

XCOPY: Boosting Weak Links for Reliable LoRa Communication

Xianjin Xia¹, Qianwu Chen¹, Ningning Hou¹, Yuanqing Zheng¹, Mo Li²

¹ The Hong Kong Polytechnic University, Hong Kong, China

² Nanyang Technological University, Singapore

{xjxia,qwchen,ning2hou,yqzheng}@polyu.edu.hk,limo@ntu.edu.sg

ABSTRACT

LoRaWAN suffers dramatic performance degradation over a long communication range due to signal attenuation and blockages. To ensure reliable data transfer, LoRaWAN adopts retransmission mechanism where an unacknowledged packet is retransmitted multiple times in the hope of successfully delivering the packet at least once over harsh wireless channels. This retransmission mechanism is ill-suited for LoRa: 1) unsuccessful retransmissions lead to high power consumption for battery-powered LoRa nodes, and 2) a retransmission at another time does not necessarily improve the signal strength over harsh wireless channels.

This paper presents the design and implementation of XCOPY, which effectively improves the signal strength by coherently combining retransmitted packets received over weak links that would otherwise be thrown away. XCOPY develops and puts together novel algorithms to 1) accurately identify the signal copies of the same packet over multiple retransmissions in the presence of interfering packets, and 2) precisely align the signal copies (in time, frequency, and phase) to ensure constructive combining, which turns out to be very challenging in ultra-low SNRs, but made possible by XCOPY. Evaluations show that XCOPY can deliver significant SNR gains and yield higher packet reception ratio and longer lifetime of LoRa nodes.

CCS CONCEPTS

• Computer systems organization → Embedded systems; • Networks → Network reliability.

ACM Reference Format:

Xianjin Xia¹, Qianwu Chen¹, Ningning Hou¹, Yuanqing Zheng¹, Mo Li². 2023. XCOPY: Boosting Weak Links for Reliable LoRa Communication. In *The 29th Annual International Conference on Mobile Computing and Networking (ACM MobiCom '23)*, October 2–6, 2023, Madrid, Spain. ACM, New York, NY, USA, 15 pages. <https://doi.org/10.1145/3570361.3592516>

1 INTRODUCTION

Recent years have witnessed a rapid development of Low Power Wide Area Networks (LPWANs) which emerge as a new platform to connect Internet-of-Things (IoT) [3, 5, 26]. LPWANs complement traditional IoT technologies (e.g., Wi-Fi, ZigBee, 5G, etc.) with novel physical layer (PHY) designs capable of offering low-power and long-range communication. For example, LoRa works in unlicensed bands and can be powered by batteries for years. LoRa advertises that a single gateway can cover tens of km², connecting thousands of IoT devices for long-term operations [41, 48, 51]. Given the remarkable performance, LoRa applications are anticipated to expand over 50 % per year [10, 20].

Despite the proliferation of LoRaWANs, many deployed LoRa nodes suffer from weak link connection and degraded performance in practice [7, 8, 29, 45]. For instance, in urban scenarios, the coverage of a LoRa gateway can decrease dramatically due to signal blockage. The IoT nodes deployed in the shadowing regions of gateways inevitably suffer weakened signal strength which leads to high packet loss rates and shortened battery life. A dense gateway deployment can potentially increase coverage but incur high hardware costs. Even for a well-planned network with adequate coverage, LoRa nodes may still suffer packet losses due to interference and channel errors. Besides, if LoRa nodes have to be deployed deep inside buildings (e.g., smart-metering use cases), the signals can attenuate severely after passing through concrete walls and may not reliably reach any gateways even using the highest transmit power.

The current LoRaWAN acknowledges a packet if the packet is received with no bit error, otherwise it discards the corrupted packet and waits for retransmission. However, retransmission over weak links can consistently experience poor signal quality, which seldom delivers retransmitted packets with good enough Signal-to-Noise Ratios (SNRs). The latest work (NELoRa [25]) can relax SNR requirements of LoRa packet demodulation using AI enhancement. Recent

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org. *ACM MobiCom '23*, October 2–6, 2023, Madrid, Spain

© 2023 Copyright held by the owner/author(s). Publication rights licensed to ACM.

ACM ISBN 978-1-4503-9990-6/23/10...\$15.00

<https://doi.org/10.1145/3570361.3592516>




Figure 1: (a) XCopy combines signal power from multiple retransmissions; (b) XCopy fills the SNR gaps between packet detection and decoding.

multi-antenna multi-gateway designs (Charm [7], MALoRa [17]) improve signal strength for weak packets with antenna diversities. But the SNR improvement for retransmitted packets can still be insufficient to cross the SNR thresholds of packet reception in harsh environments (e.g., < -20 dB).

By examining corrupted packets, we find that although the signals of every retransmission remain in low SNRs insufficient for packet reception, the weak signals can be sampled by a LoRa radio and provide partial information about the packet. If we can combine the weak signals of multiple retransmissions, rather than discarding them as the current LoRaWAN does, we may increase the signal quality. As more packets are retransmitted and combined, the SNRs can gradually exceed the threshold needed for successful packet reception as illustrated in Figure 1(a).

In this paper, we present *XCopy* – a system that allows LoRa nodes to combine signal power across multiple retransmissions for reliable communications. At a high level, *XCopy* exploits the fact that a LoRa radio requires sufficiently high SNRs to decode a packet, but can detect the packet at much lower SNRs. Although a detected packet cannot be decoded yet, *XCopy* can aggregate signal power from multiple retransmissions by coherently combining their signal copies. Theoretically, we can expect 3 dB SNR gains as the number of transmissions doubles. A sender can retransmit a packet sufficient times until the SNR of combined signals exceeds the threshold of packet decoding. Each transmission can be effectively used and contribute to higher SNRs. Ideally, once a packet can be detected, the packet can be successfully received eventually, as illustrated in Figure 1(b). The SNR gain can be further boosted, if multiple antennas can be used to collect retransmitted packets and effectively combined. By leveraging M antennas to receive N retransmitted packets, *XCopy* can potentially combine $M \times N$ packets and deliver near-linear improvement (i.e., from $M \times 1$ to $M \times N$) over the latest multi-antenna multi-gateway designs.

However, it entails substantial challenges to implement *XCopy* in practice. Unlike multi-antenna systems that synchronously receive multiple copies of the same packet, retransmitted packets arrive at different times and may overlap with the (re)transmissions of other nodes. It is non-trivial to identify retransmitted signals of the same packet, as the node IDs of packets cannot be decoded in ultra-low SNRs. Moreover, due to clock drifts of LoRa nodes and time-varying

channels, signal copies of retransmitted packets may have heterogeneous time, frequency and phase shifts, which can lead to incoherent signal combinations and adversely affect signal SNRs. It is essential but challenging to synchronize time, frequency and phase among retransmitted signals to ensure constructive combining especially in ultra-low SNRs.

XCopy develops a novel technique that is able to aggregate signal power of an entire packet to combat low SNRs. We exploit the fact that retransmitted signals of the same packet would have high similarities in corresponding chirp windows, despite that they may differ in frequency and phase because of channel variations and clock drifts. *XCopy* detects similarities between two signals by applying a conjugate multiplication. If the two signals are from the same packet and well-aligned in time, an *identical tone-frequency* would be detected from different chirp windows of the packet, which can accumulate signal power of the whole packet into a single frequency bin. Otherwise, different tone-frequencies are produced in different windows and the signal power will be dispersed. *XCopy* searches for the highest power peak in frequency domain to identify retransmitted signals of the same packet and align frame timing among multiple signal copies. While two retransmitted signals are aligned in time, the frequency and the phase of the highest peak indicate the frequency and the phase differences between the two signals. *XCopy* estimates frequency and phase differences from the highest peak (i.e., with signal power of the whole packet) and compensates accordingly to synchronize the frequency and the phase among retransmitted signals.

In addition, *XCopy* needs to handle a unique challenge related to the heterogeneous Sampling Timing Offsets (STOs) of retransmitted LoRa packets termed *STO heterogeneity*. STOs add packet-variant symbol-dependent phase shifts to received LoRa chirps. *XCopy* mitigates the impact of STO heterogeneity by leveraging the fact that the STO of a packet is upper-bounded by the sampling time interval of a receiver. *XCopy* increases the signal sampling rate and refines frame timing alignment with over-sampled signals to reduce the span of STOs among retransmitted packets. According to our results, the impact of STO heterogeneity can be effectively mitigated when LoRa packets are sampled at a sampling rate higher than $3\times$ of the signal bandwidth.

We implement *XCopy* on software define radio platforms (i.e., USRPs and RTL-SDR dongles) and build a testbed consisting of one gateway and 40 commodity LoRa nodes deployed in an urban environment. We collect data from more than 200 links to evaluate *XCopy* in diverse channel conditions. Results show that *XCopy* can yield >10 dB SNR improvement and $>40\%$ coverage extension for LoRaWAN in urban settings. *XCopy* can be integrated with LoRaWAN stacks and work on top of the latest multi-antenna multi-gateway designs to deliver extra SNR gains at reasonable computation and storage overheads.

In summary, this paper makes the following contributions:

- We introduce PHY retransmission as a new paradigm for SNR enhancement of weak LoRa connections by constructively combining weak signals of retransmitted packets that would otherwise be wasted.
- We devise strategies for LoRa packet detection, signal calibration and synchronization under ultra-low SNRs.
- XCOPY complements the state-of-the-art and can improve reliability with >10 dB SNR gains when existing techniques can no longer reliably connect.

2 A PRIMER ON LORA

LoRa PHY. LoRa is a physical layer (PHY) technique that adopts Chirp Spread Spectrum (CSS) for modulation. A chirp signal has time-varying frequencies, where the frequencies increase (*up-chirp*) or decrease (*down-chirp*) at fixed rates to sweep a spectrum. CSS uses frequency-modulated upchirps to represent symbols. A modulated chirp signal can be represented as follows.


$$S(t, f_{sym}) = e^{j2\pi(\frac{k}{2}t - \frac{BW}{2}t)} \cdot e^{j2\pi f_{sym}t} = C(t) \cdot e^{j2\pi f_{sym}t}. \quad (1)$$

Here, $C(t) = e^{j2\pi(\frac{k}{2}t - \frac{BW}{2}t)}$ is a *base chirp* whose frequency increases from $-\frac{BW}{2}$ to $\frac{BW}{2}$, k denotes the frequency increase rate and f_{sym} is the frequency shift of a modulated symbol. To demodulate a chirp signal, we need to first *dechirp* the signal by multiplying with the conjugate of a base chirp (*i.e.*, $C^{-1}(t)$), followed by a Fast Fourier Transform (FFT) to extract the encoding frequency of a symbol (*i.e.*, f_{sym}).

The LoRa PHY uses two parameters, *i.e.*, spreading factor (SF) and bandwidth (BW), to adapt data rate and communication reliability. Generally, a large SF and narrow bandwidth modulate symbols with long chirps, which can improve resilience to poor channel conditions at the cost of long transmission time and high energy consumption.

Packet structure. A LoRa packet is composed of a preamble, an optional header, a payload and a two-byte CRC calculated from the payload data. The preamble consists of a varied number of base chirps, followed by a two-chirp synchronization word and 2.25 down chirps as a Start Frame Delimiter (SFD). A LoRa radio detects a packet by detecting the presence of LoRa preamble and uses the preamble for frame synchronization (*e.g.*, frame timing, frequency and phase calibration) [4, 44]. Upon detecting a preamble, a receiver detects SFD to identify the start of payload and next demodulate and decode the payload. The CRC will be checked by a receiver to validate the integrity of received data.

Retransmissions in LoRaWAN. The LoRaWAN MAC supports confirmed-data as well as unconfirmed-data messages [30]. Unconfirmed messages require no acknowledgment, which is designed for good channel conditions and loss tolerant scenarios. If communication reliability is of concern, confirmed message type shall be adopted. A receiver acknowledges a confirmed-data message if the received data pass CRC check. Otherwise, the message needs to be retransmitted. LoRaWAN recommends a maximum of eight retransmissions.



(a) Packet Reception Ratio (PRR) (b) Packet transmission times
Figure 2: Performance of LoRaWAN retransmission in various SNRs (SF12, BW 125 kHz): MAC layer retransmission cannot benefit weak links (*e.g.*, SNR < -20 dB).




Figure 3: SNR thresholds of packet decoding and detection with different lengths of detection windows (in duration of multiple chirps).

3 WHY PHY RETRANSMISSION?

The LoRaWAN MAC acknowledges confirmed-data messages at a per-packet basis. Any received packet failing to pass CRC checks will be discarded and requires retransmission. This MAC-layer mechanism requires that all payload data of a packet must be correctly received from a single transmission. It can be effective when the physical channel keeps in relatively good SNRs or when a channel varies between good and poor SNRs over time (*e.g.*, mobile scenarios) where transmission in good SNRs is likely to deliver packets successfully. However, for some weak links with severe signal blockages (*e.g.*, inside buildings and urban settings), the physical channel can continuously experience extremely-low SNRs (*e.g.*, < -20 dB). Retransmission at the lowest data-rates (*i.e.*, SF12) with the maximum number of attempts would still result in zero packet reception as shown in Figure 2. MAC-layer retransmissions in such cases seldom improve packet reception performance, which simply waste radio resources, *e.g.*, battery power, duty-cycle quota, *etc.*

An interesting observation is that though packet transmission over weak links may not yield a successful packet reception, the packet signals can still be sampled and detected by the receiver radio. We empirically measure the minimum SNRs required by a LoRa radio (*i.e.*, Semtech SX1276) for packet detection and decoding, respectively. As displayed in Figure 3, LoRa radio is capable of detecting packets at lower SNRs than the SNR thresholds of packet decoding. In particular, if we exploit the unique signal structure of LoRa preamble to detect with longer windows, it may detect packets at even lower SNRs as shown in Figure 3.

We propose to aggregate multiple retransmissions to improve LoRa communication over weak links. The intuition is that, although the signal strength of each transmission can

be weak and insufficient for packet decoding, we may add up retransmitted signal copies to increase SNRs. As such, every retransmission of a packet can contribute SNR gains to the successful decoding of the packet. As more transmissions are coherently combined, it could potentially yield an SNR gain of >10 dB improvement! Moreover, as the sensitivity of packet detection increases with the window length of preamble detection, we may configure a LoRa packet with a long preamble and enlarge the detection window to push the SNR threshold of packet decoding downwards to the physical limit and further increase the SNR gain. It is promising to break the SNR barrier of LoRa communications and boost the data transmission reliability over weak links.

4 XCOPY DESIGN

4.1 Overview

XCOPY is a PHY retransmission strategy for LoRa that is capable of aggregating signal power of multiple transmissions to achieve reliable communication over weak links. A sender retransmits the PHY signals of a LoRa packet and configures the packet with a long preamble to assist packet detection. The length of preamble and the number of retransmissions are adapted to strike a balance between reliability and communication overhead. At the receiver side, XCOPY coherently combines the received signal copies to enhance SNRs and decode the packet with enhanced signals.

Without loss of generality, we model the received signal of a symbol $S(t, f_{sym})$ as below.

$$y(t) = \alpha e^{\varphi} \cdot e^{-j2\pi f_{cfo} t} \cdot S(t + \tau_{sto}, f_{sym}) + n(t), \quad (2)$$

where α and φ represent the amplitude attenuation and phase change of signals caused by communication channel, f_{cfo} is the central frequency offset (*i.e.*, CFO) between a sender and a receiver, τ_{sto} denotes the time gap between signal arrival and signal sampling of a receiver (termed *Sampling Timing Offset (STO)* [44]), and $n(t)$ denotes noises. Suppose a packet is transmitted through a weak link for n times. XCOPY combines n signal copies of the symbol to increase SNRs, which can be represented in the following.

$$Y(t) = \sum_{i=1}^n \omega_i \cdot y_i(t), \quad (3)$$

where $y_i(t)$ stands for the received signal of the i^{th} transmission and ω_i is a weight adjusted according to channel condition and signal quality of $y_i(t)$.

Note that different transmissions generally differ in CFO, STO and channel conditions, due to frequency drifts of LoRa radios and channel variations. Retransmitted signals may have frequency and phase shifts which can lead to *destructive incoherent signal combining* and harm the SNRs of $Y(t)$. XCOPY needs to estimate the frequency drift and phase difference among the raw signals of n transmissions and compensate for the frequency and phase misalignment. Ideally,

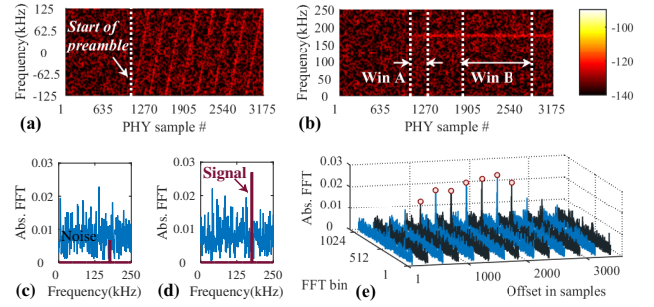


Figure 4: Detect presence of packets: (a,b) Spectra of a LoRa preamble and dechirped results (SF8, SNR: -15 dB); (c,d) FFT of signals in windows A and B; (e) Slide a 4-chirp long detection window to detect preamble.

we expect that compensated signals shall have the same frequency and phase so that the signal power of n transmissions can be combined constructively for the highest SNRs.

XCOPY addresses unique challenges in signal combining across retransmissions. Firstly, as retransmissions take place at different times, the retransmissions of a packet may overlap with the (re)transmissions of other packets. XCOPY needs to correctly identify the retransmitted signals of the same packet and strictly synchronize frame timing, frequency and phase among multiple signal copies. However, it is challenging to detect and synchronize retransmitted packets in the presence of interference of other transmissions, especially in ultra-low SNRs (addressed in §4.2). Secondly, as retransmitted LoRa packets may experience time-varying channels with heterogeneous frequency and phase drifts, how to mitigate their adverse effects on signal combination is critical (handled in §4.3). We investigate the problems in detail and present novel solutions in the following.

4.2 Coherent Combining across Retransmissions

(1) Detecting presence of packets. XCOPY takes advantage of long LoRa preambles to detect the presence of weak packets. The basic idea is that as LoRa preamble consists of several identical base-chirps, a preamble after being dechirped would produce the same tone-frequency in consecutive chirp windows as shown in Figure 4(b). By performing FFT at different windows, we can detect repetitive power peaks pinpointing the same frequency, which indicates the presence of a packet. If the signal power of a single chirp window cannot be detected, we may enlarge the detection window (DW) to aggregate more samples of multiple chirps, which can accumulate more signal power and generate a higher FFT peak (Figure 4(c) and (d)). By sliding the enlarged detection window over received signals, we can detect the desired signal patterns of LoRa preamble even in ultra-low SNRs as displayed in Figure 4(e). We note that the length of LoRa preamble is configurable. The maximum number of identical base-chirps in a preamble supported by a commodity LoRa




Figure 5: Coarse estimation of frame timing.




Figure 6: (a,b) Deviations of detected frame timing with different detection windows (DW); (c) Histogram of timing deviations (DW=4x).

radio (e.g., Semtech SX1276 [37]) can be up to 65,535 chirps. By default, XCOPY configures the length of preamble and detection window to 8 chirps and 4 chirps respectively, which balances the detection sensitivity and the communication and computation overhead.

(2) Coarse estimation of frame timing. Once a packet is detected, XCOPY needs to detect the frame timing to ensure correct chirp-level combination of retransmitted signals. Intuitively, we can use the frequency detected from preamble to infer the frame timing of a packet, according to the relationship between chirp frequency and time. As illustrated in Figure 5(a), suppose a detection window deviates Δn samples from the chirp edge of preamble and the sampling rate is F_s . The signal in the detection window is $C(t + \Delta n \cdot \frac{1}{F_s})$ which, after dechirping, gives a frequency $f = k \cdot \Delta n / F_s$, where k is the frequency increase rate of chirp. We can obtain f from the FFT results of a detection window, and infer the offset between detection window and chirp edge as $\Delta n = f \cdot F_s / k$, from which the timing of chirp edge can be estimated.

We use the above method to detect chirp edges of the preamble shown in Figure 4. We see that some edges are detected before the start of preamble as displayed in Figure 5(b). This is because as we slide a detection window (e.g., four-chirp long) across received signals, early power peaks can be produced when a preamble is partially included in the detection window (e.g., <4 chirps). While the detection window totally moves in preamble, the peaks can get higher. Ideally, the first peak with the highest power indicates the start of preamble (see Figure 5(b)). We may intuitively detect the starting point using a power threshold. However, our empirical study shows that threshold based detection can mistakenly detect early/late peaks as the first base-chirp in the preamble, leading to chirp-level misalignment (i.e., chirp errors). Besides, we observe that frequency detection of preamble chirps in ultra-low SNRs is vulnerable to noise distortions, which can cause sample-level deviations to the

estimated chirp edges (i.e., sample errors). It may even produce heterogeneous frame timing errors for different retransmitted packets. Simply enlarging detection windows (e.g., DW=8x or 12x) does not reduce such detection errors as shown in Figure 6(a). Despite that, the results give coarse estimations of the frame timing. As plotted in Figure 6(b,c), most estimations are within ± 2 chirps of the groundtruth. We can launch refined searches in a confined range to find the correct frame timing.

(3) Synchronizing signal timing, frequency and phase.

While it can be hard to improve frame timing detection with a single packet, can we leverage the retransmission opportunity to mitigate timing errors? A key insight is that as XCOPY retransmits a packet with identical PHY signals, the received signal copies of multiple retransmissions should have strong correlations with each other. Ideally, a high correlation peak can be detected when the frame timing of two signal copies are aligned. We try to use a conventional correlation detection method to align frame timing for retransmitted signals. However, we can hardly detect correlation peaks in time domain when signal SNRs are ultra-low. Even for signals with good SNRs, the detected correlation peaks still deviate from the real alignment points by a number of samples. This is because retransmitted packets may experience different channels, CFOs and STOs [4, 44] in practice, due to channel variations and clock drifts, which can add heterogeneous frequency and phase changes to retransmitted signal copies. The frequency misalignment of retransmitted signals would alter the detected locations of correlation peaks (i.e., timing errors), owing to the inherent relationships between frequency and time of LoRa chirps.

XCOPY presents a new method that can best utilize the signal structure of retransmitted packets to not only align frame timing but also synchronize frequency and phase across retransmissions in ultra-low SNRs. We leverage the fact that if two signal copies of the same packet are aligned in time, high similarities between signals of corresponding windows can be detected from any parts of the packet (e.g., preamble, SFD and payload). Without loss of generality, given two signal copies $y_1(t)$ and $y_2(t)$, XCOPY uses a conjugate multiplication to detect similarities between the two signals as follows:

$$y_1(t) \cdot y_2^*(t) = \|h_1 \cdot h_2^*\| \cdot e^{-j2\pi\Delta f t} \cdot e^{j\Delta\phi}, \quad (4)$$

where $(\cdot)^*$ denotes a complex conjugate and background noises are omitted. h_1 and h_2 represent the channel impacts of two transmissions, while $\|h_1 \cdot h_2^*\|$ indicates the magnitude of the multiplied signals. Δf and $\Delta\phi$ represent the frequency and phase differences between $y_1(t)$ and $y_2(t)$ because of CFOs, STOs, and frequency and phase drifts.

Note that if two signal copies of the same packet are aligned, the conjugate multiplication in Eq.(4) can totally remove the common chirps and produce a tone-frequency as illustrated in Figure 7(a). We can next take an FFT to aggregate signal power from not only preamble, but also other parts of a packet (e.g., SFD and payload) into a single

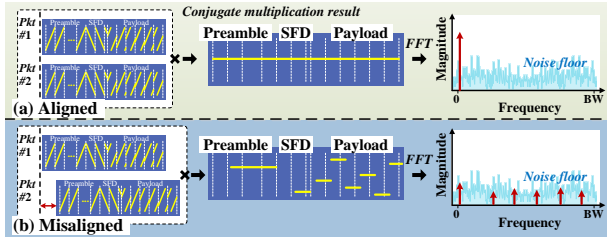


Figure 7: (a) Signal power of the whole packet aggregates into a high peak when retransmitted packets are timing-aligned; (b) Signal power dispersed into multiple frequency bins when signal timing is misaligned.




Figure 8: Alignment of frame timing: (a) Similarity detection with sliding windows; (b) The highest peak indicates timing alignment.

frequency bin, which produces the highest power peak. In contrast, if two signals are misaligned by a few samples, the signal power of misaligned samples would leak out from the main tone-frequency, resulting in decreased power peaks. As the timing misalignment enlarges, diverse tone frequencies will appear in different chirp windows depending on the frequency difference between the misaligned chirps. In this case, the signal power of a packet is dispersed into multiple frequency bins and cannot be accumulated to surpass the noise floor as shown in Figure 7(b). In summary, the highest frequency peak indicates an alignment of frame timing.

In practice, XCopy uses a sliding window to calculate Eq.(4) with different time offsets of two signal copies, which can produce a result as displayed in Figure 8. XCopy then searches for the highest frequency peak to precisely align frame timing among signal copies of retransmitted packets. Compared to a conventional correlation based detection, the conjugate multiplication method is able to produce the highest power peak which indicates the correct timing alignment of two signals even when there are frequency shifts between the two signal copies because of CFOs. Besides, Eq.(4) provides additional information (*i.e.*, Δf and $\Delta\phi$), which can be used to estimate frequency and phase differences between the two signals. We can extract Δf and $\Delta\phi$ from the highest peak shown in Figure 7(a) and compensate accordingly to synchronize frequency and phase between the two signals. By compensating for the frequency and phase differences, we can ensure that the combined signals will add constructively and thereby effectively increase SNRs of the weak packet.

(4) Packet grouping and combining. In practice, the transmissions of different LoRa nodes can interleave in time, which may interfere the signal combining of multiple retransmissions for a target packet. To ensure coherent combining,

XCopy needs to identify and combine the retransmitted signals that belong to the same packet (*i.e.*, *packet grouping*). If node IDs of packets were available, we can separate interleaved packets and group the packets according to their node IDs. However, this method requires relatively high SNRs to decode node IDs and may fail in ultra-low SNRs.

XCopy takes advantages of the PHY signal features of LoRa packets to separate interleaved transmissions in low SNRs. A key insight is that the packets of different nodes would generally carry different chirps in the payload (*e.g.*, different node IDs and encoded data). If we apply Eq.(4) to the signals of different packets, diverse tone frequencies would appear in different chirp windows, similar to the situation displayed in Figure 7(b), even when the frame timing of two packets are aligned. As such, it would be less likely to detect high power peaks above the noise floor if two signals are from different packets. Said differently, if a high power peak is detected, we can ensure that two signals are from the same packet with a high confidence. Based on this observation, XCopy uses the detection results of Eq.(4) to classify the interleaved transmissions of different nodes into packet groups, where each group contains retransmitted signal copies of the same packet. Signal combination is performed within each group to enhance SNRs of received packets.

4.3 Mitigating STO Heterogeneity

As a receiver samples wireless signals at discrete time, the time offset between signal arrival and signal sampling termed *Sampling Timing Offset (STO)* would add identical frequency shift but heterogeneous phase shifts to different symbols of the same packet [44]. The frequency shift of STO is coupled with CFO. We can use Eq.(4) to estimate and compensate the overall frequency difference of STO and CFO in between retransmitted packets. However, as the phase shift of STO differs across symbols, we cannot use the same method to uniformly remove the phase difference of STO. XCopy has to address the heterogeneous phase impacts of STOs to ensure good SNR gains for LoRa signal combination.

As STO can differ across retransmissions in general, the STO difference between two packets determines how their signals are combined (in-phase or out-of-phase). We empirically investigate the impacts of STO difference on LoRa signal combination and examine SNR gains on a per-chirp basis as plotted in Figure 9. We see from Figure 9(a) that chirp combination generally has high SNR gains when the STO difference of two packets is small. While STO difference enlarges, the SNR gains can vary dramatically across different chirps of a packet. For instance, when STO difference is $0.75\times$ of a sampling interval, combining two copies of symbol #122 can get near 3 dB SNR gain. Whereas for symbols #16 and #213, the SNR gains become negative (see Figure 9(b)). Interestingly, we can always get high SNR gains (close to 3 dB) across all payload chirps as STO difference decreases to $0.125\times$ of a sampling interval as shown in Figure 9(b).




Figure 9: Impacts of STO difference on chirp combination: (a) Statistics of per-symbol SNR gain; (b) Heterogeneous impacts on different symbols.




Figure 10: Illustration of signal combination for LoRa symbols in various STO differences ($\Delta\tau_{sto}$).

Theoretically, an STO in τ_{sto} can alter the phase of a chirp by $\varphi_{sto} = 2\pi f_0 \tau_{sto}$, where f_0 represent the initial frequency of the chirp. Let $\Delta\tau_{sto}$ denote the STO difference between two packets. The phase difference of STOs between the corresponding chirps of two packets can be estimated as below.

$$\Delta\varphi_{sto} = 2\pi f_0 (\tau_{sto1} - \tau_{sto2}) = 2\pi f_0 \Delta\tau_{sto}. \quad (5)$$

We see that the phase difference (*i.e.*, $\Delta\varphi_{sto}$) is determined by both f_0 and $\Delta\tau_{sto}$. If $\Delta\varphi_{sto}$ is small, two signal copies of a symbol can be combined coherently with the optimal SNR gains. Figure 10 illustrates the joint effects of $\Delta\tau_{sto}$ and f_0 on signal combination of different LoRa chirps according to Eq.(5). Generally, large STO difference ($\Delta\tau_{sto}$) produces large phase difference. A chirp with large initial frequency (f_0) can amplify the phase difference of STOs. As f_0 is bounded within $-\frac{BW}{2} \sim \frac{BW}{2}$, we have the following inequalities.

$$-(\Delta\tau_{sto} \cdot BW) \times \pi \leq \Delta\varphi_{sto} \leq (\Delta\tau_{sto} \cdot BW) \times \pi \quad (6)$$

Therefore, if the STO difference between retransmitted packets (*i.e.*, $\Delta\tau_{sto}$) is reduced, the phase difference can be bounded within a small range. For instance, when $\Delta\tau_{sto}$ decreases to $0.125 \times$ of $\frac{1}{BW}$, the maximum phase difference would be less than 22.5° which can ensure coherent combination for all symbols of the packet, as illustrated in Figure 10. It encourages us to reduce STO difference among retransmitted packets to avoid the adverse effects of STO heterogeneity.

Note that STO is caused by the discrete ADC sampling of a receiver. The STO difference between any two packets is physically bounded within one sampling interval. Based on this observation, we can increase the physical sampling rates of receiver to limit the upper-bound of STO difference. Theoretically, if sampling rate is higher than 1.5 times of LoRa bandwidth (*i.e.*, $1.5 \times BW$), the STO difference ($\Delta\tau_{sto}$) is smaller than $\frac{1}{1.5 \times BW}$ and STO-induced phase differences (*i.e.*, $\Delta\varphi_{sto}$) would be less than 120° for any symbols according to




Figure 11: (a) Statistics of STO difference and (b) Per-symbol SNR gains in various sampling rates.

Eq.(6), which is sufficient to ensure non-destructive add-up of signal power for all symbols of the packet. While sampling rate exceeds $3 \times BW$, the maximum phase difference is less than 60° which can result in constructive add-up of symbol power for most cases.

To investigate the effects of sampling rates (F_s) on reducing STO heterogeneity, we receive 1,000 LoRa packets (SF8, BW250 kHz) with different sampling rates. We align frame timing of packets and measure their STO difference. As plotted in Figure 11(a), the measured STO differences distribute in $0 \sim \frac{1}{F_s}$ for all cases. As sampling rate increases, the STO difference decreases proportionately. Decreased STO differences can effectively reduce phase difference between signal copies of symbols, which is helpful to constructive add-up of signal power. As displayed in Figure 11(b), high sampling rates help reduce the variance of SNR gains across symbols. We can get consistently high SNR gains for all symbols as sampling rate exceeds $3 \times$ of LoRa bandwidth.

Note that commodity LoRa radios (*e.g.*, SX1276) sample signals at a rate higher than the bandwidth of LoRa packets [37]. In practice, we can adequately re-sample the raw signals (*e.g.*, using fractional resampling) and apply Eq.(4) on the over-sampled LoRa signals to align frame timing of packets in fine-grained scales, which can bound the STO difference between packets within desired small gaps.

4.4 Putting All Together

Figure 12 illustrates the workflow of XCopy which involves three key steps, *i.e.*, *packet detection*, *signal combination* and *packet decoding*.

An XCopy empowered receiver senses wireless channel with high sampling rates and uses large detection windows to detect the presence of a LoRa preamble. From the power-accumulated frequency peak of a large detection window, XCopy is able to estimate the coarse frame timing and SNRs of a received packet. Signals of non-LoRa transmissions will be filtered out at this step. If the SNR of a LoRa packet is high enough, XCopy will skip signal combination and try to decode the packet directly. Otherwise, XCopy waits for the retransmissions of the packet and proceed to the next stage.

In the second step, XCopy classifies the interleaved transmissions of different packets into diverse packet groups. It ensures that only the signal copies of the same packet will be combined together for SNR enhancement. XCopy aligns

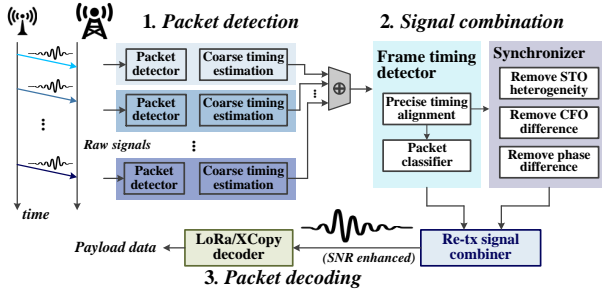


Figure 12: General workflow of XCopy.

frame timing among retransmitted packets with the over-sampled signals, which can remove the negative effects of STO heterogeneity. With the timing-aligned signals, XCopy estimates and compensates for frequency and phase differences among signal copies. Eventually, XCopy can adjust the weights of multiple retransmitted signals (*i.e.*, ω_i) according to their coarsely-estimated SNRs, and coherently combine the PHY signals of multiple retransmissions according to Eq.(3). The SNR-enhanced signals are lastly fed into a standard LoRa decoder for packet decoding (*i.e.*, Step 3), which completes the reception of a weak LoRa packet.

Integrating with LoRa stack. XCopy performs retransmission at a PHY layer. A transmitter stores the signal samples of a packet before receiving an ACK and retransmits identical samples to facilitate signal combining in XCopy. According to the LoRaWAN specification [30], a packet will be retransmitted at different frequency channels. If a packet remains unacknowledged after two retransmissions, LoRaWAN will adapt to lower data rates (*e.g.*, increment SF). As XCopy performs signal combining in the baseband, retransmissions in different frequency channels can be combined together for SNR gains since their signals will be uniformly down-converted to the baseband. XCopy adapts the default retransmission scheme of LoRaWAN by enabling more retransmission attempts with the same SF before changing to a slow data rate. By combining multiple retransmissions, XCopy can gradually converge the effective data rate to what the wireless channels can support. Specifically, to handle ACK losses over a weak link that will not be retransmitted in LoRaWAN, a maximum retransmission attempt for XCopy is set based on historical link SNRs.

5 EVALUATION

5.1 Methodology

Implementation & Devices. We implement XCopy on Software Defined Radio platforms with high-end USRPs (N210) and low-cost RTL-SDR dongles, based on an open-source LoRa project *gr-lora* [11]. We use USRPs as gateways to receive signals from LoRa nodes and forward raw PHY samples to a workstation running XCopy for offline processing. Besides, we connect RTL-SDR dongles to a laptop and move to different locations to collect data traces in diverse SNRs.

We build a testbed composed of 40 commodity LoRa nodes and one gateway. LoRa nodes consist of Dragino LoRa shields that are embedded with Semtech SX1276 radio. We connect LoRa shields to Arduino Uno boards to configure the radio chips. We deploy the testbed in our campus covering a 1.8 km \times 2.1 km area in an urban environment. The gateway is mounted at the roof of a tall building to serve LoRa nodes deployed in the campus.

Experiment setup. We collect data traces from more than 200 links connecting the nodes to the roof-mounted USRPs as well as the mobile RTL-SDR dongles. The links cover wide ranges of channel conditions. We perform extensive experiments to evaluate XCopy over two months, aiming to answer the following three questions: (1) How much performance gain can XCopy bring to LoRa packet decoding (§5.2)? (2) What are the capabilities of XCopy on decoding weak packets (§5.3)? and (3) How does XCopy perform in practical networks (§5.4)? Unless otherwise specified, we configure LoRa packets with default parameters as: central frequency 915 MHz, Spreading Factor (SF) 10, bandwidth (BW) 250 kHz, Coding Rate (CR) 4/8 and payload size 64 chirps. The PHY sampling rate of LoRa gateway is 1 MSps by default.

Benchmark. We compare XCopy with three benchmarks: (1) *LoRaWAN*: the standard protocol with MAC layer packet retransmission which is supported by existing commodity LoRa devices; (2) Multi-gateway *Charm* [7]: a state-of-the-art strategy that combines signals of multiple gateways for weak packet reception; (3) Multi-antenna *MALoRa* [17]: the latest research that leverages the MIMO capability of a multi-antenna gateway to improve weak packet reception.

Metrics. We evaluate the performance of XCopy with three key metrics: (1) *Symbol Error Rate (SER)* of packet decoding; (2) *Packet Reception Ratio (PRR)* of link communication; and (3) *SNR threshold* defined as the minimum SNRs of a packet that can be decoded by XCopy, which characterizes the decoding capability of XCopy.

5.2 Basic Performance

Retransmission gain. We first examine the performance improvement of PHY retransmission on packet decoding. In this experiment, we set up three links with SNRs of -20 dB, -25 dB and -30 dB, respectively. We transmit a packet over the three links with different numbers of retransmissions varying from 0 to 15 to investigate the impacts of number of signal copies on packet decoding performance. The results of zero retransmission are used for baseline comparison.

Figure 13(a) presents the Symbol Error Rates (SERs) of packet decoding in different settings. We see that symbol errors decreasing in general as more retransmissions are used, due to the power benefits of PHY signal combination. For instance, a single retransmission can reduce SERs from 48.2% to 8.8% for the link with -20 dB SNR. Such symbol errors can be tolerated and corrected by error correction codes (*i.e.*, Hamming codes). When a packet is retransmitted more than three times, we can correctly decode almost all symbols as




Figure 13: (a) Performance of XCOPY with different retransmissions; (b) Comparison between XCOPY and standard LoRaWAN; (c) Comparison between XCOPY, Charm and MALoRa with the same number of signal copies; (d) Packet decoding performance (SNR=-30 dB) when XCOPY and MALoRa are jointly used.

shown in Figure 13(a). As expected, we need more retransmissions for links with lower SNRs (e.g., 5 retransmissions for -25 dB and >15 times for -30 dB). The number of retransmissions shall adapt to the channel conditions of a link to ensure good performance on packet decoding.

Next, we evaluate the performance gain of PHY retransmission over existing MAC-layer retransmissions adopted by the LoRaWAN standard. We vary link SNRs from -35 dB to -5 dB by changing the transmit power and locations of LoRa nodes. We transmit a packet repeatedly eight times over each link. We use two packet settings: a fixed SF (i.e., SF10) for all SNRs and Adaptive Data Rate (ADR). When ADR is enabled, we use SF10 for -5 dB and -10 dB, SF11 for -15 dB and -20 dB, SF12 for <-20 dB. XCOPY adds up the signals of eight transmissions for packet decoding. In contrast, LoRaWAN decodes the packet with the signal of a single transmission and selects the lowest SERs from eight transmissions as the decoding results.

Figure 13(b) plots the SERs of LoRaWAN and XCOPY in various link SNRs. We see that with a fixed SF10, LoRaWAN can decode packets when SNR is -15 dB, while XCOPY is able to correctly decode packets even when SNR decreases to -25 dB. Similarly, XCOPY pushes the minimum SNRs of packet decoding from -20 dB to -30 dB when ADR is enabled. It means that the combining of eight signal copies in XCOPY brings near 10 dB SNR improvement compared to the MAC layer packet retransmissions of existing LoRaWAN.

Comparison with the state-of-the-art. In this experiment, we compare XCOPY with two closely related works, Charm [7] and MALoRa [17], which combine signals of multiple gateways/antennas for SNR enhancement. For fairness, we evaluate the packet decoding performance of XCOPY, Charm and MALoRa with the same number of signal copies. We set up eight USRPs as gateways to receive. For Charm, we place the gateways at different locations and configure them with their asynchronous local clocks. We use the method of Charm to synchronize the signals of eight gateways. For MALoRa, we feed all USRPs with a common external clock source which synchronizes signal receptions of eight receivers. For XCOPY, we retransmit a packet seven times and use eight signal copies of the packet from one gateway for signal combination. We also decode the received signals of each gateway without combination and choose the lowest

SERs of eight gateways as the results of standard LoRaWAN for baseline comparison.

Figure 13(c) displays the SERs of four strategies under different SNRs. As expected, LoRaWAN has the highest symbol error rates. Charm performs better than LoRaWAN because it can combine the signals of eight gateways for SNR enhancement. However, Charm suffers inaccurate signal calibration in ultra-low SNRs, and produces higher SERs than XCOPY and MALoRa while SNRs decrease below -15 dB. As MALoRa uses synchronized USRPs to receive multiple copies of a packet, the received signals are synchronized in nature. MALoRa can perfectly combine eight signal copies to enhance SNRs and yield the lowest symbol error rates as shown in Figure 13(c). The performance of XCOPY is quite close to that of MALoRa. It indicates that XCOPY can well synchronize retransmitted signals and produce comparable performance with MALoRa.

We note that as retransmission is complementary with the multi-antenna approach of MALoRa, XCOPY can work together with MALoRa to further improve the performance on weak packet decoding. Figure 13(d) evaluates the performance of using XCOPY and MALoRa jointly to decode packets in -30 dB. We vary the number of retransmissions from 0 to 15 and use multiple synchronized USRPs to receive. When XCOPY and MALoRa are used together, we first apply MALoRa to signals of multiple USRPs for each transmission and then use XCOPY to combine signals of different transmissions. As plotted in Figure 13(d), the symbol error rates reduce continuously as more retransmissions and more receivers are combined. If a single receiver is used, XCOPY needs 15 retransmissions to deliver a 20 % SER. In contrast, we only need 5 retransmissions to a 4-antenna gateway to achieve comparable performance.

Comparison with LoRa ADR. This experiment compares XCOPY against a standard LoRa retransmission scheme with Adaptive Data Rate (i.e., LoRa ADR). We set up 20 links with SNRs varying from -30 dB to -15 dB, i.e., five links for each SNR. For LoRa ADR, we increase SF by one after two retransmissions until a packet is successfully received. The maximum retransmission attempt is eight. For XCOPY, we keep SF unchanged and combine retransmissions coherently for packet decoding. For fair comparisons, we adopt the same initial SF settings for XCOPY and LoRa ADR and examine




Figure 14: Comparisons with LoRa ADR (Adaptive Data Rate): (a) snapshot of retransmission energy when SF=8, SNR=-25 dB, (b) overall energy performance; (c) per-packet latency, and (d) link throughput.

the energy, latency and throughput performance of the two strategies under different SNRs and initial SFs.

Figure 14(a) presents a snapshot of cumulative transmission energy for XCopy and LoRa ADR when SNR is -25 dB. We see that LoRa ADR needs 7 retransmissions. A packet can be successfully delivered when SF is adjusted from 8 to 11. The overall energy consumption is 1,405 mJ. In contrast, XCopy only requires 5 retransmissions with a fixed SF8. Thanks to the short packet air-time of SF8, the total energy of XCopy is only 458 mJ. Note that LoRa ADR does not adjust SF from 8 to 11 via multiple retransmissions for every packet. Instead, it directly uses SF11 to transmit the following packets. However, due to the long packet air-time of SF11, the energy cost can still be as high as 429 mJ, which is comparable to the total energy of XCopy in six (re)transmissions with SF8.

Figures 14(b-d) evaluate the overall energy performance, packet latency and throughput of XCopy and LoRa ADR. The two strategies perform the same when SNRs are high or SFs are large (e.g., -15 dB or $SF \geq 10$) because retransmissions are not required. As link SNRs decrease to -20 dB and -25 dB, multiple retransmissions are used in XCopy to enhance SNRs for every packet. LoRa ADR uses retransmissions to adjust SF to a proper value for the first packet only and uses the adjusted SF to transmit the following packets. Though more retransmissions are used in XCopy than LoRa ADR, the air-time of a packet in LoRa ADR is several times of that in XCopy. As a result, the overall energy and latency performance are generally comparable for the two strategies.

However, as SNRs further decrease to -30 dB, LoRa ADR cannot deliver packets even with the largest SF12. XCopy can still successfully transmit packets at low throughput.

Impact of sampling rate. The physical sampling rate of a receiver (F_s) can affect STOs of received signals as well as the effect of signal combination. This experiment examines the impact of sampling rates on packet decoding. We set up a link with SNR -20 dB and vary the sampling rates of gateway from 250 kSps to 1 MSps. Figure 15 presents the SERs of XCopy in different settings. We see that for the same number of signal copies, the SERs of XCopy decrease as F_s increases. The main reason is that a high sampling rate can reduce STO difference among signal copies, which is beneficial for coherent signal combination of payload chirps. For a fixed sampling rate (e.g., $F_s=250$ kSps), the SERs of XCopy also




Figure 15: Impacts of sampling rate on XCopy. **Figure 16: Overheads on decoding a symbol.**

decrease as more signal copies are combined because more copies can provide more signal power for higher SNR gains.

Overhead analysis. Figure 16 evaluates the computational overheads of LoRaWAN, Charm, MA LoRa and XCopy. We run the four strategies on the same desktop computer and measure the total time of packet decoding. Figure 16 displays the average time spent on decoding a symbol. We see that XCopy has higher time overheads than other strategies because it needs more time to detect and synchronize signals among retransmissions in low SNRs. The average time cost of symbol decoding, including signal processing for packet detection and synchronization, is a few milliseconds. The overhead increases as more signal copies are combined. Considering the low duty cycle and long air time (second level) of LoRa symbols, the decoding latency (i.e., tens of ms) is acceptable in general.

We note that in addition to time overhead, XCopy can incur high memory overheads. Upon detecting a weak packet, XCopy needs to store incoming signals for the longest packet length in multiple retransmission attempts. A large storage volume is required. Fortunately, XCopy runs on gateways that have sufficient resources to store and process signals. Besides, it is feasible to accelerate computations with hardware and upgrade storage volumes for a gateway in practice.

5.3 Capability Study

Packet detection capability. Recall that XCopy needs to detect repetitive patterns (i.e., power peaks in the same frequency bin) of a LoRa preamble to detect a packet. The capability of XCopy on packet detection can be affected by the size of detection window, the number of detecting peaks and the length of preamble as well. We change these factors in the experiments and evaluate their impacts on packet detection performance. The results are presented in Figure 17. We have two key observations: (1) A small detection window may




Figure 17: Packet detection capability: (a,b) Impacts of detection window (DW) and peak number; (c) SNR threshold of preamble detection.




Figure 18: Packet decoding capability: Impacts of (a) Spreading Factor (SF), (b) Bandwidth (BW) and (c) Coding Rate (CR).

miss some weak packets, resulting in *false negative errors* as shown in Figure 17(a). We need to enlarge detection windows to improve detection reliability in low SNRs. (2) A large detection window (e.g., DW=4 in Figure 17(b)) may detect the same noisy peaks in successive windows and mistakenly detect the patterns as LoRa preamble, i.e., *false positive errors*. We can reduce false positive errors by increasing the number of detecting peaks slightly larger than DW, e.g., 5 peaks for DW=4, as shown in Figure 17(b).

Figure 17(c) plots the minimum SNRs of LoRa packets that can be detected by XCopy with different preamble lengths. We empirically set DW and the number of detecting peaks to the maximum values supported by a given preamble length. As expected, a packet with longer preamble can be detected at lower SNRs, because longer preambles empower XCopy to use larger detection windows which can aggregate signal power of more preamble chirps to detect packets in weaker signal strength. As shown in Figure 17(c), the SNR threshold of packet detection (SF8) decreases from -14 dB to -23 dB as preamble length increases from 8 chirps to 32 chirps. For a given preamble length, the SNR thresholds can be lower for packets with larger SFs since a larger SF corresponds to longer symbol duration which is more resilient to low SNRs.

Packet decoding capability. In this experiment, we investigate the decoding capability of XCopy with different numbers of retransmissions. We change configurations of LoRa packets, i.e., Spreading Factor (SF), Bandwidth (BW) and Coding Rate (CR), to study their impacts on decoding capability. Figure 18 presents the minimum SNRs of packets that can be correctly decoded in different settings. In general, XCopy is capable of decoding packets in lower SNRs when more retransmissions are used, as more retransmissions provide more signal copies for higher SNR improvements. For




Figure 19: Sensitivity on packet synchronization (SF8, -15 dB): (a) Sync. detection with raw signals (Pkt Len: 48 chirps); (b) Sync. detection with longer packets (240 chirps) when $F_s=0.25$ MSps and (c) higher sampling rates ($F_s=1$ MSps) when Pkt Len = 48 chirps.

Pkt Len*	SF8	SF12	F_s^*	SF8	SF12
48	-13 dB	-19 dB	0.25 MSps	-13 dB	-19 dB
120	-15 dB	-21 dB	1 MSps	-17 dB	-23 dB
240	-17 dB	-23 dB	2 MSps	-19 dB	-25 dB

*Sampling rate: 0.25 MSps *Pkt Len: 48 chirps, BW: 250 kHz

Table 1: SNR sensitivity of XCopy in various packet settings: long packets, high sampling rates help timing detection in low SNRs.

LoRa packets with SF8, the SNR threshold decreases from -12 dB to -24 dB as the number of retransmissions increases from 0 to 15. The SNR thresholds can be even lower for packets with larger SFs. For instance, the minimum SNRs of decoded packets are -31 dB for SF10 and -35 dB for SF12 as shown in Figure 18(a). By contrast, the decoding capability of XCopy does not change much across different BW settings (see Figure 18(b)). While Coding Rate (CR) affects the capability of a LoRa decoder on bit error recovery, the decoding capability of XCopy increases by 1.9 dB on average as CR changes from 4/5 to 4/8 as shown in Figure 18(c). However, the SNR gains come at a cost of extra encoding overheads, as CR 4/8 would require more symbols to encode redundancy bits than CR 4/5.

Sensitivity study. Note that the packet synchronization and grouping operations in XCopy can be affected by settings of packet length and sampling rates. We set up experiments to examine the impacts of these factors. Figures 19(a-c) show the results of synchronization detection in -15 dB with different sample offsets between the signals of two retransmitted packets. As displayed in Figure 19(a), we detect no peaks from the raw signals of a 48-chirp packet sampled at 0.25 MSps due to the low SNRs. In contrast, clear power peaks are detected as the packet length and sampling rate increase. The peak’s magnitude drops dramatically when the signals of retransmitted packets are misaligned by even a few samples, as shown in Figures 19(b,c), indicating a high sensitivity of synchronization accuracy. In Table 1, we measure the minimum SNRs with which XCopy can correctly group and synchronize retransmitted packets. We see that XCopy can function effectively in ultra-low SNRs (e.g., <-15 dB). The SNR sensitivity generally improves as packet lengths and sampling rates increase.

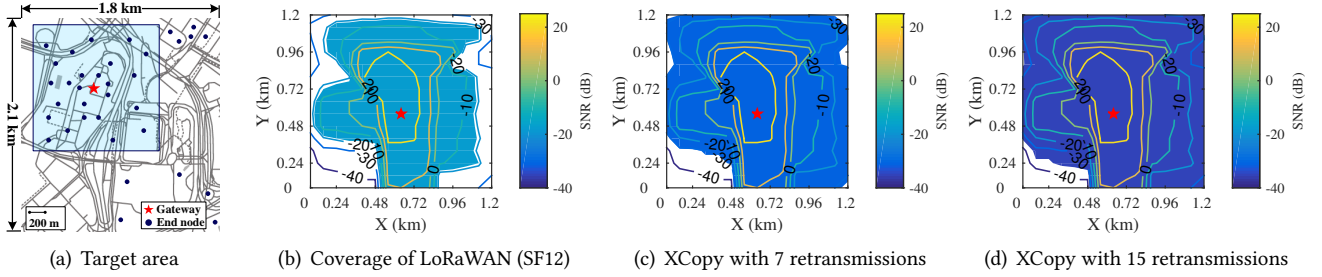


Figure 20: Coverage expanding of XCopy over standard LoRaWAN: Contour lines show the spatial distribution of SNRs, the area with reliable packet communication is marked in color blue.

5.4 Practical Performance

Communication range extension. This experiment examines the performance of XCopy on improving LoRa coverage. To obtain the spatial coverage of LoRaWAN gateway in our testbed, we divide a 1.2 km×1.2 km target area shown in Figure 20(a) into 15×15 grids and measure SNRs of LoRa signals from the gateway at each grid. The target area includes challenging indoor and outdoor environments where wireless signals are attenuated and blocked by concrete walls inside buildings and basements. We plot the measured SNRs with contour lines in Figure 20. The blue areas in Figure 20(b) represent the coverage of LoRaWAN gateway when SF12 is used for communications (packet retransmission is enabled). For XCopy, we configure LoRa nodes to send packets (SF12) with a 32-chirp long preamble and retransmit in each grid for 7 and 15 times, respectively. The gateway runs XCopy to combine all retransmitted signals for packet decoding. We present the coverage of XCopy in Figures 20(c) and (d), where a grid is marked blue if packets sent from the grid can be reliably received (*i.e.*, packet reception ratio $\geq 90\%$). Comparing Figures 20(b) and (c), we see that XCopy has expanded LoRa coverage from areas in SNRs of -18 dB to -31 dB. When 15 retransmissions are used, coverage area further extends to SNRs of -35 dB. Many weak link connections in Figure 20(b) can now communicate reliably. The overall coverage area of LoRaWAN has been expanded by XCopy for 37.2% and 43.4% respectively when 7 and 15 retransmissions are used.

Performance on weak links. To study the performance of XCopy on weak links, we examine three links from our testbed, *i.e.*, link #1 with SNRs in $-37.6\sim-24.1$ dB, link #2: $-26.7\sim-17.1$ dB and link #3: $-18.9\sim-4.3$ dB. The three LoRa nodes wake up every hour to transmit 120-Byte data with a duty cycle of 1%. The bandwidth of LoRa communication is 250 kHz. Spreading Factors are selected according to the SNRs of three links. We collect a 5-day data trace of the three links and replay the trace with XCopy and LoRaWAN respectively. XCopy uses the same packet configurations for all retransmissions; whereas LoRaWAN adapts data rate in retransmissions following the standard [30]. The maximum retransmission attempt is 8. We measure Packet Reception Ratio (PRR) and delays for three links. We also record the active time of LoRa

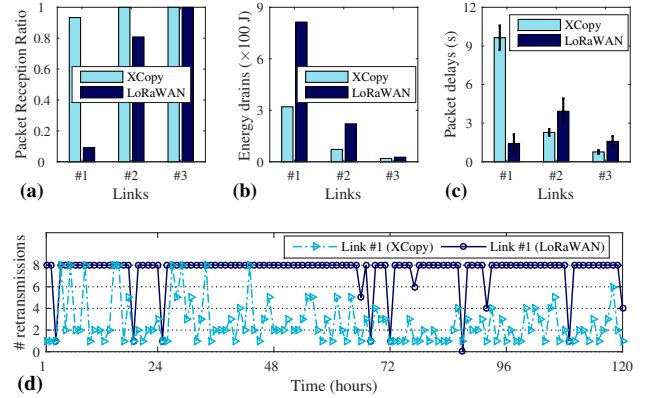


Figure 21: XCopy performance on weak links: (a) Packet Reception Ratio, (b) Energy consumption of end nodes, (c) Packet delays of correctly received packets, (d) A snapshot of retransmission times of Link #1.

radio and calculate energy consumption of end nodes based on the datasheet of Semtech SX1276 [37].

Figures 21(a-c) compare XCopy and LoRaWAN in terms of Packet Reception Ratio (PRR), energy consumption and packet delays. We see that XCopy yields higher PRRs and, in the meanwhile, consumes lower energy than LoRaWAN. For instance, the PRR of LoRaWAN is 80.8% on link #2, while XCopy improves PRR to 100%. LoRaWAN rarely receives packets from link #1 (*e.g.*, $\text{PRR} < 10\%$). In contrast, 93.3% of the packets are correctly received by XCopy on link #1. Comparing the PRRs of three links as shown in Figure 21(a), we find that XCopy is able to produce higher PRR improvements for links in lower SNRs. More importantly, as XCopy effectively uses the signals of every transmission, it can use fewer retransmissions than standard LoRaWAN to achieve even higher PRRs (*e.g.*, link #1 in Figures 21(a) and (d)). As a result, XCopy is capable of reliably delivering packets with lower energy consumption and shorter time delays than LoRaWAN as shown in Figures 21(b) and (c). In Figure 21(c), LoRaWAN exhibits shorter packet delays on link #1 because the delays of packets that cannot be received after 8 retransmissions are not included in the results.

Scalability. In this experiment, we deploy 11 nodes inside a building connecting to a roof-mounted gateway. The link SNRs vary from -27 dB to -22 dB. We configure nodes




Figure 22: XCopy performance with (a) four weak links, and (b) different numbers of interfering nodes.

with a duty-cycle of 10% (*i.e.*, 10× larger than regular LoRaWANs) to mimic the communications of a large-scale network (*e.g.*, >100 nodes). The packet retransmissions of different nodes are overlapped in time, where the overlapping (re)transmissions of different nodes can add interference to the signal combination of a target packet. We examine the performance of XCopy in the presence of interfering traffic.

Figure 22(a) presents the packet reception ratio of four nodes, while other nodes in the network perform similarly (not shown in the figure). XCopy can correctly identify the retransmitted packets of each node and mitigate the impact of interfering packets. Figure 22(b) focuses on the PRRs of one node under interference from a varied number of interfering nodes. We see that the decoding performance of XCopy is mainly affected by the number of retransmission attempts, *e.g.*, 7 retransmissions yield higher PRRs than 3 retransmissions. With fixed retransmission attempts, PRRs do not change much as the number of interference increases from 0 to 10, indicating good scalability of XCopy.

In practice, LoRa nodes shall operate with duty cycles lower than 1%. Many nodes can communicate with different SFs without interfering with each other. We believe XCopy can effortlessly support a relatively large network consisting of hundreds of LoRa nodes by assigning nodes with orthogonal SF and BW parameters.

6 RELATED WORK

Many measurement studies [29, 33, 38, 49] have shown that the performance of LoRa degrades severely in urban settings. Existing works [16, 24, 34, 35, 43, 46, 47] mainly focus on parameter optimization for LoRa (*e.g.*, spreading factor, bandwidth, channel allocations, *etc.*). Chime [8] selects the optimal channel frequency to improve SNRs. LMAC [9] reduces packet collisions at MAC layer to increase goodput. Standard LoRaWAN (ADR) [30] and some studies (DyLoRa [28]) adjust LoRa TX power and SF based on channel conditions to improve reliability and energy efficiency. However, FCC limits the maximum TX power. If communications with the highest TX power and the largest SF still suffer poor reliability, XCopy can provide further improvements.

Some works [2, 7, 17, 25] aim to improve LoRa communication over weak links. Falcon [40] extends LoRa coverage to the longest interference range by actively interfering strong LoRa transmissions in a way similar to On-Off Keying

(OOK) to convey information. OPR [2] exploits the heterogeneous bit errors at multiple gateways to recover a packet. NELoRa [25] employs deep learning techniques to enhance demodulation of weak LoRa signals. In contrast, XCopy improves signal quality before demodulation/decoding. Charm [7] combines weak signals of multiple gateways to enhance packet decoding. MAlLoRa [17] exploits the multi-antenna capability of LoRa gateways to boost weak packet reception. However, the signal processing methods in existing works cannot be applied to XCopy which addresses unique challenges involved in signal combining among retransmitted packets in ultra-low SNRs.

LoRa backscatter systems send data by backscattering frequency-shifted LoRa chirps. The backscattered LoRa signals become much weaker than active LoRa signals [14, 15, 22, 23, 32, 39]. Prior works (*e.g.*, PLoRa [32], Aloha [13], Saiyan [12], *etc.*) develop novel algorithms and customized hardware to facilitate the detection and decoding of weak backscatter signals. XCopy differs from these works in that it aims to detect and decode weak LoRa packets rather than backscattered signals.

Our work is also related to researches on reliable data transfer in wireless systems (*e.g.*, WLAN and cellular) [1, 6, 19, 31, 50]. As retransmissions introduce extra overhead, optimized strategies have been proposed to improve retransmission efficiency. Some works [27, 36, 50, 52] propose to only retransmit the symbols affected by corrupted bits to increase retransmission efficiency. PPR [21], SOFT [42], and LEAD [18] develop data correction schemes to recover error bits for corrupted packets by using advanced coding schemes or cross-layer optimization. To the best of our knowledge, XCopy is the first study on coherent combining of retransmitted LoRa packets that can work in ultra-low SNR scenarios.

7 CONCLUSION

This paper elaborates the design principle and implementation detail of XCopy, which boosts weak packet reception performance by coherently combining multiple retransmitted packets for reliable LoRa communication over harsh wireless channels. XCopy puts together a set of novel techniques to identify the signal copies of the same packet in the presence of interfering packets, and accurately align the signal copies by mitigating the impact of time, frequency as well as phase variations so as to ensure constructive coherent combining and improved signal strength. XCopy complements the latest advances in LoRa communication by exploiting rather than discarding the partial information hidden in signal copies received over weak links, thereby turning the weak signals that would otherwise be wasted into valuable signal strength improvement of more than 10 dB. The SNR gain could translate to higher throughput, longer coverage range, stronger reliability of LoRa communication.

8 ACKNOWLEDGEMENT

We thank the anonymous shepherd and reviewers for their helpful comments. This work is supported in part by Hong Kong GRF under grant 15218022 and PolyU 152165/19E, in part by the Start-up Fund for RAP under the Strategic Hiring Scheme of Hong Kong PolyU under grant P0036217, in part by the National Nature Science Foundation of China under grant 62102336, and in part by Ministry of Education Singapore MOE AcRF Tier 2 MOE-T2EP20220-0004. Yuanqing Zheng is the corresponding author.

REFERENCES

- [1] Carlos A. Astudillo, Fernando H. S. Pereira, and Nelson L. S. da Fonseca. 2019. Probabilistic Retransmissions for the Random Access Procedure in Cellular IoT Networks. In *ICC 2019 - 2019 IEEE International Conference on Communications (ICC)*. 1–7. <https://doi.org/10.1109/ICC.2019.8761468>
- [2] 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 (Toronto, Canada) (MobiSys '20)*. Association for Computing Machinery, New York, NY, USA, 247–259.
- [3] Atul Bansal, Akshay Gadre, Vaibhav Singh, Anthony Rowe, Bob Iannucci, and Swarun Kumar. 2021. Owl: Accurate lora localization using the tv whitespaces. In *Proceedings of the 20th International Conference on Information Processing in Sensor Networks (co-located with CPS-IoT Week 2021)*. 148–162.
- [4] Carolyn Bernier, François Dehmas, and Nicolas Deparis. 2020. Low Complexity LoRa Frame Synchronization for Ultra-Low Power Software-Defined Radios. *IEEE Transactions on Communications* 68, 5 (2020), 3140–3152. <https://doi.org/10.1109/TCOMM.2020.2974464>
- [5] Zhaoxin Chang, Fusang Zhang, Jie Xiong, Junqi Ma, Beihong Jin, and Daqing Zhang. 2022. Sensor-free Soil Moisture Sensing Using LoRa Signals. *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies* 6, 2 (2022), 1–27.
- [6] Wei Dong, Jie Yu, and Xiaojin Liu. 2015. CARE: Corruption-Aware Retransmission with Adaptive Coding for the Low-Power Wireless. In *2015 IEEE 23rd International Conference on Network Protocols (ICNP)*. 235–244. <https://doi.org/10.1109/ICNP.2015.12>
- [7] 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 *2018 17th ACM/IEEE International Conference on Information Processing in Sensor Networks (IPSN'18)*. IEEE, 60–71.
- [8] Akshay Gadre, Revathy Narayanan, Anh Luong, Anthony Rowe, Bob Iannucci, and Swarun Kumar. 2020. Frequency configuration for low-power wide-area networks in a heartbeat. In *17th USENIX Symposium on Networked Systems Design and Implementation (NSDI'20)*. 339–352.
- [9] Amalinda Gamage, Jansen Christian Liando, Chaojie Gu, Rui Tan, and Mo Li. 2020. LMAC: Efficient Carrier-Sense Multiple Access for LoRa. In *Proceedings of the 26th Annual International Conference on Mobile Computing and Networking*. Association for Computing Machinery, New York, NY, USA, Article 43, 13 pages.
- [10] Global Market Insight. 2022. *Industry Trends*. <https://www.gminsights.com/industry-analysis/low-power-wide-area-network-lpwan-market>.
- [11] Gr-LoRa GitHub community. 2022. *gr-lora projects*. Retrieved Jul 15, 2022 from <https://github.com/rpp0/gr-lora>
- [12] Xiuzhen Guo, Longfei Shangguan, Yuan He, Nan Jing, Jiacheng Zhang, Haotian Jiang, and Yunhao Liu. 2022. Saiyan: Design and Implementation of a Low-power Demodulator for LoRa Backscatter Systems. In *19th USENIX Symposium on Networked Systems Design and Implementation (NSDI 22)*. USENIX Association, Renton, WA, 437–451. <https://www.usenix.org/conference/nsdi22/presentation/guo>
- [13] Xiuzhen Guo, Longfei Shangguan, Yuan He, Jia Zhang, Haotian Jiang, Awais Ahmad Siddiqi, and Yunhao Liu. 2020. Aloba: Rethinking ON-OFF Keying Modulation for Ambient LoRa Backscatter. In *Proceedings of the 18th Conference on Embedded Networked Sensor Systems*. Association for Computing Machinery, New York, NY, USA, 192–204. <https://doi.org/10.1145/3384419.3430719>
- [14] Xiuzhen Guo, Longfei Shangguan, Yuan He, Jia Zhang, Haotian Jiang, Awais Ahmad Siddiqi, and Yunhao Liu. 2021. Efficient ambient LoRa backscatter with On-Off keying modulation. *IEEE/ACM Transactions on Networking* 30, 2 (2021), 641–654.
- [15] Mehrdad Hesar, Ali Najafi, and Shyamnath Gollakota. 2019. NetScatter: Enabling Large-Scale Backscatter Networks. In *16th USENIX Symposium on Networked Systems Design and Implementation (NSDI 19)*. USENIX Association, Boston, MA, 271–284. <https://www.usenix.org/conference/nsdi19/presentation/hessar>
- [16] Ningning Hou, Xianjin Xia, and Yuanqing Zheng. 2021. Jamming of LoRa PHY and Countermeasure. In *IEEE INFOCOM 2021 - IEEE Conference on Computer Communications*. 1–10. <https://doi.org/10.1109/INFOCOM42981.2021.9488774>
- [17] Ningning Hou, Xianjin Xia, and Yuanqing Zheng. 2022. Don't Miss Weak Packets: Boosting LoRa Reception with Antenna Diversities. In *IEEE INFOCOM 2022 - IEEE Conference on Computer Communications*. 1–10. <https://doi.org/10.1109/INFOCOM41043.2022.9155224>
- [18] Jun Huang, Yu Wang, and Guoliang Xing. 2013. LEAD: Leveraging Protocol Signatures for Improving Wireless Link Performance. In *Proceeding of the 11th Annual International Conference on Mobile Systems, Applications, and Services (Taipei, Taiwan) (MobiSys '13)*. Association for Computing Machinery, New York, NY, USA, 333–346. <https://doi.org/10.1145/2462456.2465429>
- [19] Jun Huang, Guoliang Xing, Jianwei Niu, and Shan Lin. 2015. CodeRepair: PHY-layer partial packet recovery without the pain. In *2015 IEEE Conference on Computer Communications (INFOCOM)*. 1463–1471. <https://doi.org/10.1109/INFOCOM.2015.7218524>
- [20] IoT Analytics. 2022. *5 things to know about the LPWAN market in 2021*. Retrieved Apr 12, 2022 from <https://iot-analytics.com/5-things-to-know-lpwan-market/>
- [21] Kyle Jamieson and Hari Balakrishnan. 2007. PPR: Partial Packet Recovery for Wireless Networks. In *Proceedings of the 2007 Conference on Applications, Technologies, Architectures, and Protocols for Computer Communications (Kyoto, Japan) (SIGCOMM '07)*. Association for Computing Machinery, New York, NY, USA, 409–420. <https://doi.org/10.1145/1282380.1282426>
- [22] Haotian Jiang, Jiacheng Zhang, Xiuzhen Guo, and Yuan He. 2021. Sense me on the ride: Accurate mobile sensing over a LoRa backscatter channel. In *Proceedings of the 19th ACM Conference on Embedded Networked Sensor Systems*. 125–137.
- [23] Jinyan Jiang, Zhenqiang Xu, Fan Dang, and Jiliang Wang. 2021. Long-Range Ambient LoRa Backscatter with Parallel Decoding. In *Proceedings of the 27th Annual International Conference on Mobile Computing and Networking (New Orleans, Louisiana) (MobiCom '21)*. Association for Computing Machinery, New York, NY, USA, 684–696. <https://doi.org/10.1145/3447993.3483261>
- [24] Zerina Kapetanovic, Deepak Vasishth, Tusher Chakraborty, Joshua R. Smith, and Ranveer Chandra. 2021. No Size Fits All: Automated Radio Configuration for LPWANs. <https://doi.org/10.48550/ARXIV.2109.05103>
- [25] 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>

- [26] Chenning Li, Xiuzhen Guo, Longfei Shangguan, Zhichao Cao, and Kyle Jamieson. 2022. {CurvingLoRa} to Boost {LoRa} Network Throughput via Concurrent Transmission. In *19th USENIX Symposium on Networked Systems Design and Implementation (NSDI 22)*. 879–895.
- [27] Tianji Li, Qiang Ni, David Malone, Douglas Leith, Yang Xiao, and Thierry Turletti. 2009. Aggregation With Fragment Retransmission for Very High-Speed WLANs. *IEEE/ACM Transactions on Networking* 17, 2 (2009), 591–604. <https://doi.org/10.1109/TNET.2009.2014654>
- [28] Yinghui Li, Jing Yang, and Jiliang Wang. 2020. DyLoRa: Towards Energy Efficient Dynamic LoRa Transmission Control. In *IEEE INFOCOM 2020 - IEEE Conference on Computer Communications*. 2312–2320.
- [29] Jansen C. Liando, Amalinda Gamage, Agustinus W. Tengourtius, and Mo Li. 2019. Known and Unknown Facts of LoRa: Experiences from a Large-Scale Measurement Study. *ACM Trans. Sen. Netw.* 15, 2, Article 16 (Feb. 2019), 35 pages.
- [30] LoRa Alliance. 2022. *LoRaWAN for Developer*. Retrieved Apr 19, 2022 from <https://loro-alliance.org/lorawan-for-developers>
- [31] Jiajue Ou, Yuanqing Zheng, and Mo Li. 2014. MISC: Merging incorrect symbols using constellation diversity for 802.11 retransmission. In *IEEE INFOCOM 2014 - IEEE Conference on Computer Communications*. 2472–2480. <https://doi.org/10.1109/INFOCOM.2014.6848193>
- [32] Yao Peng, Longfei Shangguan, Yue Hu, Yujie Qian, Xianshang Lin, Xiaojiang Chen, Dingyi Fang, and Kyle Jamieson. 2018. PLoRa: A Passive Long-Range Data Network from Ambient LoRa Transmissions. In *Proceedings of the 2018 Conference of the ACM Special Interest Group on Data Communication (Budapest, Hungary) (SIGCOMM '18)*. Association for Computing Machinery, New York, NY, USA, 147–160. <https://doi.org/10.1145/3230543.3230567>
- [33] Juha Petajajarvi, Konstantin Mikhaylov, Antti Roivainen, Tuomo Hanninen, and Marko Pettissalo. 2015. On the coverage of LPWANs: range evaluation and channel attenuation model for LoRa technology. In *2015 14th International Conference on ITS Telecommunications (ITST)*. 55–59. <https://doi.org/10.1109/ITST.2015.7377400>
- [34] 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.
- [35] Brecht Reynders, Wannes Meert, and Sofie Pollin. 2017. Power and spreading factor control in low power wide area networks. In *2017 IEEE International Conference on Communications (ICC)*. 1–6.
- [36] Eric Rozner, Anand Padmanabha Iyer, Yogita Mehta, Lili Qiu, and Mansoor Jafry. 2007. ER: Efficient Retransmission Scheme for Wireless LANs. In *Proceedings of the 2007 ACM CoNEXT Conference (New York, New York) (CoNEXT '07)*. Association for Computing Machinery, New York, NY, USA, Article 8, 12 pages. <https://doi.org/10.1145/1364654.1364665>
- [37] Semtech. 2022. *Semtech SX1276: 137MHz to 1020MHz Long Range Low Power Transceiver*. Retrieved Jul 15, 2022 from <https://www.semtech.com/products/wireless-rf/loro-transceivers/sx1276>
- [38] Madoune R. Seye, Bassirou Ngom, Bamba Gueye, and Moussa Diallo. 2018. A Study of LoRa Coverage: Range Evaluation and Channel Attenuation Model. In *2018 1st International Conference on Smart Cities and Communities (SCCIC)*. 1–4. <https://doi.org/10.1109/SCCIC.2018.8584548>
- [39] Vamsi Talla, Mehrdad Hesar, Bryce Kellogg, Ali Najafi, Joshua R. Smith, and Shyamnath Gollakota. 2017. LoRa Backscatter: Enabling The Vision of Ubiquitous Connectivity. *Proc. ACM Interact. Mob. Wearable Ubiquitous Technol.* 1, 3, Article 105 (sep 2017), 24 pages. <https://doi.org/10.1145/3130970>
- [40] Shuai Tong, Zilin Shen, Yunhao Liu, and Jiliang Wang. 2021. Combatting Link Dynamics for Reliable Lora Connection in Urban Settings. In *Proceedings of the 27th Annual International Conference on Mobile Computing and Networking (MobiCom '21)*. Association for Computing Machinery, New York, NY, USA, 642–655.
- [41] Xiong Wang, Linghe Kong, Zucheng Wu, Long Cheng, Chenren Xu, and Guihai Chen. 2020. SLoRa: towards secure LoRa communications with fine-grained physical layer features. In *Proceedings of the 18th Conference on Embedded Networked Sensor Systems*. 258–270.
- [42] Grace R. Woo, Pouya Kheradpour, Dawei Shen, and Dina Katabi. 2007. Beyond the Bits: Cooperative Packet Recovery Using Physical Layer Information. In *Proceedings of the 13th Annual ACM International Conference on Mobile Computing and Networking (Montréal, Québec, Canada) (MobiCom '07)*. Association for Computing Machinery, New York, NY, USA, 147–158. <https://doi.org/10.1145/1287853.1287871>
- [43] Xianjin Xia, Qianwu Chen, Ningning Hou, and Yuanqing Zheng. 2023. HyLink: Towards High Throughput LPWANs with LoRa Compatible Communication. In *Proceedings of the 20th ACM Conference on Embedded Networked Sensor Systems (Boston, Massachusetts) (SenSys '22)*. Association for Computing Machinery, New York, NY, USA, 578–591. <https://doi.org/10.1145/3560905.3568516>
- [44] Xianjin Xia, Ningning Hou, Yuanqing Zheng, and Tao Gu. 2021. PCube: Scaling LoRa Concurrent Transmissions with Reception Diversities. In *Proceedings of the 27th Annual International Conference on Mobile Computing and Networking (New Orleans, Louisiana) (MobiCom '21)*. Association for Computing Machinery, New York, NY, USA, 670–683. <https://doi.org/10.1145/3447993.3483268>
- [45] Xianjin Xia and Yuanqing Zheng. 2020. Connecting LoRaWANs Deep Inside a Building. In *Proceedings of the 7th ACM International Conference on Systems for Energy-Efficient Buildings, Cities, and Transportation (Virtual Event, Japan) (BuildSys '20)*. Association for Computing Machinery, New York, NY, USA, 312–313. <https://doi.org/10.1145/3408308.3431118>
- [46] Xianjin Xia, Yuanqing Zheng, and Tao Gu. 2020. FTrack: Parallel decoding for LoRa transmissions. *IEEE/ACM Transactions on Networking* 28, 6 (2020), 2573–2586.
- [47] Xianjin Xia, Yuanqing Zheng, and Tao Gu. 2020. LiteNap: Downclocking LoRa Reception. In *IEEE INFOCOM 2020 - IEEE Conference on Computer Communications*. 2321–2330. <https://doi.org/10.1109/INFOCOM41043.2020.9155224>
- [48] Binbin Xie, Yuqing Yin, and Jie Xiong. 2021. Pushing the limits of long range wireless sensing with lora. *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies* 5, 3 (2021), 1–21.
- [49] Weitao Xu, Jun Young Kim, Walter Huang, Salil S. Kanhere, Sanjay K. Jha, and Wen Hu. 2020. Measurement, Characterization, and Modeling of LoRa Technology in Multifloor Buildings. *IEEE Internet of Things Journal* 7, 1 (2020), 298–310. <https://doi.org/10.1109/JIOT.2019.2946900>
- [50] Jeongho Yeo, Sungjin Park, Jinyoung Oh, Younsun Kim, and Juho Lee. 2017. Partial Retransmission Scheme for HARQ Enhancement in 5G Wireless Communications. In *2017 IEEE Globecom Workshops (GC Wkshps)*. 1–5. <https://doi.org/10.1109/GLOCOMW.2017.8269132>
- [51] Fusang Zhang, Zhaoxin Chang, Jie Xiong, Rong Zheng, Junqi Ma, Kai Niu, Beihong Jin, and Daqing Zhang. 2021. Unlocking the beamforming potential of lora for long-range multi-target respiration sensing. *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies* 5, 2 (2021), 1–25.
- [52] Jiansong Zhang, Haichen Shen, Kun Tan, Ranveer Chandra, Yongguang Zhang, and Qian Zhang. 2012. Frame Retransmissions Considered Harmful: Improving Spectrum Efficiency Using Micro-ACKs. In *Proceedings of the 18th Annual International Conference on Mobile Computing and Networking (Istanbul, Turkey) (Mobicom '12)*. Association for Computing Machinery, New York, NY, USA, 89–100. <https://doi.org/10.1145/2348543.2348557>