BlindSpot: Efficient Single-Node Selective Jamming for LoRaWAN

Vincenz Mechler
Secure Mobile Networking Lab
TU Darmstadt, Germany
vmechler@seemoo.tu-darmstadt.de

Frank Hessel
Secure Mobile Networking Lab
TU Darmstadt, Germany
fhessel@seemoo.tu-darmstadt.de

Matthias Hollick
Secure Mobile Networking Lab
TU Darmstadt, Germany
mhollick@seemoo.tu-darmstadt.de

Bastian Bloessl
Secure Mobile Networking Lab
TU Darmstadt, Germany
bbloessl@seemoo.tu-darmstadt.de

ABSTRACT

LoRaWAN has become a widely adopted, cost-effective solution for Low-Power Wide-Area Networks (LPWANs), bridging the gap between short-range wireless protocols and high-power cellular networks. Its affordable hardware and robust physical layer make it a key enabler for Internet of Things (IoT) applications across sectors like agriculture, smart cities, and industrial automation—domains where security is of central importance. In this paper, we present BlindSpot, a novel jamming attack that enables efficient selective jamming of LoRaWAN gateways. Unlike traditional approaches that rely on creating high-power interference, BlindSpot exploits the limited number of demodulation paths in LoRaWAN gateways to continuously occupy the gateways with fabricated frames, blinding them for any other legitimate transmissions. Compared to existing approaches, this reduces the attacker's power requirements and allows them to decode the legitimate transmissions with a high probability. Selectively retransmitting these frames, the attacker has precise control over which transmissions can be decoded by the gateway. Using a Software-Defined Radio (SDR)-based LoRa transceiver, we demonstrate the effectiveness of the attack against commercial LoRaWAN gateways and propose detection and mitigation strategies to improve the security of LoRaWAN deployments.

CCS CONCEPTS

• **Networks** → **Network experimentation**; **Network security**.

KEYWORDS

LoRaWAN; Jamming; Denial of Service; Software-Defined Radio

ACM Reference Format:
Vincenz Mechler, Frank Hessel, Matthias Hollick, and Bastian Bloessl. 2025. BlindSpot: Efficient Single-Node Selective Jamming for LoRaWAN. In *18th ACM Conference on Security and Privacy in Wireless and Mobile Networks (WiSec 2025), June 30-July 3, 2025, Arlington, VA, USA*. ACM, New York, NY, USA, 6 pages. https://doi.org/10.1145/3734477.3734724

WiSec 2025, June 30-July 3, 2025, Arlington, VA, USA

ACM ISBN 979-8-4007-1530-3/2025/06
https://doi.org/10.1145/3734477.3734724

1 INTRODUCTION

LoRaWAN has become a key technology in the growing IoT landscape, providing low-power, long-range connectivity for a wide range of applications, spanning industries such as smart cities, agriculture, healthcare, and industrial automation [7]. As LoRaWAN gets embedded in critical infrastructure and production processes, security becomes increasingly relevant. In the recent past, the technology faced significant challenges, as attacks such as eavesdropping, replay attacks, and gateway compromises pose serious risks to data integrity and network reliability [5, 11].

Following this line of research, we investigate the availability of LoRaWAN networks and describe a novel selective jamming attack on LoRaWAN gateways. In contrast to existing approaches that rely on creating high-power interference [1, 3, 6], we exploit limitations of LoRaWAN gateway transceiver chips that only feature a limited number of demodulation paths. Even the most advanced gateways are only able to receive up to 16 frames in parallel [12]. By keeping these demodulation paths synchronized on fabricated frames, the gateway becomes blind to any other traffic. This enables an efficient Denial-of-Service (DoS) attack with a high success rate, since the attacker does not rely on overpowering the legitimate uplink transmissions but instead causes temporary resource exhaustion at the gateway. Therefore, the attacker's transmit power can be significantly lower than with traditional jamming attacks.

We show that half-duplex attackers can decode the legitimate transmissions with a high probability. This allows them to retransmit selected frames, providing functionality similar to a selective jammer. The ability to receive and replay legitimate transmissions also allows delaying or reordering, which serve as enablers for advanced attacks. Using an SDR implementation, we demonstrate the effectiveness of the attack against commercial LoRaWAN gateways and discuss how it can be used to enable attacks against the LoRaWAN MAC layer. Finally, we propose detection and mitigation strategies to improve the security of LoRaWAN deployments.

Our contributions can be summarized as follows:

- We present BlindSpot, a new selective jamming attack that does not rely on creating interference, increasing the efficiency and relaxing the constraints on the attacker's position.
- We implement the attack on an SDR and evaluate it against state-of-the-art commercial LoRaWAN gateways.
- We propose detection and mitigation strategies for BlindSpot to improve the security of LoRaWAN deployments.

2 RELATED WORK

Jamming LoRa is inherently challenging, as the Chirp Spread Spectrum (CSS)-based modulation is robust against noise and interference. This robustness stems primarily from the coding gain of the chirps, which depends on the Spreading Factor (SF) as the central parameter of LoRa that defines the duration of a chirp and, therefore, the robustness of the transmission. For example, a frame transmitted with SF 10 uses symbols that encode 10 bit with 2^{10} baseband samples. Increasing the SF by one adds one bit per symbol and doubles the symbol duration and, therefore, the coding gain. Using this approach, LoRa frames can be received even at negative Signal-to-Noise Ratios (SNRs), i.e., when the signal is below the noise floor.

2.1 Jamming

Several jamming attacks against LoRa have been proposed [1, 3, 6]. They show that the most efficient strategy is to replicate the CSS modulation of the signal that should be jammed. Using this approach, it is possible to corrupt the signal already at a Signal-to-Interference Ratio (SIR) of −3 dB [2, 6]. As a further optimization, it is possible to synchronize the symbol timing of the jamming signal with the targeted transmission, which further reduces this margin. This approach can yield significant symbol error rates up to an SIR of 5 dB [6]. Transmitting only the necessary amount of jamming symbols to corrupt the targeted frame can further reduce the detectability and energy consumption of the attacker [3].

Yet, these methods rely on creating interference to the targeted transmission and, thus, require the jamming signal to be strong enough to corrupt the target frame. In contrast to that, BlindSpot merely requires that the attacker can be received by the gateway. As a result, our attack can cover a larger area and is less sensitive to the placement of end devices, gateways, and the attacking node.

2.2 Truncate after Preamble

Our attack relies on resource exhaustion, keeping the demodulation paths of the LoRaWAN gateway busy decoding fabricated frames. To minimize self-interference and exposure of the attacker, we adopt the truncate-after-preamble attack of Gvozdenovic et al. [4], which was initially presented for WiFi and ZigBee. The idea is that an attacker only transmits a preamble and frame length field (set to a maximum-sized frame) but not the actual payload symbols, causing the receiver to synchronize on the frame and continue to demodulate symbols for the specified payload length. This occupies resources and wastes energy at the receiver. In [10], we showed that this attack is also applicable to LoRa receivers. Using truncated frames, we can occupy all demodulation paths of a LoRaWAN gateway and hide further transmissions from the gateway, creating a covert communication channel.

In this work, we extend this mechanism to overload a gateway continuously rather than just before a planned transmission at a known point in time. We, furthermore, optimize the process to a level, where a half-duplex node can receive the legitimate frames while overloading the gateway. This allows us to retransmit a subset of the received frames, enabling an effect similar to selective jamming.

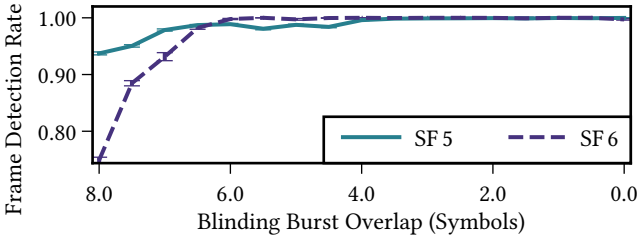

Figure 1: Detection rate of blinding frames over overlap with the end of the previous frame.

3 RECEIVE CHAIN OVERLOADING

Even the latest LoRaWAN gateway transceivers have much fewer demodulation paths than there are subchannels in typical LoRaWAN deployments. With 16 demodulation paths, the Semtech SX1302 gateway transceiver [12] has most but still significantly fewer than the 48 subchannels typically used in the EU863-870 frequency plan (eight channels with six possible SFs per channel). Transmitting 16 frames on different subchannels, therefore, blocks all available demodulation paths of the targeted gateway. Exploiting this, an attacker can continuously overload the gateway, causing any additional frames to be dropped. Combining this with the truncate-after-preamble attack, the attacker only has to transmit the preamble and physical layer headers but not the actual payload.

Since all other commercial gateway transceivers have less than 16 demodulation paths, the proposed attack signal can saturate them as well. It is, therefore, effective against all commercially available LoRaWAN gateways.

3.1 Resource Allocation Pattern

To optimize the effectiveness of the attack, modulation and timing of the blinding bursts are essential. There are two general strategies to allocate these bursts, either by channel or by SF.

If we use two channels to send blinding bursts with SF 5 to SF 12, the resulting self-interference will be concentrated on the selected channels. However, utilizing all eight SFs would require us to transmit for relatively long intervals, as the higher SFs have exponentially longer symbol durations. We, instead, allocate certain SFs for blinding and transmit blinding bursts across all channels. This minimizes the duty cycle and, therefore, exposure of the attacker. The higher, more robust SFs (i.e., SF 11 and SF 12) can maximize the probability that the gateway picks up the fabricated frames. However, this configuration again requires long transmit intervals, during which a half-duplex attacker would miss any signals.

For that reason, we prefer to allocate the lowest SFs (i.e., SF 5 and SF 6) for blinding. While this requires frequent retransmission of the blinding bursts, it minimizes burst durations at a similar duty cycle. As we show in Section 4, this enables reliable reception of frames with higher SFs.

3.2 Timing Uncertainty

The main challenge for the practical realization of the attack is to chain blinding bursts back-to-back, ensuring that no legitimate frame can be received. When the gateway frees its demodulation paths, there can be a race between a legitimate frame and the next

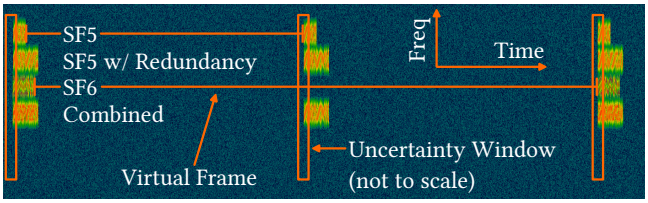

Figure 2: Spectrogram of the blinding pattern components, shown for a single channel. The *Combined* signal is the superposition of the redundant SF 5 headers and the SF 6 headers. It is sent on the eight channels, blocking 16 demodulation paths in total.

blinding bursts. While this time interval is small, it is particularly important, since once a legitimate frame is received, the blinding bursts are desynchronized, which opens a time window after the legitimate frame where the gateway is not blinded.

We noticed that the gateway transceiver exhibits some uncertainty regarding frame detection and the assignment of demodulation paths. To better understand this effect, we conducted experiments with a commercial gateway based on an SX1302 transceiver. Using an SDR, we continuously transmit blinding bursts with variable offsets and record the detection rate of the blinding frames.

The results are depicted in Figure 1. An overlap of two, for example, means that we start the new blinding frame two symbols before the current frame ends. In this and the following plots, the error bars are confidence intervals of the mean for a confidence level of 95 %. The different shapes of the curves for SF 5 and SF 6 can be explained by the different number of preamble symbols (twelve and eight, respectively). The results show that nearly the whole preamble of a frame can overlap with the previous transmission and still have a chance to get picked up by the receiver. Yet, since the plot is no step function, there is a tradeoff with regard to the spacing of the bursts: Sending the next blinding frame too early increases the chance that it might be missed, which opens a time window until the next burst for an original transmission to get through. However, sending the next blinding frame too late leads to a time window in-between bursts where the gateway is not blinded. We found that for an overlap below four symbols, detection is very reliable for both SFs. In the following experiments, we use an overlap of three symbols to the previous blinding burst.

3.3 Blinding Burst Optimizations

To maximize the blinding duration with respect to the duration of the preamble and the header of the blinding frame, we set the payload length of the SF 6 frames to the maximum of 255 Byte, the code rate to 4/8 (i.e., the highest level of redundancy), and enable the payload Cyclic Redundancy Check (CRC). As the air time of the SF 5 and SF 6 frames are not exact multiples, we set the payload length of the SF 5 frames to 210 Byte. Using this configuration, every second SF 5 blinding burst aligns with the SF 6 bursts, keeping the overall transmit duration at a minimum.

However, transmitting blinding preambles and headers on two SFs simultaneously adds interference and reduces the Signal-to-Interference-plus-Noise Ratio (SINR). Thus, some of the blinding

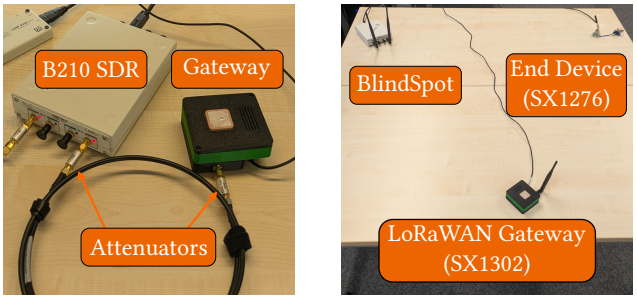

(a) Commercial gateway connected to an SDR-based attacker.

(b) Over-the-air experiment setup.

Figure 3: Evaluation setup. On the left: SX1302-based LoRaWAN gateway connected to USRP B210 SDR via cable and attenuators. On the right: commercial end device, commercial gateway, and SDR-based attacker node.

frames might not be received even in good channel conditions. To make blinding more reliable, we use the longer duration of the SF 6 headers to always transmit two SF 5 headers back-to-back, with the second header using a reduced payload length to align to the same frame end. This increases the robustness of the more sensitive SF 5, while only minimally affecting the overall transmit duration. The resulting pattern is shown in Figure 2. Since only the periodic bursts containing preamble symbols and headers have to be transmitted, it keeps the exposure of the attacker low at a duty cycle of 7.4 %.

3.4 Real-World Experiments

To evaluate the robustness of our attack, we conduct experiments with a commercial LoRaWAN gateway using the Semtech SX1302, a state-of-the-art transceiver capable of receiving 16 frames in parallel. An overview of the setup is shown in Figure 3a. Our attacker node is based on a USRP B210 SDR, with the signal processing implemented in software using FutureSDR.[1] The attacker is connected to the gateway via cable with 60 dB in-line attenuation.

As discussed in Section 3.1, we transmit the blinding bursts on all eight channels with SF 5 and SF 6. Legitimate uplink frames are generated with a fixed payload length of 16 Byte and uniformly distributed interarrival times to result in a given channel occupation.

The resulting Frame Reception Rate (FRR) of the benign frames is depicted in Figure 4, where a lower FRR implies better blinding performance. Even though the blinding frames have a reception rate close to 100 % when there is no other traffic, there is a small chance that the gateway detects a frame instead of the next blinding burst, as discussed in Section 3.2. Additionally, interference from legitimate frames can impact the blinding performance: if the beginning of the preamble of a legitimate frame overlaps with a blinding burst, there is a chance that the interference disturbs the blinding, while detection of the legitimate frame triggers only after the blinding burst in the now unblinded interval. Therefore, the blinding performance decreases with rising SIR, as can be seen in Figure 4 (a). This occurs with a higher probability for higher SFs due to the longer preamble duration. While the remainder of the legitimate

[1] https://www.futuresdr.org/

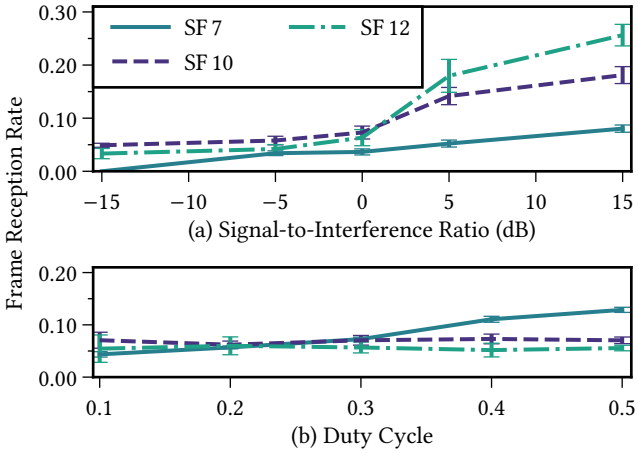

Figure 4: Reception rate of legitimate frames per SF over SIR (top, single frames) and duty cycle (bottom, 0 dB SIR) while continually blinding the gateway with SF 5 and SF 6. Lower FRR means higher blinding success.

frame may also interfere with the detection of the blinding bursts, this does not affect the performance under low channel utilization: the interfering frame has already been dropped, and if no other legitimate frame arrives until after the next interference-free blinding burst, the blinding recovers. At a very low SIR, we can, furthermore, see the effect of regular jamming, as our blinding bursts interfere directly with the SF 7 frames despite using a different SF.

However, blinding failures tend to cascade: if a legitimate frame is detected instead of a blinding frame, the respective demodulation path will stay open in the time between the end of the legitimate frame and the beginning of the next blinding burst. The probability of any frame arriving in this unblinded interval, and thus propagating the failure, rises with the channel utilization. Therefore, higher channel utilization leads to a higher detection rate of legitimate frames and, as a result, a lower blinding success rate, as shown in Figure 4 (b). Here, a duty cycle of 0.5 at SF 7 means we generate collision-free SF 7 frames on random channels with an accumulated duration of half the experiment duration. With higher SFs, we transmit exponentially fewer frames at a similar duty cycle. Therefore, the cascading effect is lower for higher SFs, yielding a better continuous blinding success rate at a similar channel utilization.

As we have no way of knowing when a legitimate frame has been detected instead of a blinding frame, we cannot mitigate this effect without allocating more logical subchannels to blinding or increasing the frequency of blinding bursts with overlapping coverage. While this would increase the overall blinding success rate, it increases the exposure of the attacker and compromises the ability to receive legitimate frames while executing the blinding attack.

In summary, we reach a total blinding success or Frame Error Rate (FER) of the legitimate traffic of up to 96 % at an SIR of 0 dB, depending on the channel utilization and SF. This is close to the maximum jamming rate of interference-based attacks, which reportedly achieve 98 % FER under controlled conditions [1]. However, unsynchronized jamming based on interference with random LoRa symbols works reliably only up to −3 dB SIR and drops to ~3.5 % FER

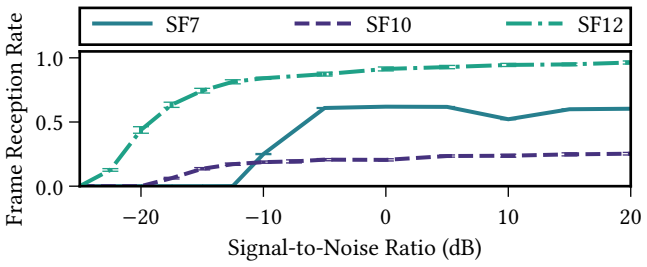

Figure 5: Reception rate of legitimate frames per SF over SNR in half-duplex mode. The signal has periodic gaps during the transmission slots of the blinding bursts.

at 0 dB SIR [6]. With tight synchronization to the targeted frames, state-of-the-art jamming attacks require less interference power but still become ineffective beyond 5 dB SIR [6]. In contrast, our attack maintains high performance even in unfavorable conditions, e.g., 74 % FER against SF 12 with an SIR of 15 dB.

4 FULL-DUPLEX FUNCTIONALITY

A jammer that can receive frames while preventing their reception at the targeted receiver is usually called a full-duplex jammer. Here, we show that we can achieve the same with a half-duplex device: as explained in Section 3.1, we transmit our blinding frames with the two lowest SFs and, therefore, only need to transmit short bursts to keep the targeted gateway in a blinded state. As a result, we can overhear most parts of the original transmission and the short gaps in the received signal can be compensated by the CSS modulation and Forward Error Correction (FEC).

To evaluate this mechanism, we conduct simulations to measure the reception rate of our attacker while blinding the gateway. The legitimate traffic is created similar to Section 3.4. To simulate half-duplex mode, we mask the signal during the periodic blinding bursts. We use our software-defined multi-channel LoRa receiver [10] to receive frames across all SFs and frequency channels.

The results for an Additive White Gaussian Noise (AWGN) channel are shown in Figure 5. We can see two effects: On the one hand, frames with a low SF have a lower probability of overlapping with a blinding burst and might, therefore, remain unaffected. SF 7, for example, starts to receive frames at the typical receive sensitivity for the SF but does not reach 100 % FRR, instead levelling out around 60 %, since any frame overlapping with a blinding burst is lost. On the other hand, higher SFs have a much larger probability of being affected by the gaps in the signal but can recover from overlapping with blinding bursts. At SF 12, the symbol duration is longer than the blinding bursts. Therefore, most frames can be received even at the low code rate of 4/5 used in LoRaWAN. The intermediary SFs like SF 10, however, have too short symbol durations to recover from the signal gaps but a high probability of overlapping with at least one blinding burst. Therefore, the FRR of SF 10 frames stays low even at high SNRs.

With increasing frame lengths, the probability that frames overlap with a blinding burst increases as well. As a result, the FRR of small SFs decreases, while the FRR of the higher SFs is less affected, given their robustness against burst interference.

5 SELECTIVE BLINDING

Combining sustained blinding with the capability to simultaneously receive legitimate uplink traffic, we can achieve an effect similar to selective jamming, which we call selective blinding. Instead of preventing the reception of only targeted frames, we start by continually blinding the targeted gateway, preventing the reception of uplink frames. At the same time, we receive the legitimate uplink traffic and decide whether we want to *jam* it or not: if a frame matches the criteria for being jammed, we simply drop it, since the gateway has already been prevented from receiving the original transmission. If the frame should not be jammed, we retransmit the captured frame after the current blinding interval.

5.1 Capture and Replay

To replay a captured uplink frame, we queue it in place of a blinding frame at the respective SF and channel as the original transmission. In LoRaWAN 1.1, these transmission parameters are part of the Message Integrity Code (MIC) and, therefore, cannot be changed without invalidating the MIC and causing the network server to reject the transmission.

In the case of unconfirmed uplinks, no further action is necessary. Replaying confirmed uplinks is, however, more complicated. Due to the delay in capturing the legitimate frame before replaying it, the Acknowledgment Frame (ACK) from the gateway is likely to miss the end device's receive window. If the ACK from the network server misses both receive windows of the end device, the end device will retransmit the uplink a configurable number of times since the rationale of using confirmed uplinks is to ensure reliable delivery. The LoRaWAN standard specifies that frame counters are not increased for retransmissions [8, Section 4.3.1.5]. Therefore, the downlink message carrying the ACK of the original confirmed uplink is still valid to acknowledge the retransmitted uplink message. We can, therefore, capture the ACK sent by the network server in response to the original uplink and replay it for the retransmission.

5.1.1 End Device Configuration. As we have seen in Section 4, capturing legitimate uplink frames works reliably only for the higher SFs. In LoRaWAN, the end devices can autonomously select the data rate and, therefore, the SFs used for uplink transmissions. While this can be statically assigned by the network operator, it is recommended for static end devices to use Adaptive Data Rate (ADR) [8], a mechanism to dynamically adapt the data rate to the channel conditions. Due to LoRa's high latency and low data rate and the resulting slow convergence of adaptive algorithms, this is not recommended for highly mobile end devices, where the most robust configuration is typically used to avoid outages.

With our attack, lower SFs cannot be reliably replayed but are still jammed. The specification mandates the following behavior in the case that requested ADR updates do not arrive at the end device: after a certain count of unacknowledged uplinks, the end device has to switch to the next lower data rate [8, Section 4.3.1.1]. Therefore, the inability of our attack to replay frames on low SFs will eventually push the end devices to use lower data rates with higher SFs, i.e., to use a configuration that allows us to perform our selective blinding attack.

5.2 Over-the-Air Experiments

We finally show that our selective blinding attack can work in practice against real-world deployments. To this end, we intercept the uplinks from a commercial SX1276-based LoRaWAN end device that uses LoRaWAN 1.0.2 to communicate with a ChirpStack v4.8.1 network server. To avoid affecting other deployments, we perform our experiments with reduced transmit power in a shielded environment. The attacker transmits at approx. −45 dBm per subchannel, whereas the end device is attenuated by 30 dB to −16 dBm. The nodes are arranged in a triangle with ∼80 cm distance, as shown in Figure 3b. We transmit 300 unconfirmed LoRaWAN uplinks on data rate 5 of the EU863-870 frequency plan, i.e., on SF 7 at 125 kHz and successfully blind 93 % of the transmitted uplinks, which is in line with our results in Section 3.4. Furthermore, all frames received at our attacker node can be played back successfully to the network server.

6 DISCUSSION

Our attack operates on a completely different principle than traditional, interference-based jammers. This yields different trade-offs, which we explore in this section.

6.1 Feasibility

Using our attack, the DoS from jamming is based on the reception of blinding frames. This is possible if the attacker can exceed the minimum SNR threshold for the respective SFs at the gateway. In this paper, we mainly focus on using SF 5 and SF 6 for blinding, given their advantages with regard to selective jamming, allowing us to overhear frames while blinding the gateway. Yet, the DoS is also possible with higher spreading factors and only requires that the attacker is in range of the gateway, i.e., blinding could just as well be done from a large distance, using SF 11 and SF 12.

This is in contrast to traditional jammers that rely on overpowering the legitimate frame. In this case, the placement of the attacker depends on the distance of the node to the gateway. If the node is close to the gateway and, therefore, achieves a high signal level at the gateway, the attacker also has to be close to create sufficient interference to lower the SINR below reception level. Although not entirely independent of the SIR, our approach provides a clear advantage, allowing a more flexible placement of the attacker node.

Shifting from jamming to selective jamming, any strategy requires the attacker to overhear the node sending the legitimate frame. While this is all that it takes for traditional jammers, our approach has the additional requirement that we receive all original transmissions while actively transmitting to blind the gateway. This is due to the inverted strategy in regard to traditional jammers, i.e., we suppress all initial transmissions and have to retransmit the frames that we want to be received by the gateway. As discussed, this mandates blinding with low SFs and, therefore, requires the attacker to be close enough to the gateway for it to receive these SFs. Yet, this drawback of our approach is mitigated when we assume full-duplex capabilities for the attacker. In this case, we could use higher SFs for blinding, as they do not disturb the reception of the original transmissions.

6.2 Impact on LoRaWAN

BlindSpot is a novel, more robust way for selective jamming of LoRaWAN networks, which enables a multitude of attacks [5]. While many vulnerabilities of LoRaWAN 1.0 have been closed in LoRaWAN 1.1, some attacks are still possible, e.g., energy depletion attacks or forging of multicast messages. Given the long lifetime of LoRaWAN nodes, many networks still operate on LoRaWAN 1.0 and are, therefore, susceptible to ACK spoofing and replay attacks, which can be enabled by our jammer.

Furthermore, BlindSpot allows the attacker to modify the payload of intercepted frames, similarly to traditional full-duplex jamming. However, since messages in LoRaWAN are cryptographically protected, the attacker is not able to forge frames, but can only replay frames that were captured previously. Since frame counters prevent the reordering of messages, intercepted frames can later be replayed in the same order but with arbitrary timing. Apart from attacks against the MAC layer protocol itself, this allows, for example, to increase the interval between reported sensor readings, reducing the update rate and introducing a growing delay in the authenticated data arriving at the network server.

6.3 Countermeasures

While traditional jammers can be countered, to some extent, by techniques for the recovery of collided LoRa frames, BlindSpot remains unaffected, since it does not rely on interfering with the legitimate signals. However, the attack can be detected by observing the patterns of failed frame receptions at the gateway, and a more restrictive configuration can decrease its efficiency.

Application Layer. In their default configurations,[2] gateways do not report frame loss due to failed checksum verification. Therefore, the LoRaWAN server never receives any information about the blinding frames. This makes it hard for the network operator to analyze the anomalous behavior and recognize the attack. We, therefore, recommend enabling logging of frames with failed checksums, at least if there are issues with the network.

Gateway Configuration. The attack requires the targeted gateway to detect and decode the blinding frames. Since LoRaWAN uses only SF 7 through SF 12 in the predefined data rates [9], gateways could be configured to ignore SF 5 and SF 6. Yet, default configurations typically enable all available SFs (see previous examples). Having SF 5 and SF 6 available for blinding makes selective jamming easier than necessary. Thus, we recommend disabling unused SFs.

Gateway Transceivers. Finally, increasing the number of demodulation paths in future LoRa gateway transceivers could effectively mitigate the attack, since it builds on resource exhaustion, occupying the limited number of demodulation paths. Even the most advanced transceivers only provide up to 16 paths, which is way smaller than the number of subchannels in LoRaWAN. Considering the EU863-870 frequency plan, for example, there are 48 LoRaWAN subchannels. SDR-based LoRaWAN gateways, like the one presented by Yu et al. in [13] or our own in [10], are not affected by our attack, since they are not limited by the number of demodulation paths.

[2] https://github.com/RAKWireless/rak_common_for_gateway/, https://github.com/Lora-net/packet_forwarder/

7 CONCLUSIONS

In this paper, we presented BlindSpot, a novel attack against the availability of LoRaWAN gateways, which provides higher flexibility in the placement of the attacker node than traditional jamming attacks. It does not rely on direct interference but instead exploits the limited number of demodulation paths in LoRaWAN gateways to continuously occupy the gateways with fabricated frames, blinding them for any other transmissions. Practical evaluation using an SDR and a commercial LoRaWAN gateway showed the high success rate of our attack. By capturing frames while denying the reception at the gateway, they can later be replayed to enable selective jamming, even with a half-duplex device. Finally, we discussed the impact of our attack and proposed practical countermeasures for real-world LoRaWAN deployments.

REFERENCES

[1] E. Aras, N. Small, G. S. Ramachandran, S. Delbruel, W. Joosen, and D. Hughes. 2017. Selective Jamming of LoRaWAN using Commodity Hardware. In *14th EAI International Conference on Mobile and Ubiquitous Systems: Computing, Networking and Services*. ACM, Melbourne, Australia, (Nov. 2017). DOI: 10.1145/3144457.3144478.

[2] D. Croce, M. Gucciardo, S. Mangione, G. Santaromita, and I. Tinnirello. 2018. Impact of LoRa Imperfect Orthogonality: Analysis of Link-Level Performance. 22, 4, (Apr. 2018). DOI: 10.1109/LCOMM.2018.2797057.

[3] A. Dossa and E. M. Amhoud. 2025. Impact of Reactive Jamming Attacks on LoRaWAN: a Theoretical and Experimental Study. cs.NI 2501.18339. arXiv, (Jan. 2025). DOI: 10.48550/arXiv.2501.18339.

[4] S. Gvozdenovic, J. K. Becker, J. Mikulskis, and D. Starobinski. 2020. Truncate After Preamble: PHY-based Starvation Attacks on IoT Networks. In *13th ACM Conference on Security and Privacy in Wireless and Mobile Networks (WiSec 2020)*. ACM, Linz, Austria, (July 2020). DOI: 10.1145/3395351.3399356.

[5] F. Hessel, L. Almon, and M. Hollick. 2023. LoRaWAN Security: An Evolvable Survey on Vulnerabilities, Attacks and their Systematic Mitigation. *ACM Transactions on Sensor Networks*, 18, 4, (Mar. 2023). DOI: 10.1145/3561973.

[6] N. Hou, X. Xia, and Y. Zheng. 2023. Jamming of LoRa PHY and Countermeasure. *ACM Transactions on Sensor Networks*, 19, 4, (Nov. 2023). DOI: 10.1145/3583137.

[7] M. Jouhari, N. Saeed, M.-S. Alouini, and E. M. Amhoud. 2023. A Survey on Scalable LoRaWAN for Massive IoT: Recent Advances, Potentials, and Challenges. *IEEE Communications Surveys & Tutorials*, 25, 3, 1841–1876. DOI: 10.1109/COMST.2023.3274934.

[8] LoRa Alliance. 2017. LoRaWAN 1.1 Specification. Std. (Oct. 2017).

[9] LoRa Alliance. 2022. LoRaWAN Regional Parameters RP002-1.0.4. Std. (Sept. 2022).

[10] V. Mechler, M. Hollick, and B. Bloessl. 2024. Beyond Sensing: A High-Performance Software-Defined LoRa Gateway. In *30th Annual International Conference on Mobile Computing and Networking (MobiCom 2024), 18th ACM International Workshop on Wireless Network Testbeds, Experimental evaluation and Characterization (WiNTECH 2024)*. ACM, Washington, DC, (Nov. 2024). DOI: 10.1145/3636534.3697317.

[11] H. Noura, T. Hatoum, O. Salman, J.-P. Yaacoub, and A. Chehab. 2020. LoRaWAN Security Survey: Issues, Threats and Possible Mitigation Techniques. *Internet of Things*, 12. DOI: 10.1016/j.iot.2020.100303.

[12] Semtech. 2020. SX1302 LoRa Gateway Baseband Processor. Datasheet DS. SX1302.W.APP. Rev 1.2. (Oct. 2020).

[13] S. Yu, X. Xia, N. Hou, Y. Zheng, and T. Gu. 2024. Revolutionizing LoRa Gateway with XGate: Scalable Concurrent Transmission across Massive Logical Channels. In *30th Annual International Conference on Mobile Computing and Networking (MobiCom 2024)*. ACM, Washington, DC, (Nov. 2024). DOI: 10.1145/3636534.3649375.

This work has been co-funded by the LOEWE initiative (Hesse, Germany) within the emergenCITY [LOEWE/1/12/519/03/05.001(0016)/72] center, as well as the German Research Foundation (DFG) in the Collaborative Research Center (SFB) 1053 MAKI (Project-ID 210487104). The authors acknowledge the financial support by the Federal Ministry of Education and Research of Germany in the project "Open6GHub" (grant number: 16KISK014).