Demo: Attacking Multi-Sensor Fusion based Localization in High-Level Autonomous Driving

Junjie Shen, Jun Yeon Won, Zeyuan Chen, Qi Alfred Chen
University of California, Irvine

I. Introduction

To enable high-level driving automation [1], the Autonomous Driving (AD) system in an Autonomous Vehicle (AV) needs to perform *centimeter-level localization* of its own global positions on the map [2]. Such localization function is highly security and safety critical in the AV context, since positioning errors can directly cause an AV to drive off road or onto a wrong way. For outdoor localization in general, GPS is the *de facto* location source, and thus a direct threat to it is GPS spoofing, a long-existing but still unsolved security problem with practicality proven on a wide range of end systems. Fortunately, AV systems today predominantly use Multi-Sensor Fusion (MSF) algorithms [3] that are generally believed to have the potential to practically defeat GPS spoofing [4]. However, no prior work has studied whether today's MSF algorithms are indeed sufficiently secure under GPS spoofing, especially in AV settings. In this work, we perform the first study to fill this critical gap. We consider the attack goal as using GPS spoofing to cause large *lateral* deviations in the MSF output, i.e., deviating to the left or right. This can cause the AV to drive off road or onto a wrong way.

To systematically understand the security property, we first analyze the upper-bound attack effectiveness and find that in the real-world trace, the majority of such upper-bound attack can only cause less than 50 cm deviation, which is far from causing the AV to drive off road. This shows that MSF can indeed generally enhance the security against GPS spoofing. Interestingly, we also observe that there still exist a few upper-bound attack results that can cause *exponential growths* of deviations. This allows the spoofed GPS to become the dominating input source in the fusion process and eventually cause the MSF to reject other input sources, which thus *fundamentally defeats the design principle of MSF*. In this work, we call it a *take-over effect*. We perform a cause analysis and find that such vulnerability only appears dynamically and non-deterministically. Leveraging this insight, we design *FusionRipper*, a novel and general attack that opportunistically captures and exploits take-over vulnerabilities with 2 stages: (1) *vulnerability profiling*, which measures when vulnerable periods appear, and (2) *aggressive spoofing*, which performs exponential spoofing to exploit the take-over opportunity.

II. Demonstration Plan

In the demo, we will show the end-to-end attack consequences of FusionRipper on Baidu Apollo [5], a production-grade Level-4 AD system, running in LGSVL [6], an industry-grade AV simulator. Specifically, the LGSVL accepts the control inputs from Baidu Apollo and produces real-time sensor inputs as the AV drives in the simulation environment. In the demo, we will run the simulation on the complete Baidu Apollo AD system with all functional modules enabled, i.e., localization, transform, perception, prediction, planning, routing, and control. We implement the attack logic as an independent component in Baidu Apollo, which intercepts GPS inputs from LGSVL and replaces them with the corresponding spoofed ones. In the demo, we will show two attack scenarios with one attacking to the left of the road and another to the right, where both have concrete safety consequences such as causing the AV to hit the road barrier or traffic sign.

Scenario 1: Attack to the left. In this scenario, the AV is driving on a road with a concrete barrier separating the opposite lanes. During the attack, the MSF localization results is taken over by the spoofed GPS inputs, causing the AV to deviate to the left and hit into the barrier.

Scenario 2: Attack to the right. In this scenario, the simulation environment has a stop sign installed on the roadside and the AV is driving on the lane next to the road curb. In the demo, we will show that the AV is deviated by the attack to drive over the road curb and finally crash into the stop sign.

In both scenarios, we will show the AV driving behaviors with and without attack to demonstrate the severity of the FusionRipper attack. To better show the take-over effect of FusionRipper, we will also visualize the real-time MSF localization outputs when the sensor inputs arrive. Attendees will be able to clearly see the two attack stages in FusionRipper as well as the immediate effect of the attack on the MSF outputs. The previews of the demo videos are available on our project website: **https://sites.google.com/view/cav-sec/fusionripper**.

Workshop on Automotive and Autonomous Vehicle Security (AutoSec) 2021
25 February 2021, Virtual
ISBN 1-891562-68-1
https://dx.doi.org/10.14722/autosec.2021.23010
www.ndss-symposium.org

Acknowledgements

This research was supported in part by the National Science Foundation under grants CNS-1850533, CNS-1929771, and USDOT UTC Grant 69A3552047138.

References

[1] SAE On-Road Automated Vehicle Standards Committee and others, "Taxonomy and Definitions for Terms Related to Driving Automation Systems for On-Road Motor Vehicles," *SAE International*, 2018.

[2] T. G. Reid, S. E. Houts, R. Cammarata, G. Mills, S. Agarwal, A. Vora, and G. Pandey, "Localization Requirements for Autonomous Vehicles," *arXiv preprint arXiv:1906.01061*, 2019.

[3] G. Wan, X. Yang, R. Cai, H. Li, Y. Zhou, H. Wang, and S. Song, "Robust and Precise Vehicle Localization based on Multi-Sensor Fusion in Diverse City Scenes," in *ICRA*, pp. 4670–4677, IEEE, 2018.

[4] J. Bhatti and T. E. Humphreys, "Hostile Control of Ships via False GPS Signals: Demonstration and Detection," *NAVIGATION: Journal of the Institute of Navigation*, 2017.

[5] "Baidu Apollo." https://github.com/ApolloAuto/apollo.

[6] "LGSVL Simulator." https://github.com/lgsvl/simulator.