The 45th Annual IEEE/IFIP International Conference on Dependable Systems and Networking
1st International Workshop on Safety and Security of Intelligent Vehicles (SSIV)

Improving Intelligent Vehicle Dependability By Means of Infrastructure-Induced Tests

Wilfried Steiner
TTTech Computertechnik AG
Vienna, Austria
wilfried.steiner@tttech.com

Ayhan Mehmed
TTTech Computertechnik AG
Vienna, Austria
ayhan.mehmed@tttech.com

Sasikumar Punnekkat
Mälardalen University
Vasteras, Sweden
sasikumar.punnekkat@mdh.se

***Abstract*—Advanced driver assistance systems (ADAS) take over more and more driving responsibilities from the human operator and, therefore, evolve into safety-critical systems. Thus, the dependability of such systems is of up-most importance. While upcoming automobiles themselves will implement fault-tolerance and robustness mechanisms, it can be beneficial to also take infrastructure measures into account when assessing the overall vehicle dependability. In this paper we discuss an example of an infrastructure measure that targets to improve the dependability of an on-board computer vision system. Based on this example we outline a cyber-physical systems (CPS) architecture for intelligent vehicles and address open research directions.**

I. Introduction

Modern advanced driver assistance systems (ADAS) make driving safer, more comfortable, and economically more efficient. Perceiving the environment with a variety of sensors, such as long and short range radars (LRR, SRR), light detection and ranging sensors (LiDAR), and video cameras (vision systems), ADAS support the driving function at different levels of automation. At the lowest level, the intelligent vehicle provides information to the driver, such as sound and visual alerts and warnings, while leaving full control of the vehicle to the driver. At an intermediate level of automation (semi-autonomous) the intelligent vehicle provides active support to the driver, such as the lane keeping-assistant, where operating the steering wheel starts to require more than the usual needed force, when the vehicle is in danger of unintentionally leaving the lane. Thus, the driver still remains in the control loop of the vehicle. At the highest level of automation (piloted and autonomous driving) the driving tasks are carried out by the vehicle itself. However, in certain cases the driver is able to overrule the system. An example for an overrulable system is Adaptive Cruise Control (ACC). Even though ACC automatically controls the speed and the distance to the front vehicle, the driver is still able to switch the system off or to set up the vehicle speed. On the other hand, in critical situations, when emergency breaking is needed, the system has to react directly without waiting for driver's activation or confirmation, thus the driver is excluded from the control loop [1].

Several OEMs have announced, that the future of intelligent vehicles will aim at the realization of highly-intelligent and fully-autonomous vehicles. However, with higher intelligence comes higher complexity and the inherent threat that this complexity may run counter to the original goal of enhanced vehicle control and vehicle safety. Thus, well-defined architectures for these cyber-physical systems (CPS) are mandatory. In order to reach a sufficient level of dependability such architectures will implement traditional dependability methods and means. For example, hardware and software redundancy to allow self-diagnostics and fault-tolerance capabilities, such as fault detection and masking, in order to ensure the correct functioning of the vehicle.

In addition to a well-defined CPS architecture, there is a variety of novel methods available to ensure the correct vehicle function and to improve the vehicle's performance. Virtual sensors, for example, such as digital maps and wireless communication enhance the environmental perception of intelligent vehicles. Thus, the data of traffic sign location, as present in digital maps, together with information of the vehicle position, as well as vehicle direction, and vehicle speed are fused together in order to improve the success rate of traffic sign recognition by the vehicle [2], [3]. On the other hand Vehicle-to-Vehicle (V2V) or Vehicle-to-Infrastructure (V2I) communication is used to enhance the environmental perception in situations when the on-board sensors exceed in range or have poor measurements due to bad environmental conditions (e.g., weather conditions, direct sunlight).

Complementary to these research branches, we introduce the concept of infrastructure-induced tests, in this paper, that aims to leverage said novel methods for the sake of dependability measurement and enhancement of the intelligent vehicle while in operation. We will discuss this concept by the example of a computer-vision monitor in the following section. In Section III we give an overview of a CPS architecture for ADAS in general and in particular for infrastructure-induced tests. Based on the computer-vision monitor example and the proposed CPS architecture we will discuss possibilities of test scenarios in Section IV. In Section V we propose various research directions in the field of infrastructure-induced tests. Finally, we summarize and conclude in Section VI.

II. The Computer-Vision Monitor Example

The computer-vision monitor (CVM) is an example of a potential infrastructure measure to continually control the correct behavior of an on-board computer-vision system (CVS). We consider the CVS to contribute to the maneuvering of the vehicle in some aspect. For example the CVS may implement a pedestrian recognition function, which may automatically initiate a braking procedure, in case of emergency. Another example is a traffic sign recognition function that detects stop signs or traffic lights and brings the vehicle to safe halt. While there certainly will be internal self-checks and monitors in place, the aim of the CVM is to use vehicle-external stimuli to conclude upon the correctness of the CVS.

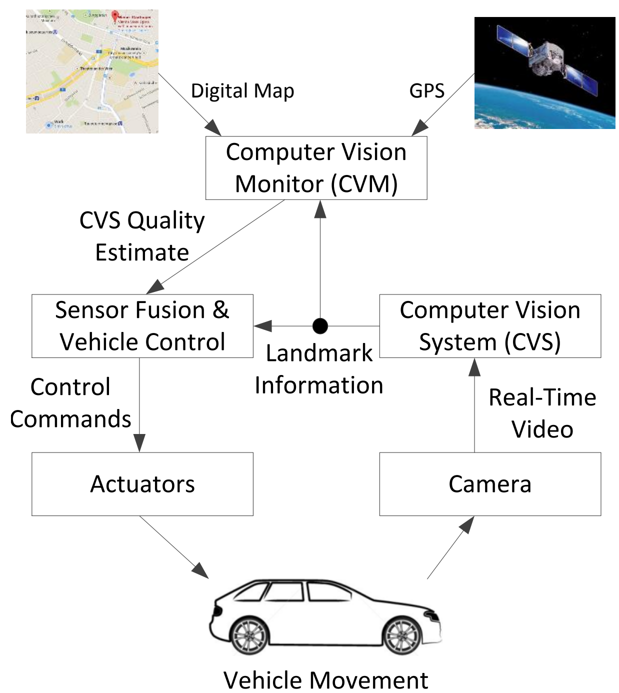

Fig. 1. Computer vision monitor - an example for infrastructure support

A. Principle of CVM Operation

An overview of the CVM concept is depicted in Figure 1. We consider an automobile in operation that perceives its environment by means of various sensors, for example video cameras (as depicted). This one or many cameras will record the vehicle's surrounding and produce real-time video data streams as an input to further processing in the CVS. Out of this video data, the CVS extracts information concerning which landmarks are present in the vehicle's surrounding. This information is potentially fused with other sensor data and used as input to the vehicle control unit, which in turn will generate control commands for the vehicles actuators. Finally, the control commands are executed and the vehicle actuators cause a certain vehicle movement. Thus, a closed control loop is formed, in which the CVS plays an important role.

Now, the concept of the CVM is to listen to the output of the CVS and to check whether this output is reasonable. For this, the CVM needs an independent source of information regarding the landmark's characteristics and position. One option to gather such information is the combination of precise digital maps and a localization procedure like GPS. The CVM can, thus, locate landmarks in the vehicle's upcoming surrounding independently of the CVS. It will then compare, whether the CVS is able to detect the landmark, or not. Depending on the characteristics of the landmark, the non-detection of the CVS can be evidence of a CVS failure. We discuss some of these landmark characteristics next.

For the purpose of infrastructure-induced tests we distinguish four types of landmarks:

- Natural landmarks are landmarks present in the environment without being installed by humans. Examples are: rivers, hills, mountains.
- Architectural infrastructure landmarks are landmarks placed in the environment by humans to achieve a purpose other than navigation support. Examples are: buildings, bridges, tunnels.
- Human-guiding landmarks are landmarks placed in the environment by humans for the sake of guiding the human operator of a car. Examples are: traffic signs, road markings.
- Vehicle-guiding landmarks are landmarks placed in the environment by humans for the sake of guiding the vehicle directly.

Naturally, different types of landmarks differ with respect to their occurrence, their implementation, and maintenance costs.

- Occurrence: the availability of the landmarks is defined according to whether they appear in urban or highway roads. Natural landmarks are mostly available near highways, whereas architectural landmarks are common for urban areas. Human and vehicle guiding landmarks appear in both, urban and highway roads. This distinction is important, since autonomous vehicles perform different functions in urban and highway roads.
- Implementation cost: while natural and architectural landmarks do not have implementation costs, human-guiding landmarks need investment of resource. However, they are already implemented in the roads for the purposes of the standard vehicles. Vehicle-guiding landmarks, on the other hand, require the highest investment, because their implementation in the roads is still in an initial stage. Nevertheless, it is likely that autonomous driving in the future will require dedicated vehicle-guiding landmarks as part of the road infrastructure to some degree. One of the reasons being the freedom and flexibility of customization for the needs of autonomous driving, e.g., their availability, frequency of occurrence, size, and so forth.
- Maintainability costs: human and vehicle-guiding landmarks are the most reasonable to maintain from the perspective of an infrastructure owner. However, it should be noted that today the process of identifying the maintenance need of human-guiding landmarks and vehicle-guiding landmarks (so they exist) is still mostly done manually.

B. Landmark Maintenance Center

While the CVM is regularly performing infrastructure-induced tests, using the landmarks, in some cases, the CVM reporting a failure could actually imply that the landmark itself is faulty (damaged, out of sight, etc.) and not the CVS.

A vehicle-local approach to mitigate this problem is to define a threshold of landmark misses that needs to be reached before the CVS is classified as being faulty. A more advanced mitigation strategy is the installation of a remote landmark maintenance center, to which CVS landmark misses (as detected by the CVM) are reported by the vehicle. An example scenario of a faulty "Landmark 2" is depicted in Figure 2.

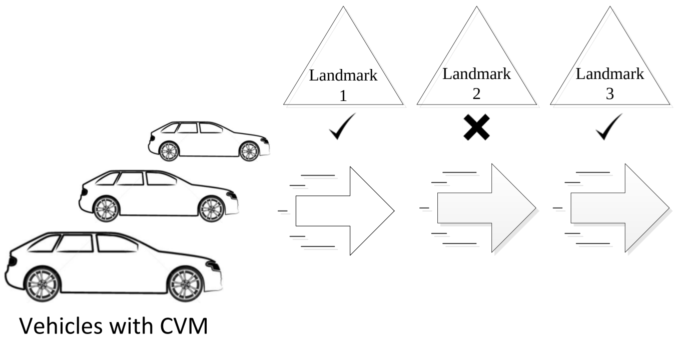

Fig. 2. An example scenario of misbehavior in CVS.

In this example, the CVM verifies the operation of the CVS at Landmark 1 and Landmark 3 but reports a CVS miss at Landmark 2. As more and more vehicles report CVS-misses of the same landmark (Landmark 2), the landmark maintenance center can correlate this information from a multitude of vehicles, classify the landmark to be faulty, and report the failure of the landmark back to vehicles in the area of the problematic landmark.

An overview of the core functions of the landmark maintenance center is depicted in Figure 3:

1) collect the CVS misbehaviors as reported by one or many computer vision monitors of one or many vehicles,
2) identify the problematic landmark, based on the collected information e.g., a landmark for which several vehicles report a CVS misbehavior can be identified to be damaged,
3) inform the CVM and vehicle control system, that the identified landmark may be damaged, and
4) trigger a maintenance activity, such as sending a repair crew to the damaged landmark's site.

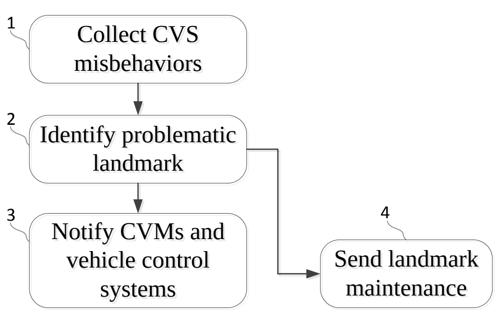

Fig. 3. An example operation of landmark maintenance center.

C. Landmark Gathering

Digital maps together with localization systems enhance the driving experience by helping the driver to reach the desired destination. Some of the data that can be found in digital maps is information about exact position of traffic infrastructure components, such as traffic signs, junctions, traffic lights. More specialized digital maps, those used for functional enhancement of ADAS systems, contain information for the road markings, such as stop lines, zebra zones and road center lines. All this mapped landmarks (traffic signs, road markings, traffic lights, junctions, etc.) on the digital map can be used by the proposed CVM as independent source of information in order to perform infrastructure-induced test.

Furthermore, human-guiding landmarks, natural and architectural infrastructure landmarks can be used for the purposes of the infrastructure-induced tests. This can be achieved by use of digital street maps. Digital street maps (e.g. Google Street View) provide a real view of panoramic images of streets and have become very popular and convenient for users to observe real-street environment. The panoramic images contain detailed information about the vehicle surrounding, which in turn can be used to extract landmarks such as rivers, hills, mountains, buildings, bridges, tunnels for the purposes of the proposed CVM.

An extreme case will be when there is a need for performing the infrastructure-induced tests within intervals of minutes or seconds (see Section IV). For that, dedicated vehicle-testing landmarks will be necessary.

III. CPS Architecture Supporting Infrastructure-Induced Tests

With the enormous pace of innovation in the automotive E/E area, functionality and the number of ECUs (Electronic Control Unit) is growing rapidly. Recently, a German OEM announced one of their top-of-the-line vehicles to comprise a hundred-and-thirty ECUs in total. This growing number of ECUs drives complexity, weight, power, and space consumption and, thus, ultimately cost. Infrastructure-induced tests as one of such automotive innovations will certainly also require non-negligible computation and communication resources. Therefore, for efficiency reasons novel architectural approaches are necessary. For this, a higher degree of integration of functions per ECU is necessary, to reduce the overall number of ECUs.

Integration has already been done today, for example using an AUTOSAR environment, but another degree of integration is necessary to compensate for emerging and future functionality in the car to reduce or, at minimum, to preserve the number of ECUs. Virtualization is one such solution for integration. Virtualization allows a single ECU to run multiple virtual machines in parallel. For example, it is possible to run AUTOSAR (and even multiple instances thereof) in parallel with rich operating systems like Linux. The hypervisor as the core element of virtualization guarantees bounded interference or even interference-freedom from one virtual machine on the respective others. Provided the right architecture being in place, it can be argued that safety-related applications can even be collocated with other, non-critical or lesser critical applications on the same ECU. Upcoming embedded multicore processors largely implement the hypervisor in hardware and only require a thin software layer for its management.

Furthermore, new E/E functions will be, to a large amount, distributed applications and it is only marginal attractive to install dedicated networks for each of these functions. It is much more cost-effective to assume a converged network solution to be in place. Luckily, automotive Ethernet that builds on the IEEE 802.1 and 802.3 standards is available for automotive use. Automotive Ethernet solutions implement AVB (Audio/Video Bridging) today and TSN (time-sensitive networking) in the near future to provide various levels of Quality-of-Service (QoS). The available Deterministic Ethernet further extends AVB and TSN to even provide Guarantee-of-Service (GoS) an hypervisor equivalent for the network.

The hypervisor technology on the ECU level together with the GoS technology on the network level establishes the foundation of vehicle-wide virtualization.

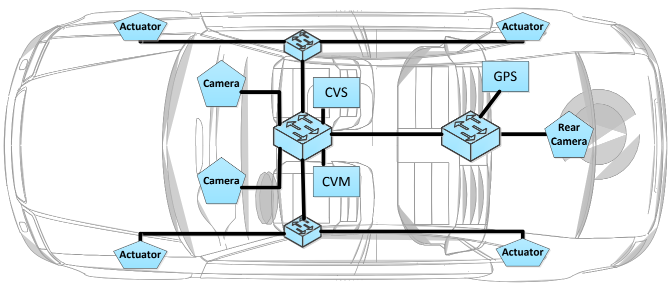

Fig. 4. Idealized setup of an in-vehicle computer-vision system and computer-vision monitor (redundancy levels are not depicted)

Figure 4 depicts an idealized setup of an in-vehicle CVS and CVM system. Here, two front cameras are installed and continuously communicate real-time video streams to a front Ethernet switch via a direct connection. The front Ethernet switch forwards the video data to the directly connected ECUs that host the CVS and the CVM functionality. Furthermore, the front Ethernet switch is connected to a back Ethernet switch towards which a rear camera is connected to. Also, the GPS receiver is connected to the back Ethernet switch. The front Ethernet switch also connects to two secondary switches left and right, that are used to deliver the control signals to the actuators.

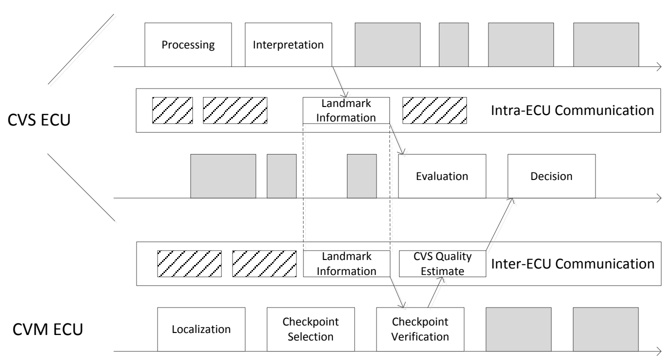

Fig. 5. Overview of CVS/CVM tasks, task to CPU/ECU assignment, and intra-ECU inter-ECU communication

Figure 5 depicts an example realization of the CVS/CVM. Here, we assume that the vehicle control system, which evaluates and decides vehicle movements, resides in the same ECU as the CVS (i.e., the CVS ECU), but on two different processing units. The CVS-specific tasks may be executed on a processor that implements GPUs while the vehicle control might be executed on a dedicated safety processor. Within the CVS ECU, the CVS and vehicle control communicate via intra-ECU communication, which could be realized as an ECU-wide Ethernet network connecting the various processing units to each other.

The CVS executes tasks like video processing and video interpretation that result in extracted features describing the landmarks surrounding the car. This landmark information is then sent (over the intra-ECU Ethernet network) to vehicle control that evaluates the landmark information and decides on the further vehicle movement (and finally generates and transmits vehicle control commands to the actuators).

As depicted, the CVM can easily be added to the system. In this example the CVM is implemented on an ECU separate from the CVS and executes (a) a localization tasks that gathers information of the vehicle's position and trajectory, (b) a checkpoint selection task that selects relevant landmarks based on the vehicle's current position and movement, and (c) a checkpoint verification task, that controls whether the CVS system actually detected the checkpoint. For this checkpoint verification task, the CVM needs input from the CVS system. In our example the inter-ECU communication is also realized by means of an Ethernet network. Thus, the input to the CVM can simply be "mirrored" from the intra-ECU communication to the inter-ECU network.

It is unlikely that systems like the CVS/CVM will be implemented as isolated systems, but they will be rather integrated via vehicle-wide virtualization, a concept that extends and integrates the concept of a virtual machine to/with a virtual network. A distributed application will, thus, be unaware that it co-exists with other applications on the same physical system. Figure 5 depicts vehicle-wide virtualization using the example of the CVS/CVM. Here CVS/CVM is not exclusively using resources, but the resources are rather shared with other tasks (boxes in grey) and other messages on the inter-ECU and intra-ECU networks (shaded boxes).

Vehicle-wide virtualization requires guarantees for temporal and spatial partitioning, which can be achieved by deterministic system design [4]. Here, both, the hypervisor and the network follow deterministic computation and communication principles. A key aspect in deterministic design is a vehicle-wide synchronized time which can be achieved by various protocols, e.g., IEEE 1588 [5], IEEE 802.1AS [6], or SAE AS6802 [7], [8]. A second key aspect of deterministic design is the scheduling of major events at design time of the system [9], [10]. Such major events are for example the execution of partitions and tasks on a processor level or the communication of messages. The synchronized time in combination with the computation/communication schedule ensures that the defined major events are executed at exactly the right points in time.

Today, the technology to build vehicle-wide virtualization using deterministic design is already available. For example a cyber-physical systems (CPS) architecture for advanced driver assistance systems (as for example the CVS/CVM) can be composed of Deterministic Ethernet [11], [12] and automotive-graded microprocessors that support virtualization (e.g., upcoming ARM architectures [13], [14]).

IV. Potential Test Scenarios

In the previous sections we have introduced the concept of infrastructure-induced tests by using the example of a camera-vision monitor (CVM). We, furthermore, discussed a next generation CPS architecture that is capable of supporting upcoming ADAS applications and in particular infrastructure-induced tests. In this section we discuss in more detail which type of tests can actually be constructed and implications on the test setting and the CPS architecture.

While existing natural, architectural, or human-guiding and vehicle-guiding landmarks could be used as stimuli for these tests (i.e., used as vehicle-testing landmarks), it can be beneficial to the test accuracy to install dedicated vehicle-testing landmarks. For example, when testing a pedestrian-recognition ADAS, a vehicle-testing landmark would be a traffic sign with a picture of a pedestrian or a picture painted on the road itself. There may even be cases in which the same landmark will serve either as vehicle-guidance, vehicle-testing landmark, or both. So the cost involved in installing vehicle-testing landmarks could partly be recovered by these multi-purpose landmarks. In the remainder of this paper we refer to landmarks as vehicle-testing landmarks, knowing that they actually could serve multiple purposes.

Since the frequency of vehicle-testing landmark occurrence is one of the key factors to characterize a test's purpose and capability, we use this frequency of landmark occurrence to guide the discussion in this section.

A. Landmark occurrence: at specific points of interest

Probably the most cost-effective way to realize infrastructure-induced tests is the execution of these tests event-driven when reaching selected points of interests. Such points could be for example: freeway entrance/exit, before stop signs, at traffic lights, etc. – points at which the driving phase changes or the vehicle stops its movement. Tests executed at such points of interest could target ADAS applications specific to the following driving phase. For example, when leaving the freeway and entering urban areas the pedestrian recognition systems could be tested. Thus, latent system failures can be identified before the vehicle enters an area in which the system becomes relevant.

With respect to the implementation in the CPS architecture, the coupling of the CVS and vehicle control to the CVM tasks would rather be loose: reasonable delays of the CVM in informing the decision making task of vehicle control are acceptable, since the vehicle has either stopped or is not in a driving phase yet in which the tested system becomes relevant (e.g. the transition from freeway to urban areas will not be instantaneously).

B. Landmark occurrence: about once per hour

A landmark occurrence of about one landmark per hour is similar to the previous point-of-interest-triggered scenarios. However, the main difference is that the vehicle will not necessarily stop nor enter a new driving phase once per hour, so the immediate dependability measurement of certain phase-specific ADAS applications will in general not be possible. This regular tests will rather allow to gather relevant diagnosis data that can be used for maintenance, either scheduled or on-demand. Furthermore, the test could be used to re-calibrate sensors, e.g., stereo cameras.

Again, the coupling of the CVS and vehicle control to the CVM tasks is rather loose: since the test is targeting only diagnostic information of the vehicle status for future maintenance actions or sensor re-calibrations within a safety envelope.

C. Landmark occurrence: about once per minute

As the landmark density increases, the infrastructure-induced tests gradually transition from a dependability measurement method towards a method to actively improve the vehicle's current dependability. For example, availability of landmarks every minute could be useful as stimuli for latent failure detection and to test the CVS-internal safety monitors (i.e., to use the infrastructure-induced tests for "scrubbing").

Coupling of the CVS, and vehicle control to the CVM tasks becomes more relevant than in the examples before, since the output of the CVM will have greater impact on vehicle control. The proposed CPS architecture supports such a coupling.

D. Landmark occurrence: about once per second

When the vehicle is maneuvering at a regular traveling speed, then the placement of vehicle-testing landmarks close to each other (within tens of meters) becomes an interesting option. For example, let us assume that a vehicle is traveling at 50 km/h, then it will move about 14 meters per second (or 28 meters at 100 km/h). In space GPS provides a worst case accuracy of 7.8 meters at a 95% confidence level. The actual quality that users experience is impacted by various environmental effects. However, *real-word data from the FAA show that high quality GPS SPS receives provide better than 3.5 meter horizontal accuracy [15].* Furthermore, the accuracy of vehicle localization can be improved by additional wireless communication functions provided by the infrastructure as well as by taking the vehicle dynamics into account.

Thus, provided that landmarks are available with a distance of about ten meters from each other, test situations can be constructed in which new identifiable landmarks become available for test within about one second intervals (or at a hundred meters distance within ten seconds intervals) at regular traveling speeds. In case that the infrastructure provides this level of landmark density, then the infrastructure-induced test even becomes a relevant factor as immediate fault-detection measure and thereby directly improves the vehicle's dependability. Even if the infrastructure provides this landmark density only at selected segments, they can be used for temporary improved dependability while transitioning through the segments, e.g., during critical road curves.

On the other hand such segments with high landmark density can be used for more sophisticated tests, i.e., tests that require to process a sequence of landmarks or for more precise calibration purposes.

In order for the infrastructure-induced tests to accept landmarks at this frequency and having an immediate effect on vehicle control, CVS and vehicle control needs to be tightly coupled to the CVM tasks. For this, CVM, CVS, and vehicle control tasks can take full advantage of the deterministic design discussed or the CPS architecture: all tasks and communication can be scheduled and tightly coupled.

V. Research Directions

In this section we discuss some specific research directions in the area of infrastructure-induced tests.

How should vehicle-testing landmarks be constructed?

Computer-vision systems perform a variety of functions ranging from detection and recognition of road markings, traffic signs, pedestrians and vehicles on the road, or to measuring the exact 3D positioning of these objects on the road. The research question arises, how can we reason upon the correct CVS operation in a real situation from a successful infrastructure-induced test. For example, can we use simple geometric test patterns (like the one depicted in Figure 6) for vehicle-testing landmarks, or are more sophisticated landmarks necessary?

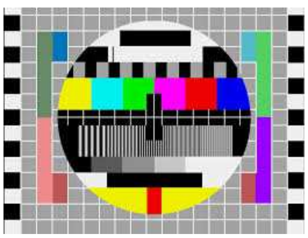

Fig. 6. Simple computer vision test pattern

How can the landmark data be safely and securely updated?

Infrastructure-induced tests rely on an accurate view of reality as for example provided by digital maps. Thus, it will be essential to maintain the accuracy of the landmark information inside the vehicle. As discussed, the landmark maintenance center can be used for this purpose. However, solutions need to be developed for it's reliable and secure information distribution.

Where are the sweet spots in cost balancing between vehicle and infrastructure investments?

Intelligent vehicles and the infrastructure they are operating in form together a system-of-systems (SoS) that needs to provide a high level of dependability. To reach this dependability, investments will be done both, in the vehicle and in the infrastructure. A relevant research direction is the detection of cost-optimized sweet spots and trade-offs of these investments.

VI. Conclusion

Intelligent vehicles rapidly improve their in-vehicle processing and communication capabilities and take benefit of upcoming CPS architectures that provide vehicle-wide virtualization. At the same time, information infrastructures become more and more available and accessible to intelligent vehicles. Advanced driver assistance systems already take advantage of these developments today to improve and extend their functions. In this paper we have proposed the complementary approach of infrastructure-induced tests, which also combine the benefits of novel CPS architectures and the availability of global information infrastructures, but rather to assess and enhance the dependability of the intelligent vehicle than to add functionality. We have discussed recent developments in CPS architectures that contribute to infrastructure-induced tests and used the computer-vision monitor as a running example. As discussed by various test scenarios, the application of infrastructure-induced tests is highly versatile and we consider infrastructure-induced test to increase their relevance as the intelligent vehicles increase their autonomy.

Acknowledgments

The research leading to these results has received funding from the ARTEMIS JU under grant agreement no.621429 (Project EMC2), from the Austrian Ministry for Transport, Innovation and Technology under the funding programme "IKT der Zukunft", grant agreement no.842567, and from the People Programme (Marie Curie Actions) of the European Union's Seventh Framework Programme FP7/2007- 2013/ under REA grant agreement no.607727.

References

[1] M. Maurer and H. Winner, *Automotive Systems Engineering*. Springer, 2013.

[2] A. U. Peker, O. Tosun, H. L. Akin, and T. Acarman, "Fusion of map matching and traffic sign recognition," in *Intelligent Vehicles Symposium Proceedings, 2014 IEEE*. IEEE, 2014, pp. 867–872.

[3] H. Jamshidi, T. Lukaszewicz, A. Kashi, A. Berghuvud, H. Zepernick, and S. Khatibi, "Fusion of digital map traffic signs and camera-detected signs," in *Signal Processing and Communication Systems (ICSPCS), 2011 5th International Conference on*. IEEE, 2011, pp. 1–7.

[4] S. Poledna, H. Kopetz, and W. Steiner, "Deterministic system design with time-triggered technology," in *Microelectronic Systems Symposium (MESS), 2014*. IEEE, 2014, pp. 1–4.

[5] K. Lee, J. C. Eidson, H. Weibel, and D. Mohl, "Ieee 1588-standard for a precision clock synchronization protocol for networked measurement and control systems," in *Conference on IEEE*, vol. 1588, 2005, p. 2.

[6] J. Teener and G. M. Garner, "Overview and timing performance of ieee 802.1 as," in *Precision Clock Synchronization for Measurement, Control and Communication, 2008. ISPCS 2008. IEEE International Symposium on*. IEEE, 2008, pp. 49–53.

[7] S. Standards, "SAE AS6802," 2011.

[8] W. Steiner and B. Dutertre, "Automated formal verification of the ttethernet synchronization quality," in *NASA Formal Methods*. Springer, 2011, pp. 375–390.

[9] S. S. Craciunas, R. S. Oliver, and V. Ecker, "Optimal static scheduling of real-time tasks on distributed time-triggered networked systems," in *Emerging Technology and Factory Automation (ETFA), 2014 IEEE*. IEEE, 2014, pp. 1–8.

[10] W. Steiner, "Synthesis of static communication schedules for mixed-criticality systems," in *Object/Component/Service-Oriented Real-Time Distributed Computing Workshops (ISORCW), 2011 14th IEEE International Symposium on*. IEEE, 2011, pp. 11–18.

[11] W. Steiner, G. Bauer, B. Hall, and M. Paulitsch, "TTEthernet: time-triggered ethernet," 2011.

[12] TTTech, "NXP and TTTech Announce Collaboration on Automotive Ethernet Switch Chip," http://www.nxp.com/news/press-releases/2013/09/nxp-and-tttech-announce-collaboration-on-automotive-ethernet-switch-chip.html [Last accessed 16/03/2015].

[13] R. Mijat and A. Nightingale, "Virtualization is Coming to a Platform Near You," http://www.arm.com/files/pdf/System-MMU-Whitepaper-v8.0.pdf [Last accessed 17/03/2015].

[14] J. Yoshida, "ARM Unveils New Architecture for Automotive," http://www.eetimes.com/document.asp?doc_id=1319873 [Last accessed 17/03/2015].

[15] N. National Coordination Office for Space-Based Positioning and Timingy, "GPS Accuracy," http://www.gps.gov/systems/gps/performance/accuracy/ [Last accessed 18/03/2015].