WIP: An Adaptive High Frequency Removal Attack to Bypass Pulse Fingerprinting in New-Gen LiDARs

Yuki Hayakawa[†], Takami Sato[‡], Ryo Suzuki[†], Kazuma Ikeda[†], Ozora Sako[†], Rokuto Nagata[†], Qi Alfred Chen[‡], and Kentaro Yoshioka[†]
[†]Keio University; [‡]University of California, Irvine

***Abstract*—LiDAR stands as a critical sensor in the realm of autonomous vehicles (AVs). Considering its safety and security criticality, recent studies have actively researched its security and warned of various safety implications against LiDAR spoofing attacks, which can cause critical safety implications on AVs by injecting ghost objects or removing legitimate objects from their detection. To defend against LiDAR spoofing attacks, pulse fingerprinting has been expected as one of the most promising countermeasures against LiDAR spoofing attacks, and recent research demonstrates its high defense capability, especially against object removal attacks. In this WIP paper, we report the progress in conducting further security analysis on pulse fingerprinting against LiDAR spoofing attacks. We design a novel adaptive attack strategy, the Adaptive High-Frequency Removal (A-HFR) attack, which can be effective against broader types of LiDARs than the existing HFR attacks. We evaluate the A-HFR attack on three commercial LiDAR with pulse fingerprinting and find that the A-HFR attack can successfully remove over 96% of the point cloud within a 20° horizontal and a 16° vertical angle. Our finding indicates that current pulse fingerprinting techniques might not be sufficiently robust to thwart spoofing attacks. We also discuss potential strategies to enhance the defensive efficacy of pulse fingerprinting against such attacks. This finding implies that the current pulse fingerprinting may not be an ultimate countermeasure against LiDAR spoofing attacks. We finally discuss our future plans.**

I. Introduction

LiDAR (Light Detection And Ranging) has been integrated into autonomous vehicles (AVs) as a critical sensor due to its capability to accurately map the surrounding environment in three dimensions. As of 2023, all 6 AV companies authorized to test on public roads in California have adopted LiDAR as a fundamental part of their perception systems. Reflecting its importance on AD systems, the security of LiDARs has been actively researched, especially for the vulnerability against LiDAR spoofing attacks [1]–[6]. These attacks involve projecting malicious lasers against LiDARs, compromising their distance measurement by overwriting the legitimate signals. Such attacks have significant safety implications for AD systems, particularly in deceiving LiDAR-based 3D object detectors. Object injection attacks [3], [5], [6] create ghost objects by injecting a set of fake points that mimic legitimate objects. Object removal attacks [4]–[6] make actual objects indetectable by erasing the legitimate points of the objects. However, existing LiDAR spoofing attacks still have a critical gap to be a real threat against recent AD vehicles, as pointed out in [6]. The majority of evaluations focus only on classic LiDAR models such as VLP-16 [7], overlooking many recent LiDARs, referred to as Next-Generation (or New-Generation) LiDARs (new-gen LiDARs) [6], [8]. New-gen LiDARs have advanced security-related features that demonstrate high defense and mitigation capability against LiDAR spoofing attacks. Among them, timing randomization shows notable defense capability that randomizes each laser emission interval of LiDAR. It can make the LiDAR scanning pattern unpredictable and thus make injecting points at designed locations virtually impossible. Despite this, Sato et al. [6] demonstrate that their HFR (High-Frequency Removal) attack is still effective even with the timing randomization, as it saturates all LiDAR measurements without needing to discern the scanning pattern. However, the HFR attack was not effective against the new-gen LiDARs with pulse fingerprinting that authenticates their lasers with an embedded fingerprint in the lasers.

In this WIP paper, we report our recent progress on the further security analysis of pulse fingerprinting in new-gen LiDARs. Our research addresses two gaps in previous studies [6]: (1) Prior work does not evaluate the attack with higher frequent pulses. As its name implies, higher attack frequency generally leads to higher attack effectiveness in the HFR attack; (2) prior work evaluated only one new-gen LiDAR with pulse fingerprinting. Their findings may not be generalizable. To address these gaps, we designed a new spoofing attack, the Adaptive HFR (A-HFR) attack, which can achieve 5 times higher frequency than the prior HFR attack even with the same laser hardware by adaptively changing the attack region. Specifically, we boost the frequency when the LiDAR is scanning the targeted objects, and we idle the attack when scanning the other areas to avoid overheating the diode emitting the attack laser. We evaluate the A-HFR attack on three widely used LiDAR models with pulse fingerprinting to further evaluate the generality of the A-HFR attack. Our preliminary results suggest that current LiDARs with pulse fingerprinting might not sufficiently counter the A-HFR attack. We finally discuss possible enhancements in fingerprint technology that could bolster defenses against such advanced spoofing techniques.

II. Background and Related Works

A. LiDAR Spoofing Attacks

The LiDAR spoofing attacks [1], [2], [5], [6] have demonstrated high attack effectiveness against LiDARs. Based on the attacker's capability, there are two types of LiDAR spoofing attacks: synchronized attacks and asynchronized attacks.

Symposium on Vehicles Security and Privacy (VehicleSec) 2024
26 February 2024, San Diego, CA, USA
ISBN 979-8-9894372-7-6
https://dx.doi.org/10.14722/vehiclesec.2024.23022
www.ndss-symposium.org

1) Synchronized Attacks: Synchronized attacks (Sync. attack) first learn the laser scanning pattern of the target LiDAR, understanding its scan timing and coverage. As the adversary can exactly predict how their attack changes the LiDAR point cloud, Sync. attack can artificially inject objects of various shapes and positions of objects such as pedestrians or vehicles [5], [6]. However, the Sync. attack prerequisites a deterministic LiDAR scanning pattern, a premise not typically valid for modern new-gen LiDARs. They often employ timing randomization, originally installed as an anti-interference feature to incorporate multiple LiDARs at close distances. Thus, all previous works can be only successful on first-generation LiDARs such as the VLP-16 [7], although these attacks can still inject hundreds of points at random locations [6].

2) Asynchronized Attacks: Asynchronized attacks (Async. attacks) operate regardless of the victim LiDAR's scan timing. There are three types of Async. attacks:

Relay attack [1] involves capturing and relaying the laser signals from the victim LiDAR. This method is theoretically effective against all LiDARs since the relayed laser is identical to the original. However, its threat to AVs is somewhat limited, as it can only create fake points further than the spoofer.

Saturation attack [2] aims to overwhelm the victim LiDAR by shooting it with a continuous wave laser. This can either mask legitimate laser signals or cause misinterpretations of the received data. The challenge with this attack lies in generating a continuous wave laser strong enough to obscure legitimate signals, which limits the attack angle and duration.

HFR attack [6] emits periodic pulses between 400 kHz and 5 MHz at the victim LiDAR. LiDARs generally prioritize the strongest of multiple valid pulses. If an attacker's pulse overshadows the genuine signal, the LiDAR is tricked into selecting the false signal. This causes the authentic point cloud to vanish, with noise replacing it across the LiDAR's range. The attack's effectiveness varies; it is potent against some new-gen LiDARs but less so against those with pulse fingerprinting.

B. LiDARs with Pulse Fingerprinting

Pulse fingerprinting [9] has shown a high defense capability against LiDAR spoofing attacks while it is also originally designed for the sake of anti-interference to operate multiple LiDARs at close distances. Pulse fingerprinting authenticates the reflected laser pulses and tries not to accept the laser reflection emitted by the other LiDARs. One of the most common implementations of pulse fingerprinting (e.g. Livox Mid-360 [10], Hesai XT32 [11], AT128 [12]) encodes its fingerprint into the interval between two consecutive pulses as depicted in Fig. 1. Thus, LiDARs with pulse fingerprinting need 2 pulses for a single distance measurement instead of 1 pulse as traditional LiDARs need. They then accept the signal only when the interval between the received pulses matches the firing interval. Otherwise, the LiDAR discards it. This mechanism enables them to be robust against interference as they can be selective for their own pulses. To account for nonidealities during operation, these LiDARs include a tolerance error time span T_α. For instance, with a pulse interval of 400 ns and a tolerance of 40 ns, the LiDAR accepts pulse pairs with intervals from 380 ns to 420 ns as valid. The interval between pulses is randomly set between a minimum (T_{min}) and maximum (T_{max}) value. The specific values for T_α, T_{min} and T_{max} vary depending on the LiDAR models.

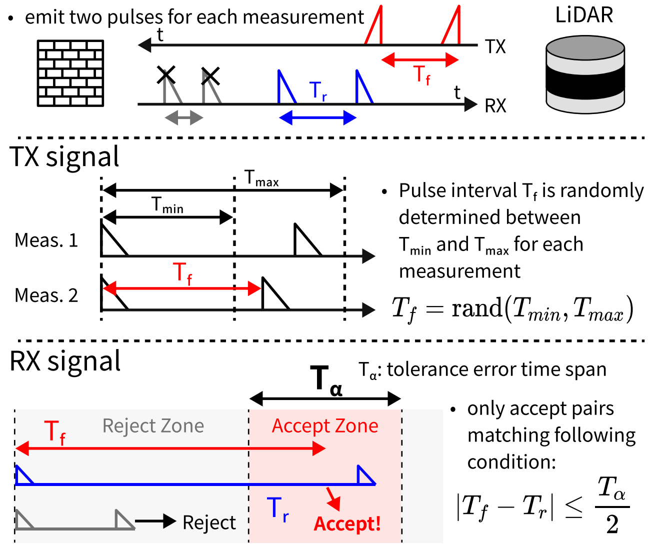

emit two pulses for each measurement
t
TX
LiDAR
T_f
t
T_r
RX
TX signal
T_{max}
T_{min}
Meas. 1
Meas. 2
T_f
Pulse interval T_f is randomly determined between T_{min} and T_{max} for each measurement
$T_f = \text{rand}(T_{min}, T_{max})$
RX signal
T_α
T_α: tolerance error time span
T_f
Reject Zone
Accept Zone
T_r
Reject
Accept!
only accept pairs matching following condition:
$|T_f - T_r| \leq \frac{T_\alpha}{2}$

Fig. 1: Operating principle of pulse fingerprinting. LiDAR fires two pulses per ranging measurement, and if the interval between the received pulses is within the tolerance error (T_α), the LiDAR accepts them as valid pulses. This interval is randomly determined between T_{min} and T_{max}.

C. Threat Model

As our A-HFR attack is an extension of the normal HFR attack [6], we follow the same threat model, which is also commonly used in other prior work [13]. Specifically, the attacker targets an AV either from the roadside or from a nearby vehicle in a nearby lane. While the object detection system used by the victim AV is treated as a black box, the model and characteristics of the LiDAR installed on the AV are assumed to be identifiable. The LiDAR scan pattern and ranging method can be easily known in its datasheet or through analysis of the emitted signal. This information is utilized by the attacker as prior knowledge. The primary attacker's goal is to compromise the LiDAR point clouds to make vehicles or pedestrians undetectable. In our preliminary analysis on the KITTI dataset [14], we found that over 95% of objects located more than 6 meters away fit within a spatial window of 20° horizontally and 16° vertically. We thus define the attack goal in this paper as removing all points within this horizontal and vertical range.

III. Methodology

A. Attack Concept to Bypass Pulse Fingerprinting

Pulse fingerprinting is a robust defense against existing LiDAR spoofing attacks due to its signal authentication that filters out malicious pulses. Yet, as Table II shows, its reliance on a single pulse interval for authentication in mass-produced LiDARs presents a vulnerability. If an attacker replicates this interval, these systems could be susceptible to spoofing.

This vulnerability is particularly pronounced in the context of HFR attacks, as illustrated in Fig. 2. Pulse fingerprinting's

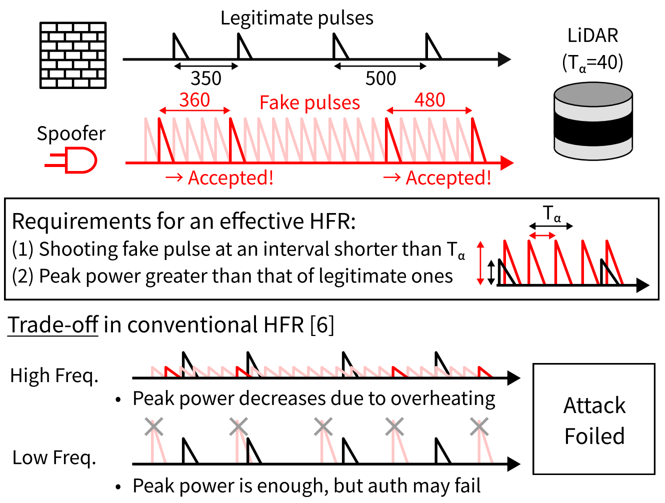
Legitimate pulses
350
500
LiDAR
(T_α=40)
Spoofer
360
Fake pulses
480
→ Accepted!
→ Accepted!
Requirements for an effective HFR:
(1) Shooting fake pulse at an interval shorter than T_α
(2) Peak power greater than that of legitimate ones
T_α
Trade-off in conventional HFR [6]
High Freq.
• Peak power decreases due to overheating
Low Freq.
• Peak power is enough, but auth may fail
Attack Foiled

Fig. 2: Our concept to bypass pulse fingerprinting by HFR attack. When authentication relies on a single pulse interval, the attack can bypass it by shooting fake pulses at intervals shorter than the tolerance error (T_α). However, conventional HFR faces limitations in simultaneously achieving high frequency and peak power due to overheating issues.

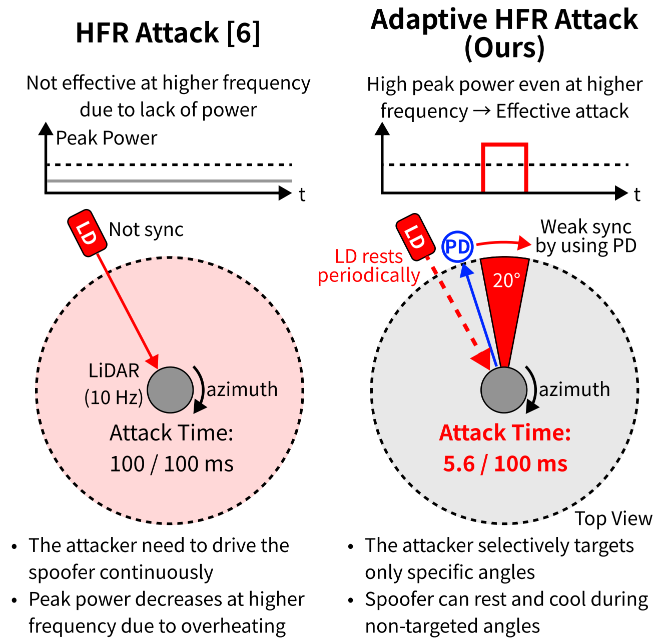
HFR Attack [6]
Not effective at higher frequency due to lack of power
Peak Power
t
LD
Not sync
LiDAR
(10 Hz)
azimuth
Attack Time:
100 / 100 ms
• The attacker need to drive the spoofer continuously
• Peak power decreases at higher frequency due to overheating
Adaptive HFR Attack (Ours)
High peak power even at higher frequency → Effective attack
t
LD
PD
Weak sync by using PD
LD rests periodically
20°
azimuth
Attack Time:
5.6 / 100 ms
Top View
• The attacker selectively targets only specific angles
• Spoofer can rest and cool during non-targeted angles

Fig. 3: Comparative illustration of the conventional HFR attack (left) versus the A-HFR attack (right). A-HFR can achieve high peak power at high frequencies by limiting the attack angle through weak synchronization.

random pulse intervals can be exploited during an HFR attack; if any two fake pulses coincide with the system's intervals, they can bypass authentication, with their success dependent on the LiDAR's tolerance error. To quantify this, we calculate the probability of a successful HFR attack, represented as p in Eq. 1. This equation holds when the attacker's pulse intervals are sufficiently shorter than the pulse interval emitted by the LiDAR (T_{max}).

$$p = \min\left(1, \frac{T_\alpha}{T_{\mathrm{HFR}}}\right) \tag{1}$$

Let T_α be the tolerance error time span of the pulse interval, and T_{HFR} be the interval between the attack pulses. Eq. 1 shows that, ideally, authentication can be bypassed in all measurements by shooting pulse light at an interval shorter than T_α. This implies that an HFR attack could effectively create a 'master key' for pulse fingerprinting by simply increasing the pulse frequency. For a spoofing attack to succeed, it's not enough to just bypass authentication; the fake pulses must have a higher peak power than that of the legitimate ones. LiDAR systems typically choose the strongest signal from multiple valid signals in a single measurement. Thus, to ensure the fake pulse is selected over the real one, attackers must ensure their pulses overpower the legitimate signals in peak power.

In summary, an effective HFR attack against pulse fingerprinting requires (1) an attack pulse interval shorter than T_α, and (2) peak power greater than that of the legitimate pulse. However, conventional HFR attack [6] faces a trade-off between pulse interval and peak power, making it challenging to fulfill both conditions simultaneously. For example, emitting high-frequency pulses (>5 MHz) reduces laser peak power due to overheating of the laser driver, failing to overcome legitimate pulses. Conversely, higher peak power attacks can only be realized at low frequencies, which do not pass authentication.

B. *Attack Design: Adaptive HFR Attack*

As shown in Fig. 3, we developed the Adaptive HFR (A-HFR) attack, which works on the existing spoofer hardware, to overcome the trade-off between pulse frequency and power. A-HFR employs *weak synchronization* to selectively target only specific angles within the victim LiDAR's scan. This selective targeting allows the spoofer to rest and cool during non-targeted angles, averting laser driver overheating. Hence, this method enables high-frequency, high-power attacks at specific angles required to bypass the fingerprinting authentication.

Unlike conventional asynchronous HFR attacks, A-HFR utilizes a photodiode (PD) for weak synchronization, to align the attack with the LiDAR scan. Facilitating the knowledge of LiDAR scanning patterns, we achieve angle-specific attacks by activating the spoofer only during desired angles. A-HFR's weak synchronization, which only approximates the LiDAR's scanning angle, tolerates minor inaccuracies, making it suitable for attacks on new-gen LiDARs with randomized intervals. This point is further elaborated in the following section.

1) Strategic Attack Angle Reduction: The perceptible horizontal angle of a rotating LiDAR's point cloud for objects like people or vehicles is just a fraction of the full 360° sweep. For objects over 6 m away, more than 95% fall within a 20° horizontal angle (§II-C). Rotating LiDARs, which stack lasers vertically within the device and rotate them horizontally, scan points with similar horizontal angles in quick succession. With a LiDAR rotating at 10 Hz, a full rotation takes 100 ms, and scanning a 20° horizontal section takes about 5.6 ms. Thus, by

TABLE I: Comparison of synchronization methods in LiDAR spoofing attacks.

	Sync [4]	Async [6]	**Weak-Sync (Ours)**
Required sync. accuracy	< 3 ns	N/A	$<$ **50 us**
Reducing attack angle	✓	-	✓
Effective to new-gen LiDARs	-	✓	✓

reducing an attack to just this 20° horizontal angle, the attack can be executed in less than 6% of the total time.

Reducing the spoofer's vertical attack angles can further reduce its operation time, as targeting the full vertical range seen by the victim LiDAR is generally unnecessary to remove objects like people or vehicles. For example, focusing on just 20° of a LiDAR's 40° vertical FoV can halve the attack duration. When combined with horizontal angle reduction, the time needed to attack both 20° horizontally and vertically is only 3% of the total LiDAR scan time.

2) Weak Synchronization Accuracy Requirement: A-HFR attack leverages a PD to capture pulses from the target LiDAR, enabling prediction of its scanning pattern. This *weak synchronization* strategy allows the attack to be timed when the LiDAR scans a specific angle. While PD-based synchronization is common in both Sync. attacks [3]–[6] and A-HFR attacks, their accuracy requirements differ significantly, as detailed in Table I. Weak synchronization is only to synchronize the scanning angle of the LiDAR, while prior synchronization is to synchronize each ranging start timing of the LiDAR. Sync. attacks demand nanosecond-level precision, with a 1 ns deviation altering the point cloud's position by 0.15 m. The PRA [4] requires sync accuracy under 3 ns for the effective removal attack. However, new-gen LiDARs with timing randomization complicate precise synchronization due to intervals randomized up to 1.3 μs [6].

On the other hand, the A-HFR attack requires less stringent synchronization, only needing approximate knowledge of the LiDAR's scan angles and timings. Its design tolerates synchronization discrepancies with the victim LiDAR up to tens of μs. Consider a rotating LiDAR that operates at 10 Hz; it scans about 0.2° horizontally in 50 μs. Since the A-HFR attack uses synchronization data primarily to limit the attack to a certain horizontal angle, synchronization errors within the range of tens of μs do not significantly impact the attack's effectiveness.

C. Attack Device Setup

The device setup for the Adaptive HFR attack consists of three parts: a laser emitter, a function generator, and a receiver. We follow the same basic setup as in Sato et al. [6]: we employ a function generator of Agilent 81160A [15], receiver consisting of S6775 PIN Photodiode [16] as a photodetector and LM6171 OpAmp [17] as the transimpedance amplifier.

D. Target LiDARs

In this WIP paper, we evaluate three mass-produced LiDARs with pulse fingerprinting, summarized in Table II. Our evaluation reveals distinct differences in the ranging intervals, specifically in the minimum (T_{min}) and maximum (T_{max}) pulse intervals, across each model. For AT128 [12], the manufacturer's product description aligns with our experimental findings, confirming its pulse fingerprinting feature. However, due to limitations in our setup, we could not accurately determine the AT128's specific fingerprinting parameters.

TABLE II: Our target mass-produced LiDARs, equipped with pulse fingerprinting. ASR: Attack Success Rate.

	Livox Mid-360 [10]	Hesai XT32 [11]	Hesai AT128 [12]
			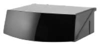
Scanning Type	Prism Rotating	Rotating	Mirror Rotating
T_{min}	1250 ns	250 ns	unspecified
T_{max}	1550 ns	450 ns	unspecified
T_α (estimated)	250 ns	42 ns	67 ns
HFR [6] ASR	1.0@4 MHz	0.20@9 MHz	0.03@9 MHz
A-HFR ASR	-	0.96@24 MHz	0.99@15 MHz

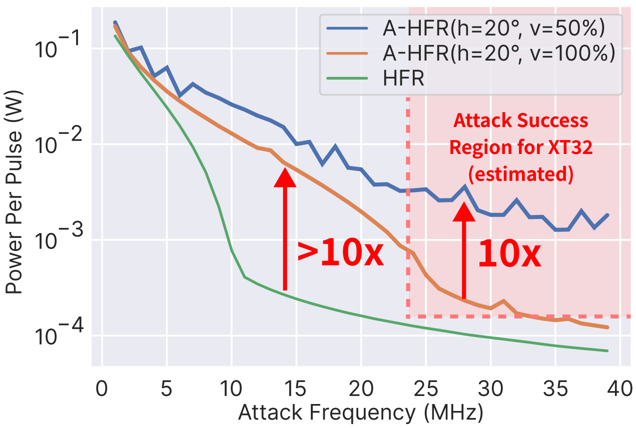

Fig. 4: Comparison of the laser pulse power at different attack frequencies. A-HFR attack reduces the attack angle to achieve higher attack frequencies with enhanced laser power. Furthermore, A-HFR can achieve higher attack capability by reducing the vertical attack angle as well.

IV. Evaluation

A. Effectiveness of Attack Angle Reduction

To assess the power enhancement in A-HFR attacks, we present a comparison of laser pulse power at different attack frequencies. This includes conventional HFR attacks and two variants of A-HFR attacks, as shown in Fig. 4. The first variant of the A-HFR restricts the attack within a 20° horizontal angle (v=100%), whereas the second variant further halves the vertical attack FoV (v=50%). In the case of low frequencies (<5 MHz), the effect of the attack angle reduction in A-HFR is minimal. However, at frequencies above 10 MHz, the pulse power of conventional HFR attacks diminishes notably. In contrast, A-HFR attacks between 10 and 20 MHz can substantially achieve higher power, with a tenfold increase. This power boost is further enhanced by reducing the vertical attack angle as well, particularly above 25 MHz, where power gains exceed

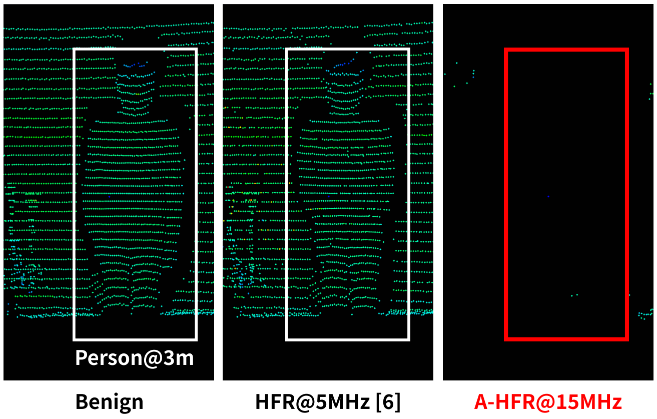

Fig. 5: A-HFR attack results on AT128. While the person point cloud is only blurred under the HFR attack (center), it is completely removed under the A-HFR attack (right).

ten times compared to scenarios without reducing. Therefore, Reducing attack angles in any direction is an effective strategy to attain high peak power at higher frequencies.

B. Evaluation Against Various Pulse Fingerprinting LiDARs

Fig. 5 compares the effectiveness of the conventional HFR attack and the A-HFR attack, targeting the Hesai AT128 equipped with pulse fingerprinting. We placed the spoofer 2 m away from the LiDAR, and the person stood 3 m away. A-HFR specifically targets a 20° horizontal angle. While the HFR attack somewhat blurs this shape, it remains identifiable. However, with the A-HFR attack, the person's point cloud is almost entirely removed. Notably, the A-HFR attack effectiveness was maintained over prolonged durations ($>$100 seconds). This demonstrates that attackers can significantly enhance their capability against LiDARs with pulse fingerprinting by employing the A-HFR attack.

Fig. 6 displays the point removal rates for both conventional and A-HFR attacks at various frequencies, tested against AT128 and XT32 LiDARs with pulse fingerprinting. For AT128, the attack was conducted over a 20° horizontal and full 25.4° vertical FoV. In comparison, the XT32 underwent evaluations in two scenarios: firstly, with a 20° horizontal and 16° vertical angle (v=50%), and secondly, with a 20° horizontal and 32° vertical angle (v=100%). The figure demonstrates that conventional HFR attacks have a maximum success rate of only 20%, limited by thermal issues. In contrast, our A-HFR attack effectively removes the target angle's point cloud by increasing the attack frequency. Its removal rate peaks at 15 MHz for AT128 and 24 MHz for XT32, eliminating over 90% of the point cloud in that angle. The variation in peak removal rates across different LiDARs can be due to differing pulse fingerprinting parameters (T_α), which influences the frequency needed to bypass the authentication.

Fig. 7 shows the efficacy of A-HFR's vertical attack angle reduction. Since the angle reduction boosts the peak power under high attack frequencies (Fig. 4), this expands the point cloud removal angle by about 20%.

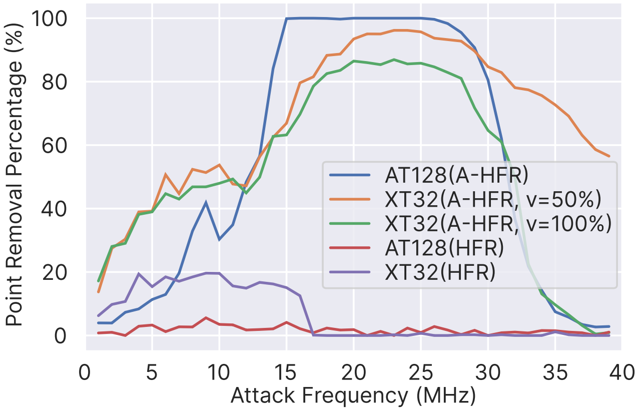

Fig. 6: Attack frequency versus point removal percentage under A-HFR attacks and conventional HFR attacks. We calculate the ratio of removed points within the angle of 20° horizontally and 16° vertically. For XT32, we experiment with two cases: one does not reduce the vertical attack angle, and another restricts the vertical angle to half.

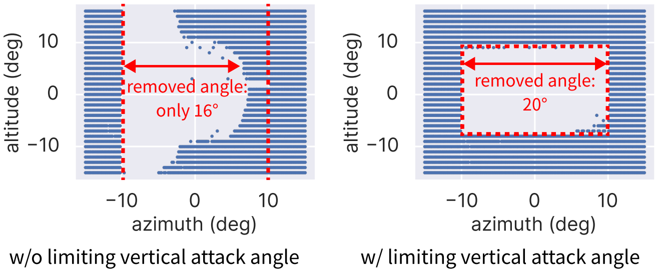

Fig. 7: 2D projection of point cloud when we conduct Adaptive HFR attack on XT32 with horizontal 20°. Blue points are the true point cloud remaining, and the red dotted line indicates the attack angle.

C. HFR Attack on Mid-360

Fig. 8 shows the efficacy of the conventional HFR attack on Mid-360, equipped with pulse fingerprinting. We found that even the conventional HFR attack can fully remove points across 20° horizontally and vertically. The Mid-360's vulnerability can be attributed to its relatively less secure authentication system compared to other LiDARs. This suggests a larger T_α as per Eq. 1 and potentially lower power in legitimate pulses, indicating that certain LiDAR models might have more lenient T_α settings and hence weaker security features.

D. Attack Capability in the Physical World

To verify A-HFR attack effectiveness in real-world conditions, we conducted an outdoor experiment aimed at removing a car from the point cloud data of XT32. The car was placed at a distance of 12 meters from the victim LiDAR, with its point cloud appearing within 10° horizontally and 8° vertically. The results, shown in Fig. 9, show the complete removal of the target car point cloud. In addition, we applied object detection to the point cloud with and without the attack, using Apollo

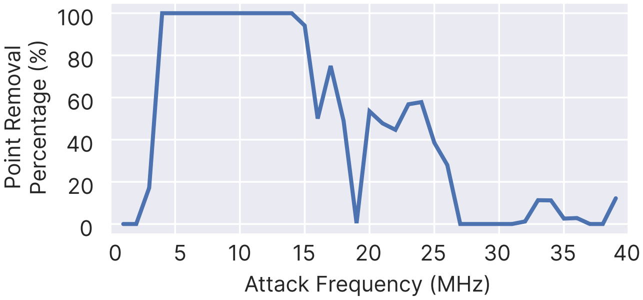

Fig. 8: Point removal percentage by HFR attack against Livox Mid-360, under the same condition as Fig. 6.

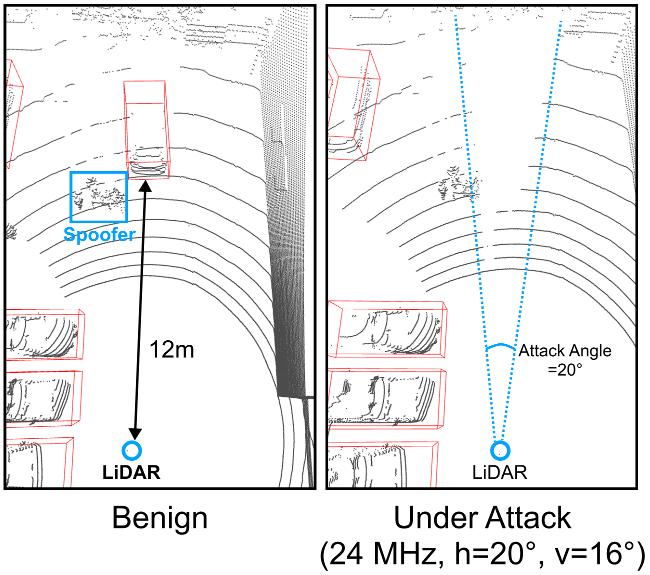

Fig. 9: A-HFR attack on Hesai XT32 to hide a real vehicle. We restrict the attack angle to 20° horizontally and 16° vertically. The red bounding box shows the detection results by Pointpillars in Apollo. When under attack, the vehicle becomes undetected with a 98% success rate over 15 seconds.

6.0 Perception [18]. As a result, while the car was detectable in the benign point cloud, but went undetected in the attacked data. This result shows that the A-HFR attack is a significant threat to real-world autonomous driving systems.

V. Discussion

A. Potential Defences against A-HFR Attack

When the pulse fingerprinting authentication relies solely on pulse interval pairs, attackers could always theoretically bypass it if higher enough frequency pulses were available. This implies that pulse fingerprinting authentication can be an effective mitigation but may not be an ultimate countermeasure. Thus, potential defenses can integrate additional pulse features, such as amplitude and width, into the authentication processes. The combination of these features can substantially increase the entropy of the fingerprinting and thus it can prevent the creation of a universal 'master key' by attackers. We plan to explore such potential countermeasures in future work since the implementation of them requires expensive modifications in the LiDAR signal processing or hardware.

B. Limitation of A-HFR Attack

1) Dependency on LiDAR Scan Parameters: A-HFR attack relies on pre-known scan patterns of the target LiDAR, but some commercial LiDARs have variable parameters which affect the spatial density of their scan. Misprediction of these parameters disrupts attack angle reduction by weak synchronization. For instance, if the rotation speed of the target rotating LiDAR is changed, the horizontal attack angle expected by the attacker will be distorted. However, this limitation is relatively minor, as external observation can allow attackers to adjust their predictions.

2) Synchronizing with Complex Scan Patterns: LiDAR systems with complex scan patterns pose challenges to the A-HFR attack. Such features hinder the attacker from accurately predicting scan timings and establishing effective attack angles. For instance, AT128 [12] has a very complicated vertical scan pattern and we could not design a vertical angle reduction for the A-HFR attack.

C. Ethical and Safety Consideration

We performed all experiments under controlled settings We fully paid attention to safety in our all experiments. We obtained official test permission on the private road we tested and kept other people out of the testing site. The experimental vehicle was always controlled by a human driver without autonomous driving features. During the experiments, we wore laser safety goggles for eye safety. and kept other people out of the testing site. We have performed a responsible vulnerability disclosure for related LiDAR manufacturers and AV companies to inform our findings before publication.

VI. Concluding Remarks and Future Plans

In this WIP paper, we identify a new LiDAR spoofing attack termed the A-HFR attack against LiDARs to assess the defensive efficacy of pulse fingerprinting implemented in new-gen LiDARs. A-HFR attack addresses the limitation of the diode overheating in the conventional HFR attack by only emitting lasers when the LiDAR scans the target object. We evaluate the A-HFR attack on three commercial LiDARs with pulse fingerprinting and find that the A-HFR attack is always effective against all three LiDARs and can successfully remove over 96% of the point cloud within a 20° horizontal and a 16° vertical angle, which can cover 95% of pedestrians and vehicles over 6 m away based on our experiment. This result poses a serious concern that current LiDARs equipped with pulse fingerprinting might not be sufficiently robust to the A-HFR attack. To address this, we plan to explore the possibility of designing more secure fingerprinting with other laser features than pulse interval.

Acknowledgements

This research was supported in part by the JST CREST JPMJCR23M4, JST PRESTO JPMJPR22PA, Amano Institute of Technology, NSF CNS-2145493, CNS-1929771, CNS-1932464, and USDOT UTC Grant 69A3552348327.

REFERENCES

[1] J. Petit, B. Stottelaar, M. Feiri, and F. Kargl, "Remote Attacks on Automated Vehicles Sensors: Experiments on Camera and Lidar," *Black Hat Europe*, vol. 11, p. 2015, 2015.

[2] H. Shin, D. Kim, Y. Kwon, and Y. Kim, "Illusion and Dazzle: Adversarial Optical Channel Exploits Against Lidars for Automotive Applications," in *International Conference on Cryptographic Hardware and Embedded Systems*. Springer, 2017, pp. 445–467.

[3] Y. Cao, C. Xiao, B. Cyr, Y. Zhou, W. Park, S. Rampazzi, Q. A. Chen, K. Fu, and Z. M. Mao, "Adversarial Sensor Attack on LiDAR-based Perception in Autonomous Driving," in *Proceedings of the 2019 ACM SIGSAC conference on computer and communications security*, 2019, pp. 2267–2281.

[4] Y. Cao, S. H. Bhupathiraju, P. Naghavi, T. Sugawara, Z. M. Mao, and S. Rampazzi, "You Can't See Me: Physical Removal Attacks on LiDAR-based Autonomous Vehicles Driving Frameworks," in *32nd USENIX Security Symposium (USENIX Security 23)*, 2023, pp. 2993–3010.

[5] Z. Jin, J. Xiaoyu, Y. Cheng, B. Yang, C. Yan, and W. Xu, "PLA-LiDAR: Physical Laser Attacks against LiDAR-based 3D Object Detection in Autonomous Vehicle," in *2023 IEEE Symposium on Security and Privacy (SP)*. IEEE Computer Society, 2022, pp. 710–727.

[6] T. Sato, Y. Hayakawa, R. Suzuki, Y. Shiiki, K. Yoshioka, and Q. A. Chen, "LiDAR Spoofing Meets the New-Gen: Capability Improvements, Broken Assumptions, and New Attack Strategies," in *Network and Distributed System Security Symposium (NDSS)*, 2024.

[7] "VLP-16 User Manual," https://velodynelidar.com/wp-content/uploads/2019/12/63-9243-Rev-E-VLP-16-User-Manual.pdf.

[8] K. Yoshioka, "A Tutorial and Review of Automobile Direct ToF LiDAR SoCs: Evolution of Next-Generation LiDARs," *IEICE Transactions on Electronics*, vol. advpub, p. 2021CTI0002, 2022.

[9] M. Yu, M. Shi, W. Hu, and L. Yi, "FPGA-based Dual-pulse Anti-interference Lidar System Using Digital Chaotic Pulse Position Modulation," *IEEE Photonics Technology Letters*, vol. 33, no. 15, pp. 757–760, 2021.

[10] "Livox Mid-360," https://www.livoxtech.com/mid-360.

[11] "XT32 — Mid-Range Mechanical Lidar — HESAI Technology," https://www.hesaitech.com/product/xt32/.

[12] "AT128 Auto-Grade Ultra-High Resolution Long Range Lidar — HESAI Technology," https://www.hesaitech.com/product/at128/.

[13] R. S. Hallyburton, Y. Liu, and M. Pajic, "Security Analysis of Camera-LiDAR Fusion Against Black-Box Attacks on Autonomous Vehicles," in *31st USENIX Security Symposium (USENIX Security 22)*, 2022.

[14] A. Geiger, P. Lenz, C. Stiller, and R. Urtasun, "Vision Meets Robotics: The KITTI Dataset," *The International Journal of Robotics Research*, vol. 32, no. 11, pp. 1231–1237, 2013.

[15] "81160A Pulse Function Arbitrary Noise Generator," https://www.keysight.com/us/en/product/81160A/81160a-pulse-function-arbitrary-noise-generator.html.

[16] "Si PIN Photodiode S6775," https://www.hamamatsu.com/us/en/product/optical-sensors/photodiodes/si-photodiodes/S6775.html.

[17] "LM6171 High-Speed, Low-Power, Low-Distortion Voltage Feedback Amplifier," https://www.ti.com/jp/lit/ds/symlink/lm6171.pdf.

[18] "Apollo Open Platform," https://apollo.baidu.com/.