern descriptors, $H \triangleq \{\eta_k\}_{k=1}^{K}$. The output is the

Algorithm 1 Hotspot Offloading Algorithm

Input: Gateway descriptors $\Gamma \triangleq \{\gamma_n\}_{n=1}^{N}$, Device descriptors $\Delta \triangleq \{\delta_m\}_{m=1}^{M}$, Antenna pattern descriptors $H \triangleq \{\eta_k\}_{k=1}^{K}$
Output: Antenna pattern assignment $\Pi \triangleq \{(\delta_m, \pi_m) : \delta_m \in \Delta \wedge \pi_m \in H\}_{m=1}^{M}$

for all $m \in \{1, ..., M\}$ **do**
  $\pi_m \leftarrow \arg\max_{\eta_k} [\max_{\gamma_n} \text{RSSI}(\gamma_n, \delta_m, \eta_k)]$
end for
$i \leftarrow 0$
$toa\_diff \leftarrow \max_{\gamma_n \in \Gamma} \text{ToA}(\gamma_n, \Delta, \Pi) - \min_{\gamma_n \in \Gamma} \text{ToA}(\gamma_n, \Delta, \Pi)$
$\Omega \leftarrow \Gamma$
while $\Omega \neq \{\} \wedge \frac{pr\_toa\_diff}{\max_{\gamma_n \in \Gamma} \text{ToA}(\gamma_n, \Delta, \Pi)} \geq C_1 \wedge i < C_2$ **do**
  $g \leftarrow \arg\min_{\gamma_n : \gamma_n \in \Omega} \text{ToA}(\gamma_n, \Delta, \Pi)$
  $Z \leftarrow \{(\delta_m, \eta_k) : \delta_m \in \Delta \wedge \eta_k \in H : \eta_k = \arg\max_{\eta_k \in H} \text{RSSI}(g, \delta_m, \eta_k) \wedge \text{RSSI}(g, \delta_m, \eta_k) \geq RS_{SF12} \wedge \eta_k \neq \pi_m$
  if $Z \neq \{\}$ **then**
    $(u, v) \leftarrow \arg\min_{(m,k):(\delta_m, \eta_k) \in Z} \sum_{n=1}^{N} \mathbb{1}_{\{\text{RSSI}(\gamma_n, \delta_m, \eta_k) \geq RS_{SF12}\}} \cdot \text{RSSI}(\gamma_n, \delta_m, \eta_k)$
    $\Pi_{new} \leftarrow (\Pi \setminus (\delta_u, \pi_u)) \cup \{(\delta_u, \eta_v)\}$
    $toa\_diff\_new \leftarrow \max_{\gamma_n \in \Gamma} \text{ToA}(\gamma_n, \Delta, \Pi_{new}) - \min_{\gamma_n \in \Gamma} \text{ToA}(\gamma_n, \Delta, \Pi_{new})$
    if toa_diff-toa_diff_new>0 **then**
      $\Pi \leftarrow \Pi_{new}$
      $toa\_diff \leftarrow toa\_diff\_new$
      $i \leftarrow 0$
    else
      $i \leftarrow i + 1$
    end if
  else
    $\Omega \leftarrow \Omega \setminus \{g\}$
  end if
end while

antenna pattern assignment, i.e., a set of ordered pairs associating each device to the respective selected antenna pattern, i.e., $\Pi \triangleq \{(\delta_m, \pi_m) : \delta_m \in \Delta \wedge \pi_m \in H\}_{m=1}^{M}$. The function RSSI$(\cdot,\cdot,\cdot)$ receives as arguments one gateway, one transmitting device, and one specific antenna pattern, and returns the respective RSSI level measured at the gateway. The ToA$(\cdot,\cdot,\cdot)$ function receives as arguments one gateway, a set of devices, and an antenna pattern assignment, and returns the sum of the Time on Air (ToA) values needed for the transmission of fixed-length packets, where each device transmits a single using the lowest possible SF – so that the received power is still greater than the respective receiver sensitivity –, and the selected antenna pattern. This function measures the load of a gateway. Constants C_1 and C_2 represent the target ToA unbalance between gateways, and the maximum tolerated number of rounds without load balance improvement, respectively. RS_{SF12} is the receiver sensitivity of SF12, i.e., the minimum power level that still allows communication in LoRaWAN. Steps 1-6 correspond to the initialization. Steps 1-3 initialize the antenna pattern assignment so that each device is assigned the pattern that maximizes the RSSI towards any gateway. Steps 4-5 initialize, respectively, the counter of elapsed rounds without load balance improvement (i), and the previous load unbalance, measured as the difference between the highest and lowest ToA values associated with the gateways (toa_diff). Step 6 initializes Ω, which is the set of gateways that can still be selected in step (8). Ω starts with all the gateways in Γ. Steps 7-24 constitute the main cycle of the algorithm, which will run while Ω is not empty, the load unbalance fraction is greater than C_1, and the number of elapsed rounds without load balance improvement is lower than C_2. In each round, in step 8, the gateway with the lowest load (g) is selected in order to become the preferential gateway of one more device (i.e., the gateway located in the direction of the main lobe of that device's antenna pattern). In step 9, the set of candidate devices, Z, is filled, comprising devices currently pointing away from g, which can change their antenna pattern to point toward g and achieve an RSSI level of at least RSS_{SF12}. Each candidate device forms a tuple with the respective antenna pattern that maximizes the RSSI level at g. At the end of this step, if Z is empty, gateway g is removed from Ω (step 22), so that the algorithm proceeds to choose the next g from among the remaining gateways in Ω. Otherwise, the algorithm selects the candidate device whose transmission using the selected pattern minimizes the sum of RSSI values at the different gateways within the range and updates its pattern selection (δ_u, π_u) (steps 11-12). In steps 13-20, the algorithm checks if the load balancing was improved, in which case it updates toa_diff, incrementing i otherwise.

A practical realization of hotspot offloading would have to be distributed, for sake of scalability. When nodes are performing their association operation, as described in Section 4.5, gateways can pass an indicator of their current capacity (e.g., based on the ToA metric). Each node can then decide to switch to a gateway with additional slack capacity even if it is not the strongest in terms of signal quality. Such distributed schemes will be explored in future work.

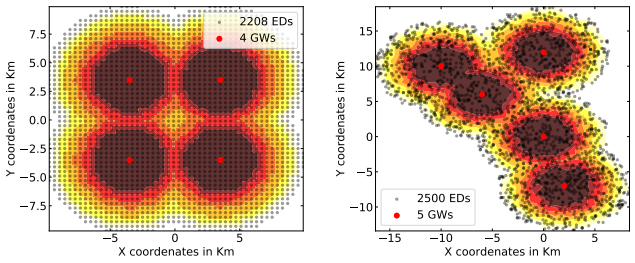

(a) Uniform Layout of 4GWs and 2208 Clients

(b) Random Non-Uniform Layout of 5GWs and 2500 Clients

Figure 8: Example of gateway and client placements used in NS3 simulations. The gradient represent the maximum achievable Spreading Factor (from SF7 to SF12)

5 EVALUATION

To evaluate network capacity, we built upon the NS3 [22, 28] LoRaWAN module [21]. We have added high-fidelity support for directional antennas, antenna arrays, and arbitrary antenna beam patterns that can be dynamically configured and adjusted at runtime[2]. These beam patterns are either based on ideal mathematical models of antennas or recorded radiation patterns from our antennas that we collected in the anechoic chamber. At the start of each simulation, each client performs an antenna gateway selection protocol searching for the antenna elements that yield the highest RSSI. In each test, we place clients at random locations and with random orientations. Clients slowly increase the rate of their traffic, which increases the Global Transmission Rate (TX) load found in many of our performance plots. The Global RX packets give an indicator for unique packets received by the network of gateways (duplicates are ignored). Figure 8 shows an example of two topologies used in our simulations. The first shows four uniformly spaced gateways where the color corresponds to different spreading factor / RSSI levels, with 2208 clients placed on a grid. The second shows a more random gateway layout with 2500 clients also randomly placed within the gateway reach. Unless specified otherwise, we use these typologies with traffic rates scaled up globally to nearly 800,000 packets per hour (222 per second) across the 4 or 5 gateways. Each packet was fixed to 23 bytes long, which corresponds to between 77.1 ms-493.6 ms of on-air time, depending on the selected spreading factor. Each experiment measures between 2,500 and 1.6M packets per test point, depending on the rate and we average the results over 5 test runs. Standard deviation is shown in some plots as shaded regions around each line.

Since our focus is on evaluating how spatial diversity impacts a network as it scales, we were mostly confined to simulation experiments. However, it is worth mentioning that both of our DoRa and NulLoRa (v1 and v2) boards worked well as stand-alone LoRaWAN clients on our campus network. Due to the high complexity, and a large number of external variables and factors we attempt to isolate each advantage of directional antennas and test their performance separately in the following sections.

[2]We plan to release the simulator as open-source software on GitHub with the required scripts to generate our graphs

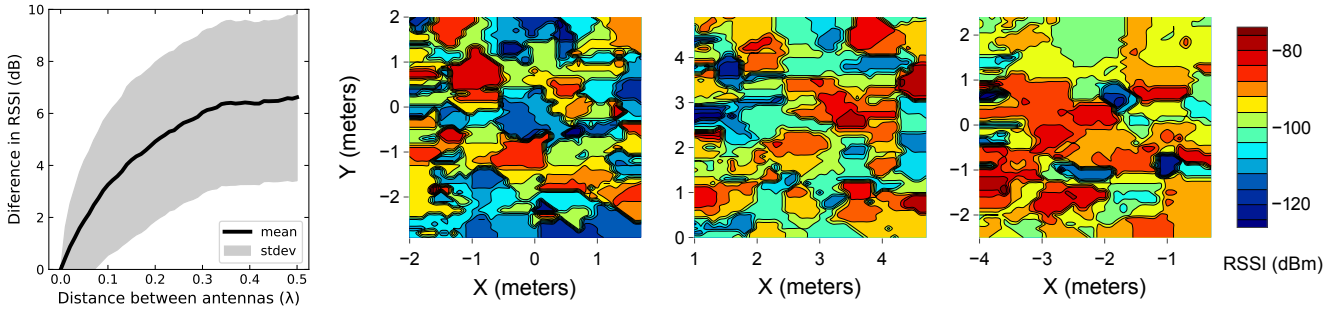

Figure 9: Mean spatial diversity gains from selectively using one of two antennas with the best signal strength. (left) High Resolution RSSI maps of three different regions around campus buildings, showing the large variation caused by fading and shadowing on LoRa networks. (right)

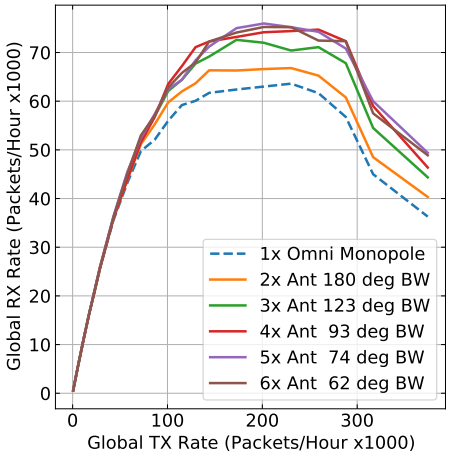

Figure 10: Impact of switchable sectors on network capacity in uniform-gateway topology with random starting orientations.

5.1 Number of Antennas and their Beamwidth

First, we answer the question of how many sectors an end node needs before it gets a diminishing return in terms of overall network capacity when using a directional patch antenna (such as the ones used by our DoRa prototype). Figure 10 shows the network capacity of the network as a function of the increasing number of sectors. As the number of sectors increases, we decrease the effective needed beamwidth of each sector ensuring a 10° overlap in coverage. Using smaller beamwidth sectors would effectively increase the gain of a ceramic patch antenna, although we normalize them to the same 0 dB gain as the omnidirectional antenna, to observe the effect of interference isolation provided by the increase in sectors. We have tested these antenna configurations using both layouts shown in Figure 8 and averaged the values over five runs. The lowest capacity is the default omnidirectional antenna (at the bottom), with 3 or 4 sectors reaching a similar maximum network capacity of up to 19% more than the baseline. In practice, we chose five sectors because they appeared slightly more stable in terms of performance, and mapped nicely to commercially available patch antennas.

5.2 Antenna Diversity Gains

With our nulling prototype (NulLoRa v1 and v2), we also have the option of simply selecting one antenna or another to benefit from a small amount of spatial diversity even with an Isotropic antenna. These gains result from the antenna moving from a peak or trough created by fading or shadowing in the environment. In prior LoRaWAN deployments, we had observed differences in signal strength as much as 19 dB from relatively small location changes within campus buildings. In order to better characterize the potential gain relative to antenna distance, we captured fine-grained maps of RSSI at 2 cm increments across 4×4 meters spaces in 6 locations. For each map, we collected 40 thousand measurements uplinks across 4 GWs located on rooftops distributed across campus. We repeated this collection process at different times of the day to get a sense of the signal stability. As one might expect, during quiet periods like the middle of the night, the signals were reasonably stable, but they can change dramatically during the day.

The right plots in Figure 9 show the highly variable nature of the RSSI signal across three different campus spaces. We simulated the placement of two fixed antenna positions across these spaces and aggregated the expected difference in RSSI between the two (simulating a choice of one antenna vs the other). We averaged the data across forty thousand points in each space and varied the distance between the antennas. The average gain and standard deviation are shown in Figure 9. We observe a substantial difference with separations between the antennas with diminishing returns starting from 0.35λ (≃12 cm at 915 Mhz). In general, one can expect almost a 6 dB gain on one antenna compared to another once you reach a large enough antenna baseline. This boost in gain is significant for nodes that might have a weak signal to their nearest gateway and hence could be more important than the community gain by using isolation techniques like nulling (described below).

5.3 Isolation Capacity Gains

Next, we evaluate the end-to-end network performance given our baseline omnidirectional antenna, our five-sector patch antenna, our quad pattern nulling antenna, and an omnidirectional antenna with an average 6 dB of gain to represent what could be achieved through antenna diversity selection. Figure 11a shows the performance of ideal simulated antennas of each class in our uniform

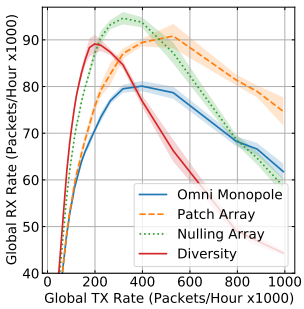

(a) Simulation Results: Isolation capacity gains in a uniformly distributed network (mean of 5 runs with standard deviation range).

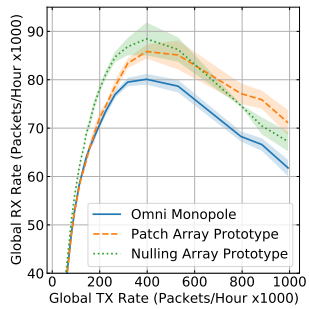

(b) Prototype Hardware Results: Isolation capacity gains in a uniformly distributed network (mean of 5 runs with standard deviation range).

Figure 11: Network capacity given a uniform gateway and client topology (4GWs 2208 EDs).

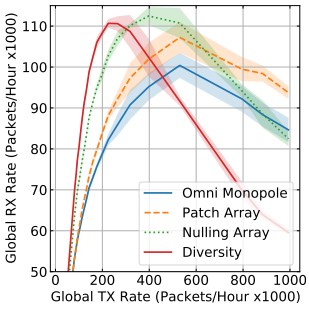

(a) Simulation Results

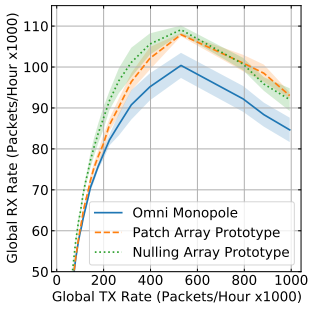

(b) Prototype Hardware Results

Figure 12: Network capacity given a non-uniform gateway and client topology (5GWs 2500 EDs).

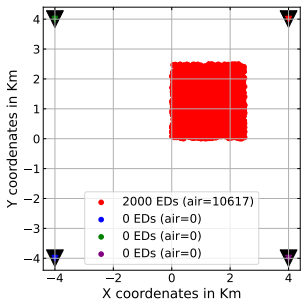

(a) Default Greedy Allocation

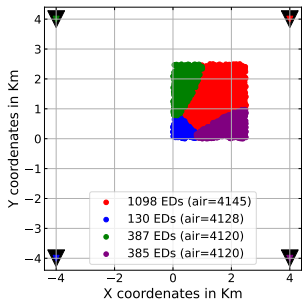

(b) Balanced Airtime Allocation

Figure 13: Comparing End-device gateway allocation with and without Load Balancing Suggestions

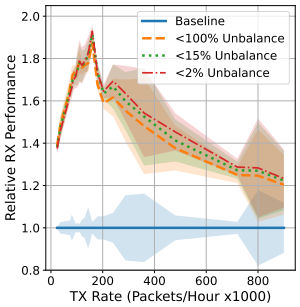

(a) Performance Gains

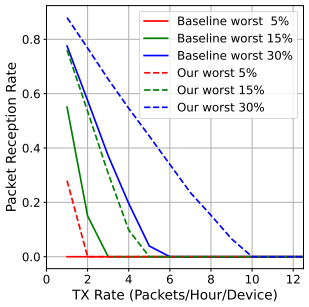

(b) Worst Performers

Figure 14: Performance of Simple End-device gateway allocation with and without Load Balancing Suggestions

environment. We see that Diversity (selecting one of two antennas with a 0.35λ separation) provides a 13% boost over the omnidirectional baseline. This peak occurs early, which implies that, as load increases, there are likely many packets lost in retransmissions. This is to be expected, given that each node generates a large amount of interference with a strong signal arriving at the target gateway. We see that our patch design and the nulling design provide a 13% and 20% increase in capacity, respectively, over the omnidirectional baseline. Nulling has the highest capacity in this configuration by quite a large margin. When comparing plain Diversity to the Patch sector array, we see similar peak capacity, but the Patch array seems to be losing more packets due to collisions. We attribute this to the natural fact that even though Diversity is boosting power at one gateway, it is possible that the packet is not being received as strongly by other gateways. With the sectored antenna, the energy is focused.

Figure 11b shows the performance of our four measured hardware antenna patterns. We get slightly smaller gains of between 10% and 15% with less difference in practice between the sector and null hardware options. All things equal, the nulling antenna could be made more cheaply and provides the flexibility to also leverage standard antenna selection diversity.

We see that the performance gains in terms of percentage seem consistent compared to the uniform topology. Note that the overall throughput is generally almost 10% higher, but there is an additional gateway, which indicates that per gateway performance has decreased. This can be attributed to imbalanced load and ideally, could be further improved with a hotspot load-balancing algorithm (which was not applied in this case).

5.4 HotSpot Offloading Performance

With information about the traffic load at each gateway, it is possible to use our various hardware platforms to redirect traffic to alleviate hotspot gateways. As a simple proof-of-concept demonstration, we set up a simulation topology where we place 2000 client end-devices (EDs) close to one of the four gateways to simulate a dense concentration of nodes (HotSpot). Though not formally defined in the specification, LoRaWAN normally uses a greedy approach where nodes will associate with the gateway that has the highest signal strength. Figure 13a shows this antagonistic topology where, by default, all of the clients would associate with the closer gateway (top right red triangle). This case might look contrived, but it captures a common deployment scenario where many devices are placed around a single gateway, for example, sensors

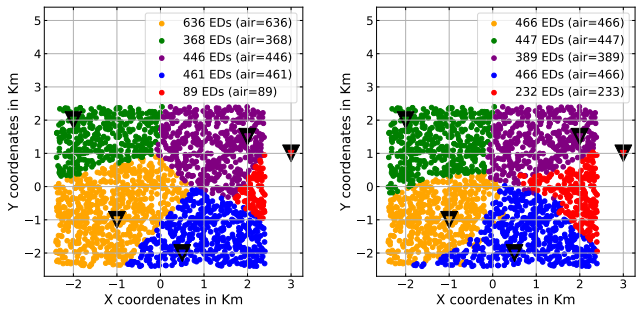

(a) Default Greedy Allocation **(b) Performance Results.**

Figure 15: Common End-device gateway allocation with and without Load Balancing Suggestions

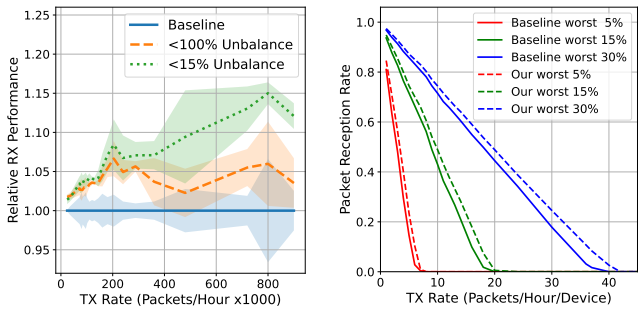

(a) Performance Gains **(b) Worst Performers**

Figure 16: Performance of Common End-device gateway allocation with and without Load Balancing Suggestions

in a large skyscraper. Other gateways nearby may have significantly less load and ideally should help in supporting additional traffic. We make use of the NulLoRa steering capabilities in order to better distribute the traffic across the available gateways. The distribution of clients between the Gateways is accomplished by our Load Balancing Algorithm formally described in Section 4.6 and the balanced allocation is shown in Figure 13b. We observe that the allocation of devices between gateways is proportional to the distance between the clients and their target gateway. This is because LoRa achieves a longer range by using higher spreading factors (and consequently lengthens the transmission time of a message). Figure 14a demonstrates the performance gains associated with performing the allocation. As shown in the figure we compare the baseline greedy allocation with three degrees of *unbalance*. Unbalance is defined as the difference in airtime between the most and the least loaded gateway (i.e. 100% Unbalance represents double the traffic). We can observe a relative improvement in the capacity of up to 90%. We also want to demonstrate that is in the interest of the clients to participate in the load balancing of the network, even if this requires that some nodes use a slower transmission speed and make longer use of the spectrum (as shown by the cumulative air times in figure 13). The directionality of the patterns isolates the interference and allows for better spatial reuse of the spectrum. Figure 14b shows the packet reception rate of the worst performing 5%, 15%, and 30% nodes in the network. By distributing the load across the gateways we minimize the number of collisions, which translates into a better packet reception rate. Figure 15 shows a more common device placement scenario. In this allocation, we have five gateways (shown as triangles) and 2000 devices randomly distributed across a 5×5 Km space. In this allocation, the Gateways are more densely packed and the default allocation does a better job of distributing the nodes. As shown in Figure 16b we are still able to achieve up to 15% improvement over the baseline allocation. Figure 16b demonstrates that we can achieve this without deteriorating the performance of our worst-performing devices, (i.e. all devices benefit regardless of their distance to the gateway). Thus it is fair to state that is in the interest of all network devices to cooperate for a balanced network.

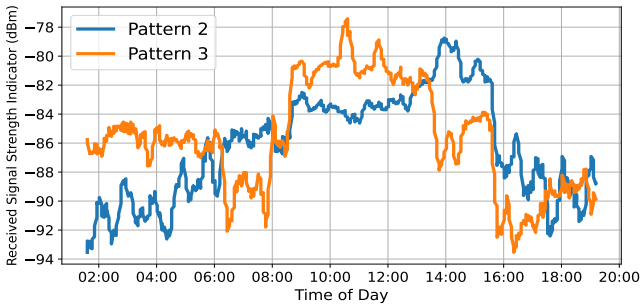

Figure 17: Changes in the Received Signal Strength Indicator of the two best performing patterns over the course of a day. The end-device is located inside a building office without direct Line-of-Sight (LoS) to a Gateway.

5.5 Load Balancing Performance in Practice

We evaluate the effectiveness of NulLoRa (v2) on a 10 sq. km. testbed in a major U.S. city across our university campus network, and inside large office buildings over the course of 28 days in 12 different locations. In order to better characterize the performance of the NulLoRa Hardware over time and in a dynamic environment, we sampled the performance of all of its patterns every 5 minutes. Figure 17 shows the changes in RSSI of our two best performing patterns (from the four available) of one of such locations. We can see that during the day there are multiple changes in the best-performing pattern. We have observed that these changes in RSSI are highly related to the dynamics of the environment, such as large reflective objects near the source or sink of the signal. Furthermore, the variability in signal strength is slowed during the weekends. We observed up to 8.5 dB difference between the two patterns during this day. The dynamism of the environment can define the frequency we need to re-tune the best pattern and the search scope. For this particular location, Patterns 2 and 3 were always the two best performing patterns for the duration of the experiment (28 days). Most of the other 12 locations demonstrated the same proprieties, except for one location where a third pattern was also optimal for a small portion (3.5%) of the total time.

6 LIMITATIONS AND FUTURE WORK

While spatial reuse techniques on LP-WAN clients seem promising, there are several open challenges that need to be addressed in practice. In terms of hardware, our nulling front-end is a proof-of-concept that the RF transmission pattern can be made to improve overall network capacity, but it is not yet suitable in terms of cost and power for most deployments. We imagine that space-efficient nulling antenna arrays can be further miniaturized by using PCB antenna structures such as the Inverted F or ceramic chip antennas. In this work, we also saw that both nulling and antenna selection for diversity provide significant gains, but they cannot both be used simultaneously with a two-antenna configuration. Ideally, a system should decide when nulling is reliable enough to use over diversity based on the current traffic load (and distribution of clients) within the network. Generally, we see that nulling provides better overall network capacity, but, in lower load, cases diversity is potentially more energy-efficient due to reducing spreading factor and retransmissions (when the gain is high).

While we do our best to model antenna radiation patterns, this paper only simulates transmission in 2-dimensions. We collected 3D radiation patterns for our antennas, but have not included this in our simulator, since the computational overhead would significantly increase. This is both an interesting area of future work and likely requires some guidance to be given to installers for antenna placement for both clients and gateways. Along the same lines, we only show a teaser of how hotspot offloading is possible, but we do not actually dive into how to make a stable online algorithm. We believe that simple heuristics or more complex techniques from game theory could easily be applied. Finally, in a real system there needs to be a gateway selection and configuration step that adjust to the network over time. This would be a natural extension to LoRaWAN's Adaptive Data Rate (ADR) configuration mechanism with a simple mark and sweep style operation like what we use in our simulations. There are still open questions about how often this process should execute and what thresholds should be used to avoid clients from thrashing between gateways and modes. We also have not yet explored algorithms that consider the redundancy of gateways as an additional parameter beyond just diversity and nulling.

7 CONCLUSIONS

This paper provides evidence that spatial diversity control, even without significant gain over omnidirectional antennas, can significantly increase LP-WAN network performance. This not only improves the number of clients, and volume of traffic, but also has implications on increasing client life. The key insight in this work is that even at sub-GHz frequencies, it is possible to create simple and low-cost spatial diversity management hardware, which, when considered at scale, provides significant benefit to the overall system.

We first perform a design exploration in a simulation that looks at the impact of the number of sectors and sector beam width on overall network capacity. Using this simulation as a guideline, we design a five-sectored antenna with miniature patch antennas and an RF switch. Through testing in an anechoic chamber, we capture the hardware's real beam pattern that we then feedback into the simulator for comparison. We see that a five-sector design is between 13% and 20% better than an omnidirectional antenna in terms of overall capacity. Since patch antenna systems are still expensive and bulky, we prototype a phase-switched nulling front-end with two offset Isotropic antennas as a simpler alternative design. Based on testing in the anechoic chamber, we see that the nulling antenna does not as closely match our idealized simulated radiation pattern, but still performs quite well in network-wide simulations. This performance stems from the fact that, even though nulling patterns can be less regular, if the client can detect the strongest main direction with deep nulls, the interference is reduced on neighboring gateways. We see that the less expensive nulling design is between 16% and 28% better compared to our baseline omnidirectional design. Finally, we show that, with network coordination, it is possible to perform load shedding from gateways to offload hotspots to dramatically increase capacity around bottlenecks (95% and beyond in ideal situations).

Given the trajectory of LP-WAN adoption in urban environments and the simplicity of adding nulling and diversity capabilities to client endpoints, we do not doubt that spatial diversity systems will play a major role in future LP-WAN deployments.

REFERENCES

[1] 2021. *[Online] MathWorks, Inc. Antenna Toolbox.* www.mathworks.com/products/antenna.html

[2] 2021. *[Online] Pulse Larsen Ceramic Patch Antenna W3215.* www.pulseelectronics.com

[3] LoRa Alliance. 2015. LoRaWAN What is it? - A Technical Overview of LoRa and LoRaWAN LoRa Alliance. [Online; accessed Dec-2019].

[4] Artur Balanuta, Nuno Pereira, Swarun Kumar, and Anthony Rowe. 2020. A cloud-optimized link layer for low-power wide-area networks. In *Proceedings of the 18th International Conference on Mobile Systems, Applications, and Services.* ACM, New York, NY, USA, 247–259. https://doi.org/10.1145/3386901.3388915

[5] Luca Catarinucci, Sergio Guglielmi, Riccardo Colella, and Luciano Tarricone. 2014. Compact Switched-Beam Antennas Enabling Novel Power-Efficient Wireless Sensor Networks. *IEEE Sensors Journal* 14, 9 (2014), 3252–3259. https://doi.org/10.1109/JSEN.2014.2326971

[6] Luca Catarinucci, Sergio Guglielmi, Luigi Patrono, and Luciano Tarricone. 2013. SWITCHED-BEAM ANTENNA FOR WIRELESS SENSOR NETWORK NODES. *Progress in Electromagnetics Research C* 39 (2013), 193–207.

[7] Marco Centenaro, Lorenzo Vangelista, Andrea Zanella, and Michele Zorzi. 2016. Long-range Communications in Unlicensed Bands: The Rising Stars in the IoT and Smart City Scenarios. *IEEE Wireless Communications* 23, 5 (2016), 60–67.

[8] Hoang Nam Dao, Monai Krairiksh, and Dinh Thanh Le. 2016. A design of switched-beam Yagi-Uda antenna for wireless sensor networks. In *2016 International Conference on Advanced Technologies for Communications (ATC).* 393–396. https://doi.org/10.1109/ATC.2016.7764813

[9] Adwait Dongare, Revathy Narayanan, Akshay Gadre, Anh Luong, Artur Balanuta, Swarun Kumar, Bob Iannucci, and Anthony Rowe. 2018. Charm: Exploiting Geographical Diversity through Coherent Combining in Low-Power Wide-Area Networks. In *Proceedings - 17th ACM/IEEE International Conference on Information Processing in Sensor Networks, IPSN 2018.* IEEE, 60–71. https://doi.org/10.1109/IPSN.2018.00013

[10] Andrzej Duda and Martin Heusse. 2019. Spatial Issues in Modeling LoRaWAN Capacity. In *Proceedings of the 22nd International ACM Conference on Modeling, Analysis and Simulation of Wireless and Mobile Systems* (Miami Beach, FL, USA) *(MSWIM '19).* Association for Computing Machinery, New York, NY, USA, 191–198. https://doi.org/10.1145/3345768.3355932

[11] Hassan Fawaz, Kinda Khawam, Samer Lahoud, Steven Martin, and Melhem El Helou. 2021. Cooperation for Spreading Factor Assignment in a Multioperator LoRaWAN Deployment. *IEEE Internet of Things Journal* 8, 7 (2021), 5544–5557. https://doi.org/10.1109/JIOT.2020.3031681

[12] Jing Feng, Che-Wei Chang, Serkan Sayilir, Yung-Hsiang Lu, Byunghoo Jung, Dimitrios Peroulis, and Y. Charlie Hu. 2010. Energy-Efficient Transmission for Beamforming in Wireless Sensor Networks. In *2010 7th Annual IEEE Communications Society Conference on Sensor, Mesh and Ad Hoc Communications and Networks (SECON).* 1–9. https://doi.org/10.1109/SECON.2010.5508256

[13] Weifeng Gao, Wan Du, Zhiwei Zhao, Geyong Min, and Mukesh Singhal. 2019. Towards Energy-Fairness in LoRa Networks. In *2019 IEEE 39th International*

Conference on Distributed Computing Systems (ICDCS). IEEE, 788–798.

[14] Robin George and Thomas Anita Jones Mary. 2020. Review on directional antenna for wireless sensor network applications. *IET Communications* 14, 5 (2020), 715–722. https://doi.org/10.1049/iet-com.2019.0859

[15] Jetmir Haxhibeqiri, Eli De Poorter, Ingrid Moerman, and Jeroen Hoebeke. 2018. A survey of LoRaWAN for IoT: From technology to application. *Sensors* 18, 11 (2018), 3995.

[16] Yong Huang, Weibo Gong, and D. Gupta. 2006. MCMSDA: a multi-channel multi-sector directional antenna wireless LAN. In *2006 International Symposium on a World of Wireless, Mobile and Multimedia Networks(WoWMoM'06)*. 7 pp.–58. https://doi.org/10.1109/WOWMOM.2006.58

[17] Haejung Kim, Sangkyung Kim, Changhwa Kim, and Chanjung Park. 2008. Directional Antennas Based MAC Mechanism for Spatial Reuse. In *2008 Second International Conference on Future Generation Communication and Networking*, Vol. 2. 334–337. https://doi.org/10.1109/FGCN.2008.175

[18] Chenning Li, Hanqing Guo, Shuai Tong, Xiao Zeng, Zhichao Cao, Mi Zhang, Qiben Yan, Li Xiao, Jiliang Wang, and Yunhao Liu. 2021. NELoRa: Towards Ultra-Low SNR LoRa Communication with Neural-Enhanced Demodulation. In *Proceedings of the 19th ACM Conference on Embedded Networked Sensor Systems* (Coimbra, Portugal) *(SenSys '21)*. Association for Computing Machinery, New York, NY, USA, 56–68. https://doi.org/10.1145/3485730.3485928

[19] Li Liu, Yuguang Yao, Zhichao Cao, and Mi Zhang. 2021. DeepLoRa: Learning Accurate Path Loss Model for Long Distance Links in LPWAN. In *IEEE INFOCOM 2021 - IEEE Conference on Computer Communications*. 1–10. https://doi.org/10.1109/INFOCOM42981.2021.9488784

[20] Lu Lu, Geoffrey Ye Li, A Lee Swindlehurst, Alexei Ashikhmin, and Rui Zhang. 2014. An overview of massive MIMO: Benefits and challenges. *IEEE journal of selected topics in signal processing* 8, 5 (2014), 742–758.

[21] Davide Magrin, Martina Capuzzo, and Andrea Zanella. 2020. A Thorough Study of LoRaWAN Performance Under Different Parameter Settings. *IEEE Internet of Things Journal* 7, 1 (2020), 116–127. https://doi.org/10.1109/JIOT.2019.2946487

[22] Jaco M. Marais, Adnan M. Abu-Mahfouz, and Gerhard P. Hancke. 2019. A Review of LoRaWAN Simulators: Design Requirements and Limitations. In *2019 International Multidisciplinary Information Technology and Engineering Conference (IMITEC)*. 1–6. https://doi.org/10.1109/IMITEC45504.2019.9015882

[23] Martin Nilsson. 2009. Directional antennas for wireless sensor networks.

[24] Gopika Premsankar, Bissan Ghaddar, Mariusz Slabicki, and Mario Di Francesco. 2020. Optimal Configuration of LoRa Networks in Smart Cities. *IEEE Transactions on Industrial Informatics* 16, 12 (2020), 7243–7254. https://doi.org/10.1109/TII.2020.2967123

[25] Z. Qin and J. A. McCann. 2017. Resource Efficiency in Low-Power Wide-Area Networks for IoT Applications. In *GLOBECOM 2017 - 2017 IEEE Global Communications Conference*. 1–7. https://doi.org/10.1109/GLOCOM.2017.8254800

[26] M. Rahman and A. Saifullah. 2018. Integrating Low-Power Wide-Area Networks in White Spaces. In *2018 IEEE/ACM Third International Conference on Internet-of-Things Design and Implementation (IoTDI)*. 255–260. https://doi.org/10.1109/IoTDI.2018.00033

[27] Usman Raza, Parag Kulkarni, and Mahesh Sooriyabandara. 2017. Low Power Wide Area Networks: An Overview. *IEEE Communications Surveys & Tutorials* 19, 2 (2017), 855–873. https://doi.org/10.1109/COMST.2017.2652320 arXiv:1606.07360

[28] George F. Riley and Thomas R. Henderson. 2010. The ns-3 Network Simulator. In *Modeling and Tools for Network Simulation*, Klaus Wehrle, Mesut Günes, and James Gross (Eds.). Springer, 15–34.

[29] Benigno Rodríguez, Javier Schandy, Juan Pablo González, Leonardo Steinfeld, and Fernando Silveira. 2017. Fabrication and characterization of a directional SPIDA antenna for wireless sensor networks. *2017 IEEE URUCON* (2017), 1–4.

[30] Wei Sun. 2020. Towards High Throughput Wireless Network with Directional Antenna. *CoRR* abs/2003.02235 (2020). arXiv:2003.02235 https://arxiv.org/abs/2003.02235

[31] Hideaki Takagi, Leonard Kleinrock, and Fellow Ieee. 1985. Throughput analysis for persistent CSMA systems. *IEEE Transactions on Communications* 33 (1985), 627–638.

[32] Aleksandra Tiurlikova, Nikita Stepanov, and Konstantin Mikhaylov. 2018. Method of Assigning Spreading Factor to Improve the Scalability of the LoRaWan Wide Area Network. In *2018 10th International Congress on Ultra Modern Telecommunications and Control Systems and Workshops (ICUMT)*. 1–4. https://doi.org/10.1109/ICUMT.2018.8631273

[33] Thiemo Voigt, Martin Bor, Utz Roedig, and Juan Alonso. 2017. Mitigating Inter-Network Interference in LoRa Networks *(EWSN '17)*. Junction Publishing, USA, 323–328.

[34] Juan Carlos Zuniga and Benoit Ponsard. 2016. Sigfox System Description. *LPWAN-IETF97, Nov. 14th* 25 (2016).