SAVIOR: Securing Autonomous Vehicles with Robust Physical Invariants

Raul Quinonez[1], Jairo Giraldo[4], Luiz Salazar[2], Erick Bauman[1], Alvaro Cardenas[2], Zhiqiang Lin[3]

[1]The University of Texas at Dallas
[2]University of California Santa Cruz
[3]Ohio State University
[4]University of Utah

USENIX Security 2020

Autonomous Vehicles

- Autonomous Vehicles (AVs) include aerial, sea, and ground vehicles
- Levels of automation range from 0 to 5
- AVs evaluate their environment with a variety of sensors

[Gon17]

Current Problem

GPS Spoofing Mystery Affirms Need for Protection

The spoofing of GPS systems at the Geneva Motor Show gave us an unfortunate example of how vulnerable vehicles are to "spoofery."

Roi Mit | Apr 23, 2019

[Mit19]

Ars Technica

Researchers trick Tesla Autopilot into steering into oncoming traffic

Researchers trick Tesla Autopilot into steering into oncoming traffic. Stickers that are invisible to drivers and fool autopilot. Dan Goodin - 4/1/2019, ...

Apr 1, 2019

[Goo19]

Autonomous vehicles can be fooled to 'see' nonexistent obstacles

BY YULONG CAO, Z. MORLEY MAO | MAR 06, 2020

[YC20]

Mysterious GPS glitch telling ships they're parked at airport may be anti-drone measure

Elizabeth Weise, USATODAY Published 1:41 p.m. ET Sept. 26, 2017 | Updated 3:03 p.m. ET Oct. 3, 2017

[Wei17]

AVs Are Vulnerable to Sensor Targeted Attacks

Main Problem

- AVs rely on sensors to evaluate and interact with their environment
- Sensors are susceptible to GPS spoofing and transduction attacks that manipulate environmental physical signals

Previous Research Has Exposed Sensor Vulnerabilities

- Camera [DWJ+16, PSFK15, YXL16]
- LiDAR [PSFK15, SKKK17, CXC+19]
- RADAR [YXL16]
- Inertial Measurement Unit (IMU) [SSK+15, TWX+17, TLLH18]
- GPS [NKS+19, HLP+08, TPRC11, ZLS+18]

Transduction Attacks and GPS Attacks Cannot be Addressed with Classical Security

Transduction Attacks

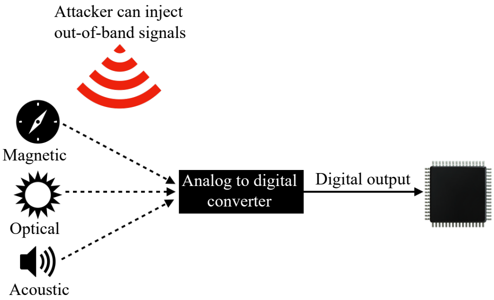

GPS Attacks

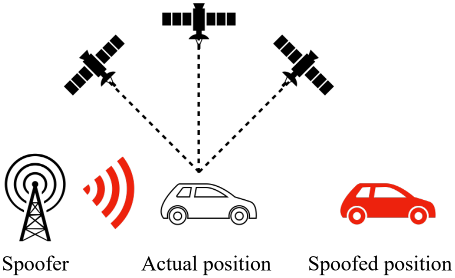

Insights of Our Work

SAVIOR

We introduce our SAVIOR (Securing Autonomous Vehicles wIth rObust physical invaRiants) framework contributing to the following:

1. We use well-known nonlinear dynamic models for aerial and ground AVs
2. We introduce a stronger stealthy attacker
3. We implement a Cumulative Sum (CUSUM) algorithm that improves detection performance over previous defenses that keep track of anomalies using time windows
4. The implementation is done in real vehicles including including an Intel drone, and our autonomous car

Sensors and Movement Variables

Drones

- 3 axes: roll, pitch, yaw
- Sensors: accelerometer, gyroscope, magnetometer, and GPS (lat, lon, alt)

Ground AV

- 2 axes: pitch, yaw
- Sensors: line data (angle, position) and speed

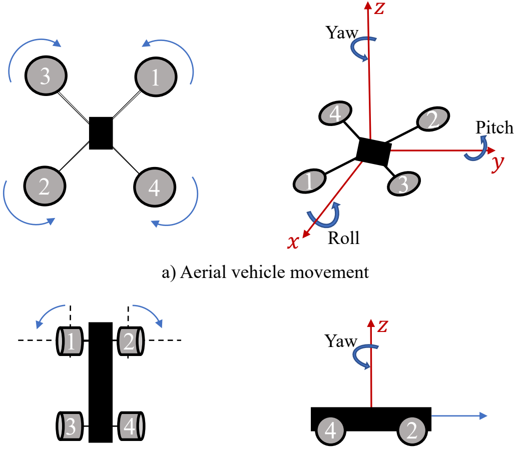

a) Aerial vehicle movement

b) Ground vehicle movement

Nonlinear Models

Dynamics of a Quadcopter [CFCH14, Luu11]

$$\dot{\phi} = \omega_\phi$$
$$\dot{\theta} = \omega_\theta$$
$$\dot{\psi} = \omega_\psi$$
$$\dot{\omega}_\phi = \frac{U_\phi}{I_x} + \dot{\theta}\dot{\psi}(\frac{I_y - I_z}{I_x})$$
$$\dot{\omega}_\theta = \frac{U_\theta}{I_y} + \dot{\phi}\dot{\psi}(\frac{I_z - I_x}{I_y})$$
$$\dot{\omega}_\psi = \frac{U_\psi}{I_z} + \dot{\phi}\dot{\theta}(\frac{I_x - I_y}{I_z})$$
$$\dot{x} = v_x$$
$$\dot{y} = v_y$$
$$\dot{z} = v_z$$
$$\dot{v}_x = \frac{U_t}{m}(\cos\phi \sin\theta \cos\psi + \sin\theta \sin\psi)$$
$$\dot{v}_y = \frac{U_t}{m}(\cos\phi \sin\theta \sin\psi - \sin\phi \cos\psi)$$
$$\dot{v}_z = \frac{U_t}{m}\cos\phi \cos\theta - g$$

Dynamics of a Car [KPSB15]

$$\beta = \tan^{-1}(\frac{l_r}{l_r + l_f}\tan(\delta))$$
$$\dot{x} = v\cos(\psi + \beta)$$
$$\dot{y} = v\sin(\psi + \beta)$$
$$\dot{\psi} = \frac{v}{l_r}\sin(\beta)$$
$$\dot{v} = a$$

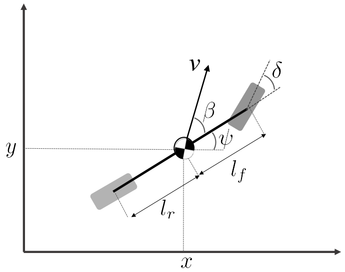

SAVIOR Design

- Online sensor pre-processing to convert raw data into usable form
- Offline pre-processing stage to learn physical invariants and a build model
- Online stage to predict measurements and compare observe values
- Anomaly detection will raise an alert if the anomaly is persistent

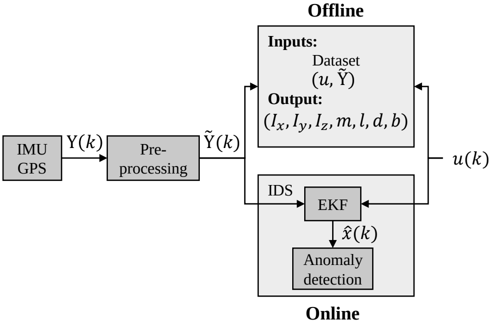

Online Stage

- An Extended Kalman Filter (EKF) [RG14] is used to predict AV's physical behavior by estimating unknown parameters from noisy sensor input
- The algorithm is divided into two main routines: prediction and correction
- The prediction will be compared against the observed data to be analyzed for sensor tampering

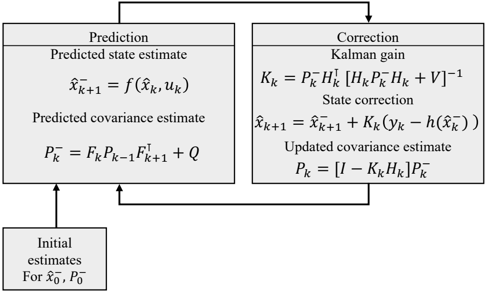

Anomaly Detection

- The residual associated with each sensor is calculated (1)
- A Cumulative Sum (CUSUM) algorithm is then used to detect persistent attacks (2)
- An alarm is raised if the residual difference is larger than a predefined threshold (3)

CUSUM Algorithm

1. $r_i(k) = \widetilde{Y}_i(k) - \hat{Y}_i(k)$
2. $S_i(k+1) = (S_i(k) + |r_i(k)| - b_i)^+$
3. $S_i(t_k) > \tau_i$

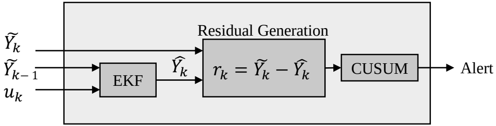

Implementation

- Controllers follow a publish-and-subscribe architecture to provide inter-process communication via topics
- We are interested in the following topics for aerial AVs: sensors_combined, vehicle_magnetometer, and vehicle_gps_position
- Anomaly detector is situated right before the control signals are being sent to the actuators
- The code runs in its own module in parallel with the controller

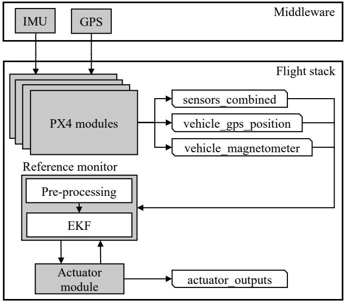

Evaluation

- Aerial AV: Intel Ready-To-Fly drone using PX4 flight controller (v1.9.2)
- Ground AV: Custom build on top of a Traxxas Ford Fiesta ST Rally chassis using ROS Kinetic Kame controller

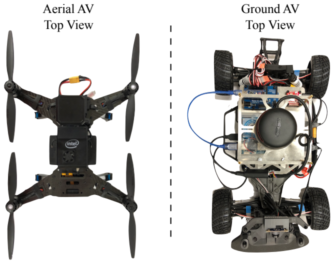

Ground AV Camera Attack and Detection Video

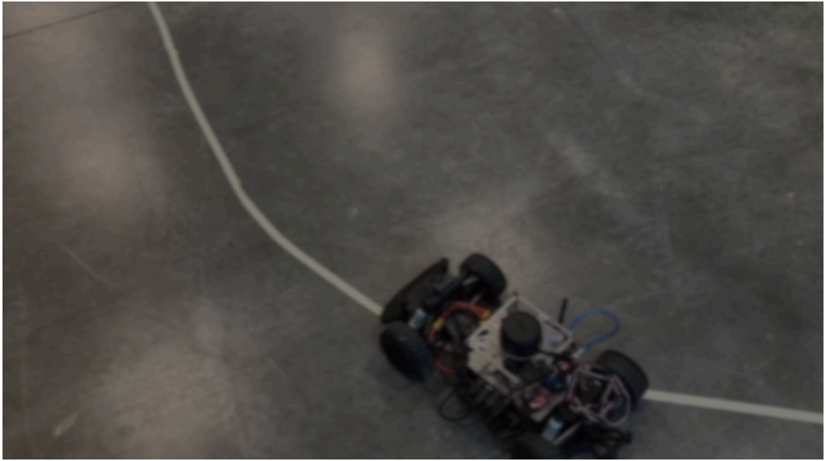

Videos available: https://www.youtube.com/watch?v=Ljrbtfo0gvM&list=PLmicm3IoL28eLU5v1FH3Z0FSn5Nl0uQLG

Aerial AV GPS Attack and Detection

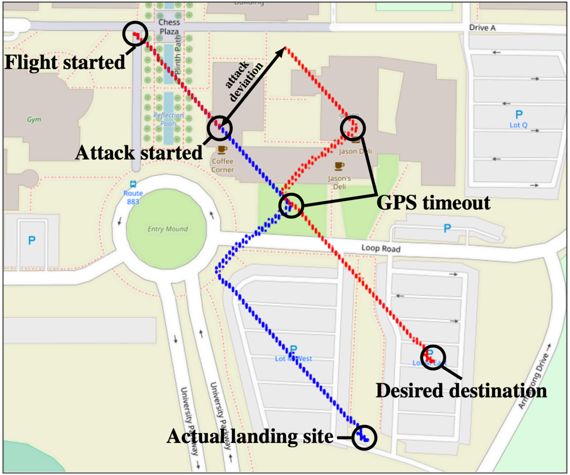

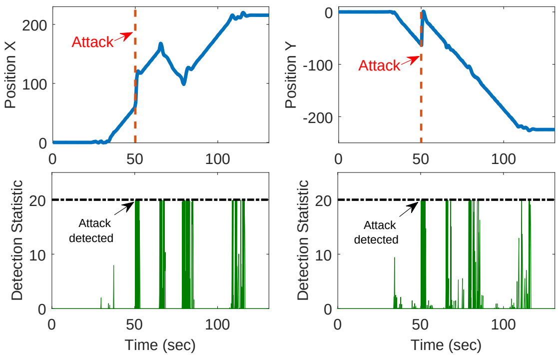

Comparison of SAVIOR with Baseline

- SAVIOR uses a nonlinear model for predicting the observations, and a CUSUM algorithm for anomaly detection (NLC)
- We will use Choi et al.'s [CLA+18] algorithm as a baseline since their anomaly detector was the current state-of-the-art
- Choi et al.'s [CLA+18] algorithm uses linear models for predicting observations and a Time-Window algorithm for anomaly detection (LTW)
- Our results show that our algorithms outperform state-of-the-art detection tools for AVs by detecting more attacks, detecting attacks faster, and having less false alarms

Linear (LTW) vs Nonlinear (NLC) Prediction Comparison

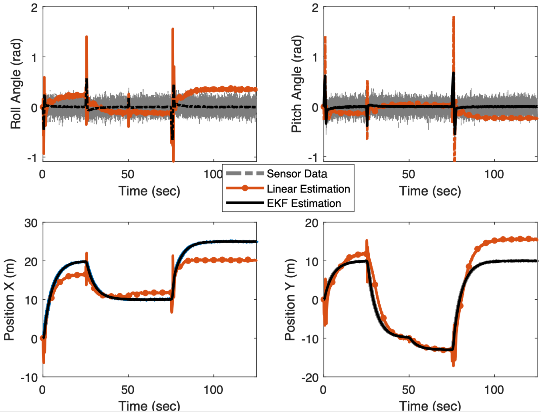

Window (LTW) vs CUSUM (NLC) Detection Time and ROC Curves

Drone

- NLC detects attacks faster
- NLC has a better ROC curve than LTW

Ground AV

- Detection is better for both, drones and ground vehicles

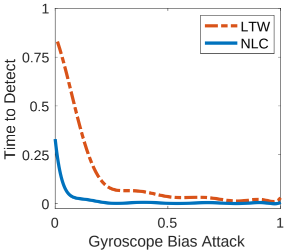

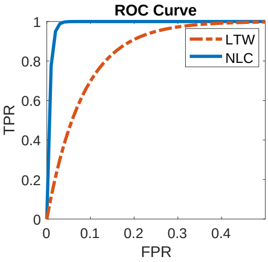

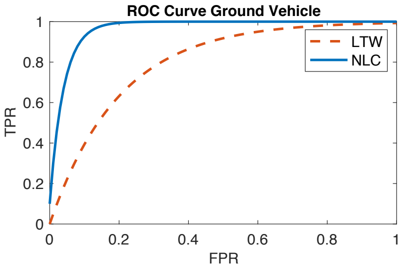

Stealthy Attacks

- We want to maximize the value of sensor tampering without raising any alarms
- The goal is to maximize deviation without increasing the added discrepancies
- This stealthy attack allows us to consider the worst case scenario of our PBAD system, where an attacker is not detected while it persistently injects the maximum amount of false information in the system

Purely Stealthy Attacks Against NLC Have Less Impact Than LTW

- NLC (blue) is able to follow the signal closer while the attacker performed an stealthy attack on the gyroscope and GPS
- LTW (orange) allows more tampering which ends up deviating the final destination more than NLC (blue)

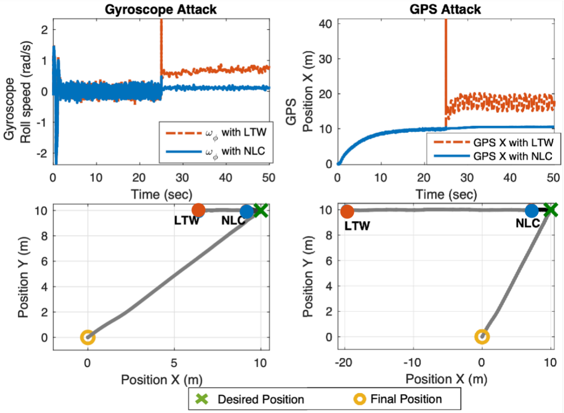

Performance Overhead

Drone

- On average, SAVIOR consumes 5.4332% of CPU resources on Intel Aero

Ground AV

- On average, SAVIOR consumes 2.2501% of CPU resources on Traxxas Ford Fiesta ST Rally

Module	Armed	Hovering	RC
Idle	30.1444%	29.4379%	30.6056%
mavlink_if1	16.0183%	15.6195%	15.8956%
EKF2	14.3242%	14.3779%	14.3006%
logger	6.8647%	7.1288%	6.8752%
mc_att_control	5.4349%	5.4007%	5.3425%
reference _monitor	**5.3572%**	**5.4332%**	**5.5093%**
tap_esc	4.4742%	4.4357%	4.4285%
sensors	4.2744%	4.4792%	4.5200%
hpwork	2.5077%	2.4462%	2.4750%
mavlink_if0	2.3323%	2.1384%	2.2667%
mc_pos_control	1.4911%	2.4727%	1.4693%
commander	1.4824%	1.4478%	1.4448%
gps	0.3662%	0.3323%	0.3077%

Module	Line Following	CA
lidar_collision_avoidance	12.6886%	13.0694%
elp_cam_bridge	11.0179%	15.6009%
process_line	10.3861%	11.7353%
image_processing	6.0726%	7.8523%
reference_monitor	**2.5192%**	**1.9809%**
arduino_node	2.4150%	2.5133%
line_follower	1.0097%	1.0488%
low_level_controller	0.7948%	0.4503%
perot_demo	0.6990%	0.6589%
roslaunch	0.4541%	0.2678%
rplidarNode	0.3074%	0.3020%
rosmaster	0.2973%	0.1569%
rosout	0.0658%	0.0250%

Conclusion

The Key Elements of Our Proposal

1. Use of well-known physical invariants
2. The use of offline system identification
3. The use of CUSUM algorithms
4. Evaluating the effectiveness of the anomaly detection tool with stealthy attacks that attempt to maximize the damage to the system

SAVIOR Source Code

https://github.com/Cyphysecurity/SAVIOR.git

Videos

https://www.youtube.com/watch?v=LjrbtfoOgvM&list=PLmicm3IoL28eLU5v1FH3ZOFSn5N1OuQLG

Thank You

Contact

Raul Quinonez	rxq100020@utdallas.edu
Jairo Giraldo	jairo.giraldo@utah.edu
Luiz Salazar	luedsala@ucsc.edu
Erick Bauman	exb131030@utdallas.edu
Alvaro Cardenas	alacarde@ucsc.edu
Zhiqiang Lin	zlin@cse.ohio-state.edu

References I

- Anežka Chovancová, Tomáš Fico, L'uboš Chovanec, and Peter Hubinsk, *Mathematical modelling and parameter identification of quadrotor (a survey)*, Procedia Engineering **96** (2014), 172–181.
- Hongjun Choi, Wen-Chuan Lee, Yousra Aafer, Fan Fei, Zhan Tu, Xiangyu Zhang, Dongyan Xu, and Xinyan Xinyan, *Detecting attacks against robotic vehicles: A control invariant approach*, Conference on Computer and Communications Security (CCS), ACM, 2018, pp. 801–816.
- Yulong Cao, Chaowei Xiao, Benjamin Cyr, Yimeng Zhou, Won Park, Sara Rampazzi, Qi Alfred Chen, Kevin Fu, and Zhuoqing Morley Mao, *Adversarial Sensor Attack on LiDAR-based Perception in Autonomous Driving*, Conference on Computer and Communications Security (CCS), 2019.
- Drew Davidson, Hao Wu, Robert Jellinek, Thomas Ristenpart, and Vikas Singh, *Controlling UAVs with sensor input spoofing attacks*, Workshop on Offensive Technologies (WOOT), USENIX Association, 2016, pp. 221–231.
- Gloria Gonzalez, *Autonomous vehicles, drones offer new insurer risks and opportunities*, `https://www.businessinsurance.com/article/20171207/NEWS06/912317799/Autonomous-vehicles,-drones-offer-new-insurer-risks-and-opportunities#`, 2017.
- Dan Goodin, *Researchers trick tesla autopilot into steering into oncoming traffic*, `https://arstechnica.com/information-technology/2019/04/researchers-trick-tesla-autopilot-into-steering-into-oncoming-traffic/`, 2019.

References II

- Todd E Humphreys, Brent M Ledvina, Mark L Psiaki, Brady W O'Hanlon, and Paul M Kintner, *Assessing the spoofing threat: Development of a portable gps civilian spoofer*, Radionavigation Laboratory Conference Proceedings, 2008.
- Jason Kong, Mark Pfeiffer, Georg Schildbach, and Francesco Borrelli, *Kinematic and dynamic vehicle models for autonomous driving control design*, Intelligent Vehicles Symposium (IV), IEEE, 2015, pp. 1094–1099.
- Teppo Luukkonen, *Modelling and control of quadcopter*, Independent research project in applied mathematics, Espoo **22** (2011).
- Roi Mit, *Gps spoofing mystery affirms need for protection*, `https://www.wardsauto.com/industry-voices/gps-spoofing-mystery-affirms-need-protection`, 2019.
- Juhwan Noh, Yujin Kwon, Yunmok Son, Hocheol Shin, Dohyun Kim, Jaeyeong Choi, and Yongdae Kim, *Tractor beam: Safe-hijacking of consumer drones with adaptive gps spoofing*, ACM Transactions on Privacy and Security (TOPS) **22** (2019), no. 2, 12.
- Jonathan Petit, Bas Stottelaar, Michael Feiri, and Frank Kargl, *Remote attacks on automated vehicles sensors: Experiments on camera and lidar*, Black Hat Europe 11 (2015).

References III

- Sławomir Romaniuk and Zdzisław Gosiewski, *Kalman filter realization for orientation and position estimation on dedicated processor*, acta mechanica et automatica **8** (2014), no. 2, 88–94.
- Hocheol Shin, Dohyun Kim, Yujin Kwon, and Yongdae Kim, *Illusion and dazzle: Adversarial optical channel exploits against lidars for automotive applications*, International Conference on Cryptographic Hardware and Embedded Systems (CHES), Springer, 2017, pp. 445–467.
- Yun Mok Son, Ho Cheol Shin, Dong Kwan Kim, Young Seok Park, Ju Hwan Noh, Ki Bum Choi, Jung Woo Choi, and Yong Dae Kim, *Rocking drones with intentional sound noise on gyroscopic sensors*, USENIX Security Symposium (USENIX Security), USENIX Association, 2015.
- Yazhou Tu, Zhiqiang Lin, Insup Lee, and Xiali Hei, *Injected and delivered: fabricating implicit control over actuation systems by spoofing inertial sensors*, USENIX Security Symposium (USENIX Security), USENIX Association, 2018, pp. 1545–1562.
- Nils Ole Tippenhauer, Christina Pöpper, Kasper Bonne Rasmussen, and Srdjan Capkun, *On the requirements for successful gps spoofing attacks*, Conference on Computer and Communications Security (CCS), ACM, 2011, pp. 75–86.

References IV

- Timothy Trippel, Ofir Weisse, Wenyuan Xu, Peter Honeyman, and Kevin Fu, *Walnut: Waging doubt on the integrity of mems accelerometers with acoustic injection attacks*, European Symposium on Security and Privacy (EuroS&P), IEEE, 2017, pp. 3–18.
- Elizabeth Weise, *Mysterious gps glitch telling ships they're parked at airport may be anti-drone measure*, https://www.usatoday.com/story/tech/news/2017/09/26/gps-spoofing-makes-ships-russian-waters-think-theyre-land/703476001/, 2017.
- Z. Morley Mao Yulong Cao, *Autonomous vehicles can be fooled to 'see' nonexistent obstacles*, https://gcn.com/articles/2020/03/06/lidar-spoofs-autonomous-vehicle-hack.aspx, 2020.
- Chen Yan, Wenyuan Xu, and Jianhao Liu, *Can you trust autonomous vehicles: Contactless attacks against sensors of self-driving vehicle*, DEF CON **24** (2016).
- Kexiong Curtis Zeng, Shinan Liu, Yuanchao Shu, Dong Wang, Haoyu Li, Yanzhi Dou, Gang Wang, and Yaling Yang, *All your GPS are belong to us: Towards stealthy manipulation of road navigation systems*, USENIX Security Symposium (USENIX Security), USENIX Association, 2018, pp. 1527–1544.