Autonomous Vehicles and New Cyber Physical Challenges

White Paper submitted to the 2014 NSF National Workshop on Transportation Cyber-Physical Systems

Mario Gerla, Peter Reiher, Mevlut Turker Garip, Chuchu Wu
Computer Science Dept, UCLA

Introduction

The urban fleet of vehicles is rapidly evolving from a collection of sensor platforms that provide information to drivers and upload filtered sensor data (e.g., GPS location, road conditions, etc.) to the Internet Cloud; to a network of autonomous vehicles that exchange their sensor inputs among each other in order to optimize a well defined goal. This goal, in the case of autonomous cars, is the prompt delivery of the passengers to destination with maximum safety and comfort and minimum impact on the environment.

When the human control is removed, the autonomous vehicles must efficiently cooperate to maintain smooth traffic flow in roads and highways. Visionaries predict that the vehicles will behave much better than drivers allowing more traffic on highways with lower delays, less pollution and certainly better driver and passenger comfort. However, the complexity of the distributed control of hundreds of thousands of cars cannot be taken lightly. If a natural catastrophe suddenly happens, say an earthquake, the vehicles must be able to coordinate the evacuation of critical areas in a rapid and orderly manner. This requires the ability to efficiently communicate with each other and also to discover where the needed resources are (e.g., ambulances, police vehicles, information about escape routes, images about damage that must be avoided, etc.). Moreover, the communications must be secure, to prevent malicious attacks that in the case of autonomous vehicles could be literally deadly since there is no standby control and split second chance of intervention by the driver (who may be surfing the web).

This efficient communications and distributed processing environment can be provided by a new network and compute paradigm specifically designed for vehicles - the Vehicular Cloud. The Vehicular Cloud is the place where the Physical World and the Cyber World come together. The Cloud receives inputs from the various on board sensors and radios, computes efficient control strategies and “actuates” these strategies in the vehicles. At the first cyber layer, the Cloud offers several basic services, from routing to content search, spectrum sharing, dissemination, attack protection, etc. These basic services are leveraged by the next cyber layer where the autonomous vehicle applications reside and are shared via standard, open interfaces by all auto manufacturers. Below we outline a few representative vehicular applications and services, and highlight the new challenges posed by autonomous vehicles.

Beacons and Alarms

One important application built within the vehicular cloud is ``Beaconing and Alarms". Recall that the Autonomous Vehicle (AV) sensors (from optical to Lidar) do most of the work in the attempt to keep the vehicle and its passengers out of trouble. Sensors alone, however, are not sufficient to maintain stable operations in high speeds and extremely reduced inter-vehicle spacing. This is particularly true in truck platoons. In this case, it was found that communications from front to rear trucks in the platoon are essential to avoid the onset of oscillations. Likewise, V2V (Vehicle-to-Vehicle) communications are necessary to avoid the formation of shock waves in a long column of AVs when an accident occurs. Intersection collisions will not be so critical when most of the cars are autonomous, since the AVs (unlike human drivers) abide by the signals and speed limits and approach intersections with caution. However, V2V communications will still be required among lead cars facing 4-stop intersections in order to implement the ``smart traffic light". The electronic light schedules mini-platoons of cars across the intersection, like a real traffic light would do, dramatically reducing delays. AVs will also learn about road conditions ahead via V2V in order to make the drive more comfortable for the passengers.

Intelligent Transport

The introduction of autonomous driving will greatly enhance Intelligent Transport. The AVs will be able to use the existing highway network much more efficiently than manually-operated cars because they can be packed in compact platoons and convoys. They can also make efficient use of preferred (or pay-per-service) lanes, by maintaining a ``train on wheels" configuration on such lanes, and by allowing efficient in-and-out lane transitions using a combination of sensors and V2V communications, in a much safer way than human can (especially at sustained speeds). The AVs can also manage automatic charges. They can participate in auctions and bids on behalf of customers if necessary, and can enforce the fees by detecting and reporting non-complying vehicles. As for safety, the AVs can become aware of other mobiles sharing the road, say pedestrians and bicycles. They can track them with their sophisticated sensors/Lidars and can share the information of ``bike ahead" with vehicles behind and one of two lanes across through V2V communications.

Infrastructure Failure Recovery

The AVs depend on the infrastructure (e.g., WIFI access points, DSRC RSUs, and LTE) for several non-safety functions such as advanced sensor data processing and intelligent transport. In the case of a major infrastructure failure caused by an earthquake, say, some of the AV functions that strictly depend on the infrastructure must be taken over by human drivers. However, there is a gray interval between infrastructure failure and human take over. During this interval, AVs must fight the problems on their own. This is a very critical window because the AVs only know about their immediate neighbors. After the disaster, they have lost knowledge of the neighbors beyond “visual” reach which was provided by the Internet ITS server. To avoid a second disaster, caused by the AVs going out of control, it is important to maintain a V2V-supported propagation of traffic conditions and congestion state on adjacent roads. This background ``crowdsourcing" of traffic will allow the AVs to make intelligent routing decisions (to avoid obstacles or blocked roads in case of earthquakes) so that the human drivers can progressively take over with confidence.

Protection from Attacks

Besides the common security requirements like privacy, confidentiality, Distributed Denial of Service (DDoS) protection and authentication, the AV is very vulnerable to vicious attacks that may, say, disable the steering or the brakes system. The latter attacks are of concern with normal cars when a human driver is in control. They are extremely more dangerous for AVs because the driver cannot react instantly. For this reason, the protection from attacks - both external (from access points or from conventional vehicles) as well as internal (from other AVs) - must be designed with stricter standards. Yet, access to the cars' internal mechanism and possibly to On-Board Diagnostics (OBD) and CAN bus must be allowed when the AV is out of control, because of either internal malfunctioning or a malicious attack.

The more fully integrated and the more dependent on the network the AVs are, the greater the risk of successful cyber attacks upon them. Such attacks might be no different in character than the typical attacks seen on other networked systems, but we should also expect attackers to target AVs for their unique characteristics. An AV has significant control and impact on its physical environment, affecting everything from the congestion of traffic in its vicinity to the safety of pedestrians and other vehicles. Imaginative criminals are sure to find alarming uses for compromised AVs.

Great efforts should be made to secure the network and computing systems of these autonomous vehicles, of course, but all experience with network-connected systems suggests that even the most determined effort to protect a resource on the network sometimes fails. Thus, we can expect to see successful attacks against AVs' computers even if we work hard to prevent them. Thus, while no reasonable effort should be spared to prevent cyber attacks on such vehicles, we must also be prepared to deal with compromised autonomous vehicular systems.

Important elements in protecting these systems will include:

- Strong compartmentalization of the various computing and control elements in the vehicle, so that compromises of one subsystem will not necessarily cause complete loss of all systems
- Sanity checking of behavior of both individual vehicles and the overall system, to ensure that misbehavior is detected
- Registration of all vehicles and suitably strong cryptographic authentication of vehicles
- Safe defaults to allow acceptable AV behavior in the face of both general and targeted denial of service attacks on the communications system

Proper development of defenses against compromise and misuse of AVs will require better understanding of their vulnerabilities and what benefits attackers can gain from controlling such vehicles. Such understanding will allow us to focus our defensive efforts on the easiest attacks that offer the most benefit to the attackers.