Paper number ITS-TP2189

TPEG3: Dynamic Information for Automated Vehicles

Matthias Unbehaun[1], Teun Hendriks[2], Martin Pfeifle[3], Pieter Pauwels[4]

[1] TISA – Traveller Information Services Association, Belgium, m.unbehaun@tisa.org

[2] TISA – Traveller Information Services Association, The Netherlands, t.hendriks@tisa.org

[3] NNG, Hungary, martin.pfeifle@nng.com ,

[4] Be-Mobile, Belgium, pieter.pauwels@be-mobile.com

Abstract

The TPEG2[1] protocol suite for traffic and traveller information is developed and maintained by the Traveller Information Services Association (TISA) for delivering real-time services to support drivers. With the advent of automated vehicles, the role of the drivers and their need for information are changing. As the vehicle will be controlling its own trajectory, it will need more detailed and precise external Information for Automated Driving (I4AD) to do this safely. This paper focusses on traffic-related aspects of I4AD; supporting AD vehicles and passengers to stay safe while driving. This paper outlines technical solutions that solve relevant technical and business issues in context with automated driving, addressing also that not all vehicles are automated. The paper describes example use cases and 'TPEG3 as a product' in terms of functional, non-functional and business requirements. Basic ideas and concepts that address these requirements are presented.

Keywords:

Traffic Information, Automated Vehicles, TPEG

Introduction

Real-time information to support Automated Vehicles (AD) is very relevant for AD vehicles to be able to properly drive through non-standard situations, such as roadworks where roadworkers may move about and dynamic changes in road layouts are to be expected. Adding connectivity to automated vehicles provides significant potential to extend what is possible with in-vehicle driving automation alone (see TISA Business White Paper [1]). Connectivity can dramatically improve situation awareness (well in extent of the sensor range), and thus safety of autonomous vehicles. Supplemental information, provided from an external source, may help offset (performance) limitations of the on-board sensors systems, and provide also information from, and about, other road-users. Connectivity can also provide current (possibly dynamic) traffic regulations and traffic control strategies, and support the safe, yet efficient, usage of available road infrastructure.

TPEG3 is targeted at both supporting and augmenting the in-vehicle situation awareness; this to support an autonomous vehicle's conformance with applicable traffic control strategies and policies. An automated driving vehicle operates principally following a *sense – plan – act* paradigm (the middle, grey block in Figure 1). A suite

[1] Transport Protocol Experts Group

of sensors scans the environment, records vehicle parameters, and creates an internal model of the world. This model of the world encodes the situation awareness of the vehicle and is supplemented typically with the (semi-)static information from an on-board digital road map.

Based on this internal model of the world, the “brain” of the autonomous vehicles determines then the best path to follow (with SAE level 2 or below still supervised or supported by humans). This is then transformed into actuation parameters to control the vehicles speed and steering systems for the automated driving to take effect.

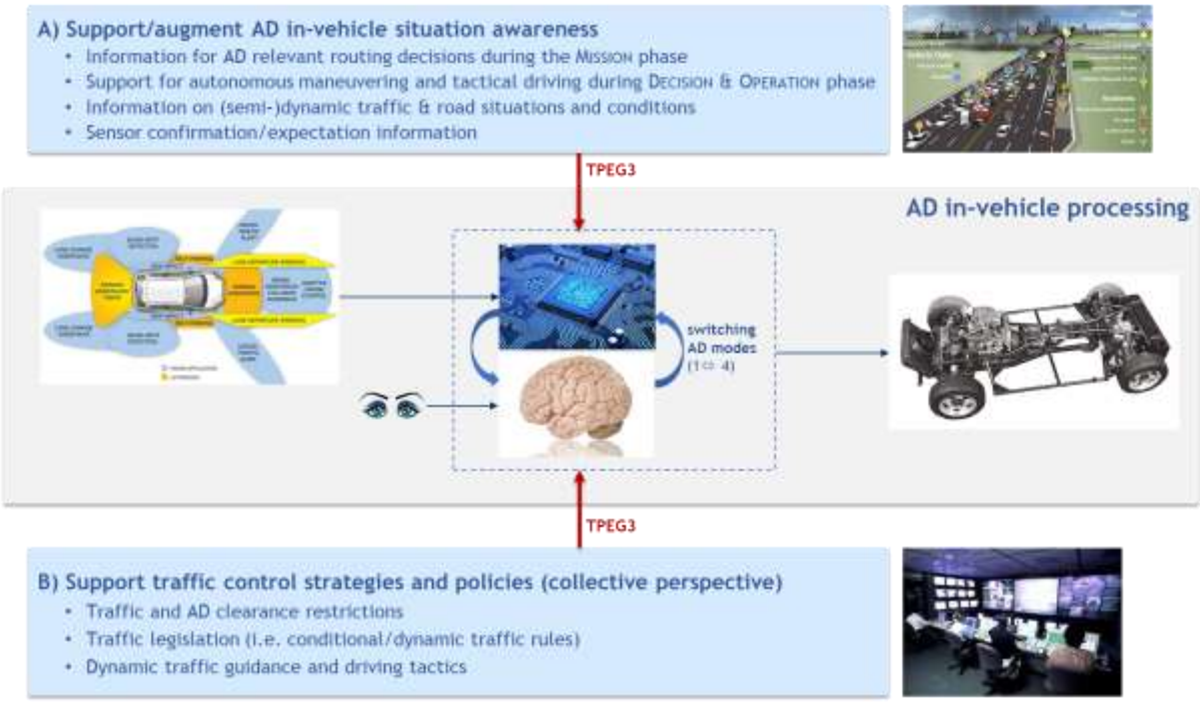

Figure 1 Positioning of TPEG3 with respect to Information for Autonomous Driving

Such AD systems work well on regular roads, in regular conditions. However, non-standard situations and/or non-standard/static traffic control strategies or policies may be beyond the vehicle’s ODD (Operational Design Domain), hence automation needs to be turned off in those situations. Here, TPEG3 is aimed at providing supplemental information, early and timely, to support the AD systems driving task where possible, or to provide necessary information well in advance of sensor range to allow an AD vehicle to smoothly hand-over driving back to the human (in an SAE level 3 AD system).

TPEG3 is geared towards supplementing on-board AD systems with supplemental information, both to support and augment situation awareness and to support conformance with traffic control strategies and policies.

The principal design features of TPEG3 as a connectivity product are the following.

- **Dynamic information provisioning to support Autonomous Driving through non-standard situations**. TPEG3 focuses on (dynamic) information provisioning to support autonomous driving while driving through non-standard situations or with non-standard traffic control strategies and policies (see Figure 2). This information about non-standard situation supports AD system to either extend their ODD (Operational Design Domain), or be able to timely warn the human to take over control in an SAE level 3 system,
- **Map attachment modus operandi**. TPEG3 has a focus on (dynamic) information provisioning. In doing so,

the principal modus operandi is that of a map attachment, i.e. temporary information supplementing or correcting the static information (i.e. map) timely and efficiently, hence more effective than with incremental map updates.

- **Safe and secure by design**. TPEG3 information targets usage in an AD system for safety-critical operations, i.e. operations whose failure or malfunction could result in death or serious injury to people, or loss or severe damage to equipment. Therefore, TPEG3 targets to be safe and secure by design, and incorporate relevant meta-information on quality, reliability, and provenance.
- **Generic, interoperable information delivery**. TPEG3 targets interoperable information delivery via different channels and delivery mechanisms. TPEG3 content will be deliverable via unidirectional or bi-directional channels, and support interaction as the information usage may re
- **Support towards highly dynamic information**. For the safety and control of AD vehicles, information timeliness is key, especially for that information in the immediate vicinity of the vehicle that may affect the imminent control actions of the AD vehicle. TPEG3 content encoding and delivery mechanisms allow fast local information delivery to support highly dynamic AD use cases.

An illustrative Road Works zone example

Road Works zones are known worldwide and can vary between simplified cases up to very complicated layouts. The subject 'road works' is actually of no interest to an AD vehicle. It only knows of actions that need to be taken, such as to slow down according to a given speed limit. The reason of this deceleration is information that could be mentioned to the passenger to explain why certain actions are taken by the vehicle; however, it is irrelevant for the vehicle itself.

To support Automated Vehicles to drive through Works Zones, the external information for an AD splits the information for a such work zone into different smaller components each focussing on supporting a specific action a vehicle must take.

Consider the very simple Road Works Zone example which is given in Figure 2, with three actions to be taken:

1. decelerate to match the given speed limit, prepare for sudden stop at the end of queue
2. approach the workers controlling the traffic with 'go / stop' signs and follow their signalling
3. when 'go' is signalled, pass the road works zone (on the opposite lane for traffic moving east-bound) around the traffic cones delimiting the drivable area while road construction machinery is moving on the east-bound lane, potentially protruding into the west-bound lane or creating fumes or soil spills, requiring additional manoeuvres by passing cars to avoid such impairments of the drivable area

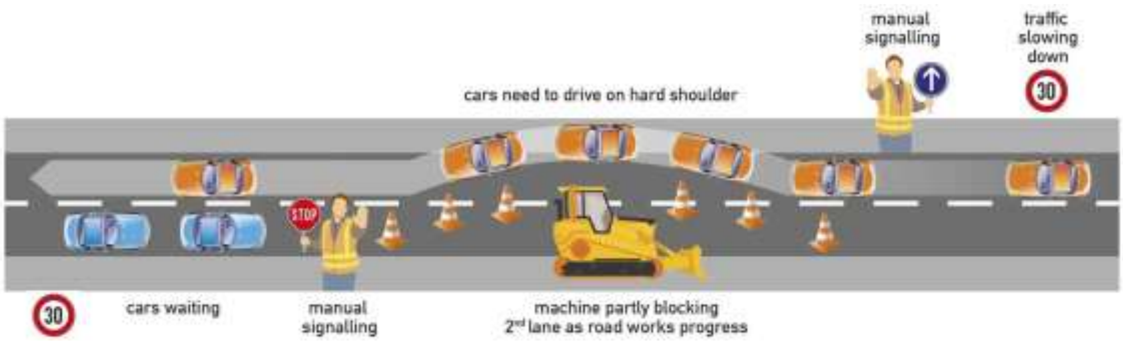

Figure 2 Example Work Zone

When approaching from far, getting closer, and finally passing the road works (or any similar situation), an AD vehicle needs different types and granularity of information. More general, we can expand this concept to all types of dynamic incidents or changes in traffic and separate three distinct phases as follows:

1. **Mission**: before the vehicle starts driving, a route is calculated to the desired destination. During the route, driving conditions will change (regulations, traffic situations, weather, ...) and a continuously recalculation of the route
2. **Decision**: once rerouting is no longer an option, a plan is created by the vehicle to take the necessary actions
3. **Operation**: once the decided location is reached, the actions must be performed

In the *Mission* phase, no fine-granular information is needed on cones, road signs, etc. The focus is more on travel time and whether AD is allowed and feasible on a specific route. Additional parameters such as driving comfort (e.g. uninterrupted AD driving) or fuel economy can be taken into account.

The *Mission* phase can start with a pre-trip planning, where an AD vehicle performs a route planning and needs to know whether relevant traffic situations are to be expected on the route. Then, the vehicle would after starting the trip and before approaching a given traffic situation, monitor the Mission information channel and receive more specific information.

In the *Decision* phase, rerouting is no longer an option. A plan is needed for vehicle to take the necessary actions. In the example, a deceleration must be performed before a defined location and other signage controlling the traffic obeyed. If traffic regulations state that AD is not allowed in the road works zone, the passenger must be alerted well ahead of time in order to get an orientation on the driving situation ahead and prepare for taking over the driving task.

Once the locations have been reached that were decided in the *Decision* phase, the actions should be performed. This is the *Operation* phase. In the example, deceleration is regulated at a specific location and heavy breaking should be avoided; the deceleration is executed smoothly in time. The same applies to all following actions while driving through the road works zone.

In both the *Decision* and *Operation* phase, the necessary information must be of a very high granularity. The locations where lanes are closed do not have much room for inaccuracy, for obvious reasons.

After the road work ends, driving can be considered as a ‘normal’, as defined in the *Mission* phase. Actions have been performed and the vehicle can continue its path. The *Decision* phase will become active again for the next traffic situation on the route.

Functional Concept

The TPEG3 functional concept is built around the following three phases (illustrated in the example) that are defined for each action a vehicle must take:

MISSION: Planning the trip, deciding on which route to take (e.g. fast/economic/all-AD), avoiding any issues on the route depending on the drivers/passengers’ preferences. Monitoring the route choice ‘where & what’. It is intended for ‘judging the situation’ and if necessary ‘wake up the driver’

DECISION: Getting closer, ‘coarse description of the scene’, medium-size information. It is intended for ‘preparing a strategy for managing the situation’

OPERATION: Situation directly ahead, ‘very fine-grained description of the scene’, heavy information, and intended for ‘detailed tactics how to negotiate obstacles’. It includes bounding boxes and free-space descriptions; to correlate with sensor inputs to validate the situation and increase reliability / minimize system error

The granularity of the information differs in each of the phases, and this for multiple ‘dimensions’. For TPEG3 we distinguish the following dimensions:

- location
- time, and applicable time window
- cause: nature of event
- effect: the consequence of this event, and impact on automated driving
- advice: non-binding recommendations and support
- metadata: information on quality, reliability and source of information
- rationale: explanation to the ‘passenger’ of the automated vehicle why the vehicle may drive differently than expected or customary on this road

This is explained in more detail in Table 1.

Table 1 - Phased approach with scalable information density

	Mission	**Decision**	**Operation**
Location	*for routing purposes* signpost locations, restriction zones, merging zones etc. lane topology + lane width	*meter accuracy* signpost locations, lane geometry, lane markings	*cm accuracy* lane markings
Time	evolving time window applicable for my route dynamic/moving work zone that requires info about the dynamics	dynamic/moving work zone with a higher precision	less relevant as location is essential here
Cause	only for display to the passenger if AD feasible and travel time impact	type of road furniture (e.g. cones, barriers) moving people/road workers	type of road furniture (e.g. cones, barriers) moving people/road workers
Effect	lane reductions narrow lanes and restrictions AD clearance (per lane)	longer braking distance slippery surfaces degraded road	degraded road sensor impact such as dust
Advice	traffic mgmt. centre advice *(e.g. dangerous goods avoid)* lane choice	where in zipper-zone to merge preferred trajectory Minimal Risk Condition (MRC)	staggered driving patterns in lane no overtaking position in lane

	Mission	Decision	Operation
	lane speed(s) safe distance	info	shockwave reduction emergency lane Minimum Risk Condition (MRC)
Metadata	Unsupervised AD operation confidence level quality on location lane width + variation(s)	quality on location lane width + variation(s)	additional information for the sensors
Rationale	'Roadworks upcoming' reason for lane change & speed reduction	success rate of prior AD vehicles	information to the driver about situation-awareness and intentions of the AD vehicle

The *dimensions* of the information should in this context be understood as components of a service or message. Each message contains the basic set *Location – Time – Cause – Effect – Advice – Metadata – Rationale* in always the same order such that the way to encode and decode it is implicitly defined. If a component is not needed, it will be empty. With increasing detail/granularity from Mission to Operation, the components will grow in size and may contain more and more sub-components.

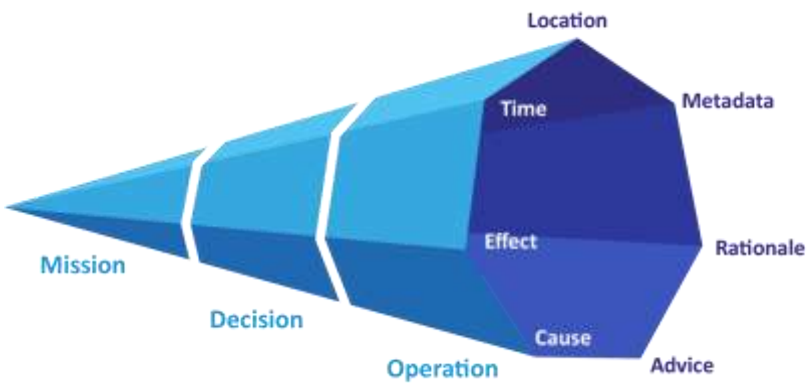

Figure 3 The MDO paradigm for hierarchical information encoding

The required granularity of the information necessary in each phase increases when the vehicle reaches the location of the action required. In the MISSION phase more information is needed on multiple traffic situations on different areas but the specific details for each of them would be an immense amount of data, useless at that moment. Reaching the specific locations (DECISION and OPERATION) would require more detailed information but only for this local situation.

The trigger for the information also differs between these phases.

MISSION is mostly **time-triggered**. A vehicle requests with a specific time-interval updated information on the global situation around the route, and to obtain a renewed view on the global situation yet monitoring of specific situations for impact on the plan is **location-triggered.**

DECISION and OPERATION are **location-triggered**. The definition of a lane ending, or a regulated speed reduction is performed in the space-domain. The speed of the vehicle is clearly an important factor in this. Figure 3 visualizes this MDO concept, showing that all information components (Location, Time, Cause, Effect, Advice/recommendation, Metadata and Rationale) are present in all phases.

Real-time aspects of information for AD vehicles

It is important to consider bandwidth, frequency and latency when discussing the supply of information to AD vehicles. For purposes of this discussion, we distinguish the following five categories:

a) infrequent, bandwidth-demanding types, such as map updates for larger areas (e.g. a region of an entire country) that are typically done when the vehicle is stationary
b) map tile updates, which consume less bandwidth since smaller areas are involved (1x1 or 2x2km) and compression can be applied, such that these types of map updates can be transmitted and ‘patched onto the base map’ while the vehicle is moving
c) near real-time information in bandwidth-efficient streams or packets with low latency (minutes to seconds) and regional context
d) real-time information in small chunks, with a limited, local context and very low latency (seconds to milliseconds)
e) real-time remote control of an AD vehicle by an operator in case of emergency

TPEG3 targets the categories c) and d), i.e. it will not provide map updates, nor will it allow some sort of remote control of the vehicle. It is rather an additional sensor for the AD vehicle with long (category c) to medium (category d) range, however providing information *beyond* the reach of the on-board sensors. A TPEG3 design principle hence is a *map attachment* modus operandi: provision of non-permanent dynamic information, to be attached to a map dynamically, but to be forgotten after the drive. This is expected to be at least *an order of magnitude* more efficient than incremental map updates, while providing for high refresh rates (seconds to sub-seconds) for the most local scope information. it is targeted to support easy filtering and easy on-board processing of the information.

TPEG3 technology considerations

In order to implement the aforementioned concepts, a number of technology choices have to be made, some of which will be elaborated in this section.

Overcoming Limited Visual Horizon of onboard sensors

Sensors currently used for AD systems typically use electromagnetic signals (radar, LiDAR, camera) to provide an image of what is around them. The range of these sensors is limited; their functioning can be inhibited by environmental factors, e.g. dense fog, rain, or obstructions from buildings, or other cars around the vehicle. The aim for TPEG3 is to extend this horizon; information provided by TPEG3 can be seen as a set of extra sensors for the vehicle’s ADS.

Qualifying Information

For information to be used by an AD systems, it has to be clear how trustworthy and applicable the information is in the local context. An AD system needs to use analytic means for assessing the quality of the information. The information captured by TPEG3 thus has to include sufficient meta-data so that an efficient algorithmic judgement of the information quality, integrity and trustworthiness is supported. Such meta-data will therefore be an integral part of TPEG3.

Actionable

TPEG3, which is to be processed by the on-board computer performing the Dynamic Driving Tasks (DDT), must be actionable. It therefore needs to contain more detailed and different types of information than TPEG2, where the driver 'translates' the information and performs the DDT manually:

- A higher location accuracy is required: lane-level and sub-lane-level accuracy will be required
- Level of detail of the information of the road ahead with respect to blocked lanes, road condition, etc. These can have an impact on the sensor's capabilities to traverse a road segment.

The TPEG2 model of communication is the distribution of a large set of messages every few minutes to all cars that can receive it. For TPEG3, a more flexible approach is required. Vehicle should be able to request for relevant information, as well as subscribe to such changes in relevant information.

Flexible Open Source Toolchain

TPEG3 aims at a wide adoption of its concept to AD research and development. It must therefore be easy to use and deploy TPEG3 technology, both in AD vehicles and in the supporting cloud infrastructure. To facilitate this, TPEG3 will rely as much as possible on open source toolchains.

Accurate Location referencing

Using TPEG3 technology, AD vehicles shall to be informed for example on which lane to take, where objects could obstruct the driving path, or about trajectories that other AD vehicles passing an incident successfully followed. For this, they need much more accurate information on road geometry, landmarks and other traffic-/driving-related objects. To support this, new location referencing methods need to be integrated for accurate location referencing as part of the TPEG3 development activities.

Outline towards a Technical Solution

This section summarises the current outline towards a technical solution for brevity. Further details and more detailed considerations are given in the TISA technical White Paper 1.

A Hybrid Interaction Paradigm

Multiple interaction paradigms can be envisioned for the information interchange between AD vehicles and road infrastructure or service providers. Early version of traffic information relied on a broadcast-type paradigm. Communicating with AD vehicles requires a more differentiated approach of bi-directional connectivity. Two principal (bi-directional) message exchange patterns apply:

1. **Request-response**: a message exchange pattern in which a requestor (typically the vehicle) sends a request message to a replier system (typically an infrastructure/service provider) which receives and processes the request, ultimately returning a message in response.
2. **Publish-subscribe**: a message exchange where senders of messages, called publishers, do not target the messages to be sent directly to specific receivers, called subscribers, but instead categorize published messages into classes without knowledge of which subscribers, if any, there may be. Similarly, subscribers express interest in one or more classes ('topics') and only receive messages that are of

interest to them, without knowledge of the publishers.

Comparing modes to the four phases identified in the MDO paradigm, we note the following:

- The *request-response* pattern fits well with *Mission* phase, which is (mostly) time-triggered;
- The *publish-subscribe* pattern fits well with the *Decision* and *Operation* phases, which are (mostly) location-triggered.

This suggest a hybrid approach. The *request-response* interaction is preferred for time-triggered interaction (and service discovery) when longer times between updates and higher information age/latency are acceptable. The *publish-subscribe* interaction is preferred for the mostly location-triggered interaction, where low information age/latency is important, yet much detail may stay unchanged.

Location referencing

Location referencing for TPEG3 is critical. Unlike TPEG2, it does not select one out of several possible methods for a given service. While TPEG2 compiles information based on type (e.g. incident, flow, weather), TPEG3 uses an approach, where information from different domains is aggregated into a ‘situation’, e.g. ‘roadworks’ or ‘accident’. These contain a compilation of information concerning obstacles blocking the e.g. the driving path, road surface, road layout and lane structure changes.

Multiple location referencing methods, tailored to the phase in the MDO model are under considerations, ranging from e.g. geo-hashes 3 for the *coarse, yet efficient, referencing* of areas for certain information types (Mission phase); on-the-fly location referencing schemes for *situations and objects that are not embedded on HD maps* (e.g. temporary road geometries or lane layouts, or temporary road furniture) during the Decision and Operation phases; HD map-like schemes for *identifying objects that are embedded on HD maps* (e.g. permanent road furniture, default road geometry, road signage) during the Decision and Operation phases.

Serialisation

Encoding information destined to AD vehicles needs to be efficient, as well as should be supported by a large community to help adoption. Contrary to TPEG2, TPEG3 will consider well-established off-the-shelf serialisation schemes such as Google Protocol Buffers [4] **,** Universal Binary JSON [5], or Efficient XML Interchange (EXI) [6]. Next to these, further options (see also Wikipedia [7]) identified are: ASN.1, CBOR, CDR, ZSERIO, the latter originated by the NDS community [8].

Over-the-wire protocol

The interaction paradigms chosen for TPEG3 implies a different choice for an over-the-wire protocol. TISA opts for a widely deployed, open source and standard solution. In particular, the following data distribution protocols are considered: MQTT [9] the *Message Queuing Telemetry Transport* publish-subscribe-based messaging protocol; and AMQP [10], a flexible publish-subscribe messaging protocol which also supports other communication patterns, e.g. request/response. Lastly, Data Distribution Services (DDS) [11], an OMG standard is an option, as is http. However, http does not support publish-subscribe properly, it is not very lightweight.

Conclusion

Having an information service to support automated vehicles is key to achieve automation Level 4 and higher, a finding that has been validated through projects such as AUTOPILOT [12] and echoed by platform such as

Connect Automated Driving [13]. The value that TPEG3 adds to the AD ecosystem reflects as well on road safety, specifically in mixed traffic situations and in context with complex situations that are challenging for AD vehicles to interpret and where external information can augment on-board sensor data to increase the probability of a correct interpretation and safe actions derived from these interpretations.

In this paper, we have outlined the need for extern information to support Automated Driving and described TISA's approach and outline towards a TPEG3 protocol to provide such information, targeting concepts that are scalable and implemented. TISA's approach and protocol proposal is open by nature, to be easily adopted by other parties, both inside and outside TISA. Further work is ongoing at TISA to detail and validate TPEG3 for Information for Automated Driving.

References

1. I4AD Business Whitepaper, R. van den Berg, T. Hendriks, P. Pauwels, M. Unbehaun. TISA19015.
2. I4AD Technical Whitepaper, R. van den Berg, T. Hendriks, P. Pauwels, M. Unbehaun. TISA19016.
3. Geo-Hash, https://en.wikipedia.org/wiki/Geohash
4. Google Protocol Buffers, https://developers.google.com/protocol-buffers/
5. UBJSON http://ubjson.org/
6. Efficient Extensible Interchange Working Group, https://www.w3.org/XML/EXI/
7. Comparing Data Serialization https://en.wikipedia.org/wiki/Comparison_of_data-serialization_formats
8. ZSERIO, https://github.com/ndsev/zserio
9. MQTT, https://mqtt.org/
10. AMQP, http://www.amqp.org/
11. DDS, https://en.wikipedia.org/wiki/Data_Distribution_Service
12. AUTOPILOT, https://autopilot-project.eu/
13. Connected Automated Driving, https://connectedautomateddriving.eu/