

August 31st 2018.
T-SET Final Report

Vehicle Trust Manager for Connected Vehicles

Nicola Bezzo (nicbezzo@seas.upenn.edu)
Prof. Insup Lee (lee@cis.upenn.edu)
PRECISE Center
University of Pennsylvania



DISCLAIMER

The contents of this report reflect the views of the authors, who are responsible for the facts and the accuracy of the information presented herein. This document is disseminated under the sponsorship of the U.S. Department of Transportation's University Transportation Centers Program, in the interest of information exchange. The U.S. Government assumes no liability for the contents or use thereof.

PROBLEM

In recent years we have been witnessing an increase in autonomous vehicles: most of the cars we drive nowadays use multiple sensors to maintain constant speed (e.g., adaptive cruise control), avoid obstacles and collisions, park, move autonomously through traffic, and improve the overall driving comfort. In the presence of a cyber attack in which the received information from the sensors is compromised, safety is also compromised. Thus, our goal is to develop a resilient framework to guarantee vehicular safety in the presence of malicious cyber attacks.

To consider cyber attacks that can compromise the overall safety of a vehicle and the surrounding vehicles, we are exploiting sensor fusion and redundancy in the sensor measurements. Specifically the problem under investigation is the following:

Given a vehicle with N sensors measuring directly or indirectly a certain state, find the set of policies such that the vehicle can achieve a desired state while one or more sensor measurements are maliciously compromised by an adversarial attack.

We focus primarily on control design schemes and address attacks on sensors for autonomous ground vehicles (Fig. 1). We build upon ways to introduce redundancy within the control loop, as well as methods for attack detection and identification. We also consider communication spoofing attacks and develop a technique to measure trust among vehicles.




Fig.1: The ground vehicle used to test the attack detection/mitigation scheme developed within this project. The figure displays the sensors setup of the platform.

METHODOLOGY

To solve the problem above, we build an adaptive recursive estimator (RAE), which uses a filter approach to estimate the state while reducing the malicious effects introduced by an attacker.

Our recursive algorithm is motivated by the results found in the Kalman Filter implementation with some modifications to accommodate the possible presence of an attack in one of the sensors. Together with the *prediction* and *update* phases found in the Kalman implementation we include a *shield* procedure. If an attack is present and such that one of the measurements is corrupted, the goal is to remove it or mitigate its effect. Since the attack vector is generally unknown, the strategy we implement changes the covariance matrix associated with the measurement error in order to increase the uncertainty where the measurement is different from the predicted state estimate.

Our formulation is hierarchical and use feedback to control the motion of the vehicle and achieve the desired state. Specifically, the application focus of this work is cruise-control for ground vehicles. Each sensor measures a specific environmental variable correlated with speed. The sensor measurements are passed to a security module, which is in charge of attack detection and state estimation and outputs an estimate of the velocity of the vehicle. The estimate is sent to a controller (in our application a P.I.D. loop) that returns the control inputs to drive the actuators to the desired state.

A vehicle may also broadcast incorrect information to the surrounding vehicles. To deal also with this problem we build a crowd-sourced trust management framework. Trust is computed by executing three tasks on each vehicle: (1) whether to accept the received incident report, based on its likelihood of being accurate; (2) if accepted, compute an endorsement opinion, which signifies the level of endorsement, on the incident report message; and (3) attach the computed endorsement to the incident report and forward it down-stream. With the proposed approach, these two decisions are made based the trust score of the report originator and forwarders. The trust score is computed by a central authority by aggregating vehicle behavior history w.r.t. incident report accuracy

RESULTS

We illustrated the development framework on a design of secure cruise control of a fully electric unmanned ground vehicles (UGV) built at the University of Pennsylvania (UPenn), shown in Fig. 1. The robot was programmed in C++ on the ROS (robot operating system) environment.

Several data were collected to extract the dynamical model of the vehicle. These

tests were conducted on different type of surfaces both indoor and outdoor and, as a result, a seventh order model was extracted that capture both the electromechanical and kinematical constraints of the vehicle.

Extensive simulations run in Matlab/Simulink (Fig. 2(a)) has shown that the vehicle can reach and maintain the desired state even when one of the sensors is compromised by a malicious attack. Specifically, we showed that if less than $N/2$ sensors are under attack, we can estimate the correct state of the system and maintain the desired cruise speed.

During the hardware implementations (Fig. 2(b)) we decided to use GPS and the left and right encoders to obtain three independent speed measurements. The GPS measures time-stamped global position, thus with this information and specific transformations we can derive the speed. Similarly from the encoder we can obtain the number of counts which translate into rotational velocity and finally into linear velocity.




Fig.2: (a) Results from simulations in ROS and Matlab in which the GPS is spoofed while other two sensors are not compromised (bottom-left subfigure). (b) Hardware implementation on the UGV. The subfigure shows the GUI used during the experiments to visualize data and to launch attacks on the sensors.

Once an attack is injected in one of the sensors, a weight is added to the noise variance of the corrupted measurement decreasing its trustworthiness. During the experiments each of the sensors was attacked while the robot was moving to maintain the set cruise speed. The RAE algorithm was always able to detect and remove the sensor under attack and guarantee the desired performance.

CONCLUSION

The developed adaptive technique compares the estimated state with each of the sensor measurements and returns a higher variance of the measurement noise if a

sensor is under attack.

Within this technique we can consider noisy measurements to estimate the correct state of the system and detect attacks that act outside the noise profile of the sensors. The only limitation is that the adaptive recursive algorithm needs an accurate selection of the noise profile and weights in order to converge to the correct state. A too small bound on the error noise implies that the estimator may reject most of the measurements while a too large bound on the error can lead more attacks going through the system because they are within the error noise profile. In real applications these boundaries on the noise profiles are usually given or can be calculated through hardware testing.

Future work will focus on extending the proposed technique on multi-vehicle systems incorporating V2V (vehicle-to-vehicle) and V2I (vehicle-to-infrastructure) protocols. We are also targeting more complex hardware implementation such as adaptive cruise control, waypoint navigation, and complex attack vectors on both sensors and actuators as well as experimenting with passenger automobiles.