



Toward a Standard-Compliant Implementation for Consensus Algorithms in Vehicular Networks

Downloaded from: <https://research.chalmers.se>, 2026-04-04 07:29 UTC

Citation for the original published paper (version of record):

Cinque, E., Wymeersch, H., LINDBERG, C. et al (2018). Toward a Standard-Compliant Implementation for Consensus Algorithms in Vehicular Networks. IEEE Vehicular Technology Conference, 2018-August. <http://dx.doi.org/10.1109/VTCFall.2018.8690998>

N.B. When citing this work, cite the original published paper.

© 2018 IEEE. Personal use of this material is permitted. Permission from IEEE must be obtained for all other uses, in any current or future media, including reprinting/republishing this material for advertising or promotional purposes, or reuse of any copyrighted component of this work in other works.

Toward a Standard-Compliant Implementation for Consensus Algorithms in Vehicular Networks

Elena Cinque*, Henk Wymeersch†, Christopher Lindberg†, Marco Pratesi*

* DISIM Department, University of L'Aquila, L'Aquila, Italy

† Department of Electrical Engineering, Chalmers University of Technology, Gothenburg, Sweden

Email: elena.cinque@univaq.it, henkw@chalmers.se, christopher.lindberg@zenuity.com, marco.pratesi@univaq.it

Abstract—Cooperative Intelligent Transport System (C-ITS) applications require a continuous exchange of information between road users and roadside infrastructures. In this regard, distributed consensus algorithms can play an essential role in the definition of the information exchange rules between an ITS station and its neighbors. Although the consensus approach for networked systems is well-established, the efficiency of consensus methods under real-world vehicular communication constraints is largely unknown. This paper provides an ITS standard-compliant framework for analysis of consensus algorithms in vehicular networks with an emphasis on the role of robustness to changes in network topology in highly dynamic and dense environments. Our simulations reveal that in regular and realistic traffic conditions, the implemented consensus algorithm is able to achieve good performances in terms of both convergence time and needed consensus iterations. However, numerical results demonstrate that under dense and high-mobility traffic conditions the frequent exchange of large amounts of range information increases the Channel Busy Ratio (CBR) of the vehicular network and reduces the effectiveness of the algorithm as well.

I. INTRODUCTION

Cooperative Intelligent Transport Systems (C-ITSs) are evolving rapidly and have received significant attention in recent years due to the need to increase road safety and tackle growing emission and traffic congestion problems. The technology is ready and in many respects, existing vehicles are already connected devices. In such cooperative systems road users and traffic managers are able to share information and use it to coordinate their actions through vehicle to vehicle (V2V) or vehicle to infrastructure communications (V2I). First releases of the needed sets of specifications for C-ITSs have been defined, in Europe with the ETSI ITS G5 standard and in North America with the WAVE (Wireless Access in Vehicular Environment) architecture [1], that is based on the IEEE 1609 and 802.11-2012 families. Concerning the access technologies, both standards are based on IEEE 802.11p, specifically developed for vehicular ad-hoc networks (VANETs) in the 5.9 GHz dedicated short-range communication (DSRC) spectrum. Since C-ITS applications require reliable and accurate knowledge of the local neighborhood, cooperative awareness messages are periodically broadcast to one-hop neighbors with a frequency typically ranging from 1 to 10 messages per second. In such vehicular environments, the design of cooperative strategies can be supported on both centralized and distributed schemes. Distributed schemes reduce communication requirements us-

ing local V2X interactions, thus improving scalability, flexibility, reliability and robustness. In this context, *consensus* is a fundamental tool to be considered when designing distributed cooperative inference and optimization methods.

Consensus problems have a long history in computer science and form the foundation of the field of distributed computing [2]. The aim of *consensus* is to reach an agreement of interest for all the agents in a network, by only exchanging information locally following a predefined strategy. In particular, we consider the average consensus approach presented in [2] and exploited in [3], based on successive refinements of local estimates at vehicles and information exchange among neighbors. To date, numerous scenarios for consensus networks have been addressed in the literature for practical implementation. An emerging paradigm is cooperative localization, in which nodes help each other to determine their locations. In [3] average consensus is introduced to define an Implicit Cooperative Positioning (ICP) algorithm in which vehicle measurements and information are fused through V2V links. In [4], the authors describe several cooperative localization algorithms and quantify their performances through an extensive measurement campaign using US Federal Communications Commission (FCC)-compliant UWB radios. In [5], the authors propose a Bayesian method for distributed sequential localization of mobile networks, while in [6] a consensus-based method has been presented to enable the localization of an entire fleet of entities (i.e., vehicles, pedestrians or any objects), at each member of the fleet. In parallel, communication channel has become a pivotal issue because of its unreliability and noise, which may result in an impact on consensus. Link failures and channel noises have been discussed for distributed cooperative positioning in vehicular networks in [7] and more generally, stochastic link failures have been addressed for distributed average consensus in [8]. The performance of a consensus-based application in a VANET largely depends on the amount of neighborhood information that is available to a vehicle. For instance in a CP-oriented algorithm, the general intuition is that, the more neighbors that are in a vehicle's transmission range, the better accuracy can be achieved by CP [3]. However, the increase in the number of neighbors may, at the same time, deteriorate the network congestion conditions over the DSRC wireless channel.

In this paper, we characterize the performance of a generic

average consensus algorithm under several DSRC communication scenarios. Through high-fidelity simulations performed in a WAVE standard-compliant environment, we consider several network topologies and traffic conditions, with the aim to validate the investigated consensus method also in critical vehicle-dense and dynamic scenarios. Numerical results show that the accuracy achieved by the algorithm can be significantly impaired in dynamic topologies and critical congestion conditions, as the set of neighbors that a vehicle can actually detect is not complete.

II. INTEGRATION OF CONSENSUS-BASED METHOD INTO ITS ARCHITECTURE STANDARDS

A. Standard Average Consensus

Mathematically, average consensus is defined by a set of N_v nodes, an initial state $z_i^{(0)} \in \mathbb{R}^K$ for node i , and a set of neighbors J_i for node i . At iteration $r > 0$, agent i updates its local consensus variables by applying the following rule:

$$z_i^{(r+1)} = z_i^{(r)} + \sum_{j \in J_i} w_{i,j} (z_j^{(r)} - z_i^{(r)}), \quad (1)$$

where the weights w_{ij} are chosen to ensure convergence to the average:

$$\lim_{r \rightarrow \infty} z_i^{(r)} \rightarrow \frac{1}{N_v} \sum_{i=1}^{N_v} z_i^{(0)} \quad (2)$$

A common implementation of the weights is the so-called *Metropolis update* [8] where the weights are set to

$$w_{i,j} = \frac{1}{\max(d_i, d_j) + 1}, \quad (3)$$

where d_i and d_j denote, respectively, the number of neighbors of nodes i and j . Note, that Metropolis weights only require information exchange between neighbors to be set.

B. Consensus Implementation

In the present work, we consider the IEEE 802.11p technology, which is the basis of WAVE and ETSI ITS standards, in the USA and Europe, respectively. Specifically, in the simulations the WAVE protocol stack has been used since the network simulator (VEINS) provides a stable and accepted implementation; however, similar results can be obtained also for the ETSI standard. As mentioned above, by referring to IEEE 802.11p, cooperative awareness messages are periodically broadcast. For this purpose, in 2006 the US FCC designated the service channel (SCH) 172 (5855-5865 Hz) for V2V safety communications, i.e. to transmit the Basic Safety Messages (BSMs). In our network implementation of average consensus, only BSMs are exchanged and we encapsulate the consensus-related information into the BSM payload. The number of bits needed to represent the consensus-related information can change according to the different use-cases. However, the size of the single contribution is out of scope of the present paper and according to the SAE J2735-defined range, we will consider a BSM length of 150 Bytes in all the performed simulations. In each transmitted message, each

TABLE I: Simulation parameters setting

Parameter	Value
Frequency	5.9 GHz
Channel bandwidth	10 MHz
Modulation Rate	6 Mbps
BSM size	150 Bytes
Default BSM sending interval	100 ms

vehicle broadcasts (i) the transmitter ID, (ii) the consensus variable $z_i^{(r)}$, (iii) the number of neighbors and (iv) a bit flag to advertise the topology changes. The initial phase in which a vehicle learns the network topology is necessary for the nodes synchronization and it is hereinafter referred to as *learning period*. Thus, once each vehicle has knowledge of consensus variables in the network and associated vehicles, the consensus iterations commence, according to (1) and (3). After a consensus iteration, each vehicle broadcasts a new BSM with the updated information. The algorithm is then iterated, until one node detects a change in the network topology. In this case, all vehicles in the network start a new *learning period* and initialize the algorithm with a value $z_i^{(0)}$ equal to the last updated consensus variable. According to the specific use case, in a future work we could consider a cluster-based partitioning to avoid a *deadlock* situation that could be due to a overly wide network. In section III, the described strategy will be validated through a set of simulation trials under several vehicular scenarios.

III. PERFORMANCE EVALUATION AND SIMULATION RESULTS

In this section, we characterize the performance of the consensus method in a vehicular environment through simulation experiments. We first present the simulation setup in section III-A. Next, we evaluate the effectiveness of the implemented consensus algorithm for several non-line-of-sight (NLOS) conditions in both dense and mobile scenarios and provide numerical results.

A. Simulation Setup

After a survey on the available solutions, VEINS seems the best suited for our purposes: it is an open source inter-vehicular communication simulation framework based on an event-based network simulator (OMNeT++) and a road traffic microsimulation model (SUMO) running in parallel [10]. The primary parameters are listed in Table I. To simulate a realistic urban radio propagation environment, VEINS includes a simple obstacle shadowing model that has been calibrated and validated against real world measurements [11]. In our simulations, the DSRC transmission range is set to 300 m. Thus, we implemented a cross-road V2V scenario with a variable number of vehicles. To initialize the consensus variable, as soon as the simulation starts, a random initialization value $z_i^{(0)} \in \mathbb{R}$ is generated by each vehicle i according to a uniform distribution on the interval $[0, 100]$. We consider the following metrics: (i) convergence time, node-wise defined as the time (in seconds) until the deviation from the convergence value is




Fig. 1: Three different NLOS conditions in a V2V cross-road scenario.

within a factor of 0.15; (ii) Channel Busy Ratio (CBR), defined as the fraction of time the channel is regarded as busy; (iii) the number of sent messages before the consensus algorithm achieves convergence.

B. A Realistic SUMO-Simulated Crossroad Scenario

In the first set of performed tests, we consider a generic and static crossroad scenario and implement three basic network topologies. As depicted in Fig. 1, the network consists of 3 nodes in case A and case B, and 4 nodes in case C. Due to the obstacles' layout (shown in Fig. 1 by the grey area) and to the vehicles' positioning, each topology has a different network diameter, denoted as D_{\max} : the largest, minimal hop count over all pairs of nodes. As a benchmark, we first study the convergence speed of the implemented consensus algorithm under an ideal scenario, in which we assume: (i) at any time, each consensus agent (i.e. each vehicle) is aware of the whole network topology; (ii) each consensus agent collects the information from the neighbors isochronously. Then, we investigate the actual consensus algorithm performance in a realistic vehicular environment, according to the parameters setting described in section III-A. In the latter case, before the nodes start to perform the consensus algorithm, they collect all the necessary neighborhood information, according to the procedure described in section II. Fig. 2 shows the convergence time measured for a reference node (i.e. node 7), for three

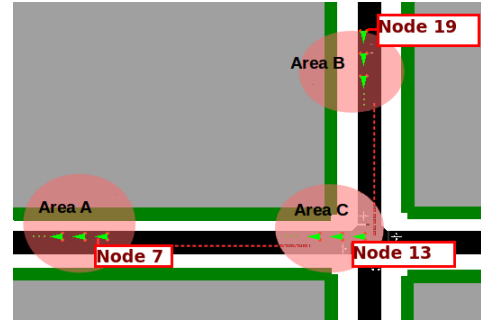


Fig. 3: Vehicles are distributed on 3 geographical areas (A, B and C), so that each area will include the same number of vehicles; the network topology has a diameter of $D_{\max} = 2$.

different values of D_{\max} in both ideal and realistic scenarios. As observed, the real-world communication constraints increase the convergence time of the algorithm by 200 ms, i.e., two additional BSM sending intervals. However, the two bars follow the same trend, and the convergence time value always remains below 600 ms. We note that the additional information introduced by node 25 in Fig. 1-c, allows node 7 to converge faster for both realistic and ideal scenarios: this behaviour is only due to the particular initialization values of the consensus variables.

C. Dense Vehicular Environment and Channel Congestion

Since the performance of a consensus-based application is expected to increase when the vehicles density grows [8] the idea to use a consensus approach is particularly attractive in C-ITSs. However, the intensive rate of information exchange over the shared DSRC control channel naturally introduces significant communication overhead into the vehicular network [9]. In this section, our simulation study will show that frequent range information exchange has a significant impact not only on the reliability of communications but on the consensus algorithm performance as well.

To assess the consensus algorithm performance in a dense environment, we consider the cross-road scenario depicted in Fig. 3, with a varying number of vehicles. As in Fig.1-b, the network diameter is $D_{\max} = 2$ and the vehicles are distributed on three geographical areas (A, B and C), so that

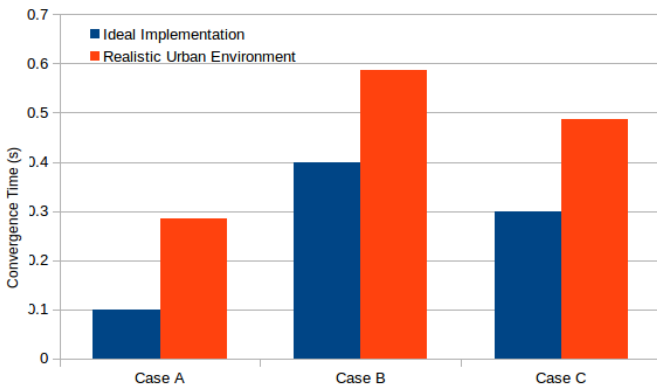


Fig. 2: Convergence time of the consensus algorithm in node 7 under different NLOS conditions.

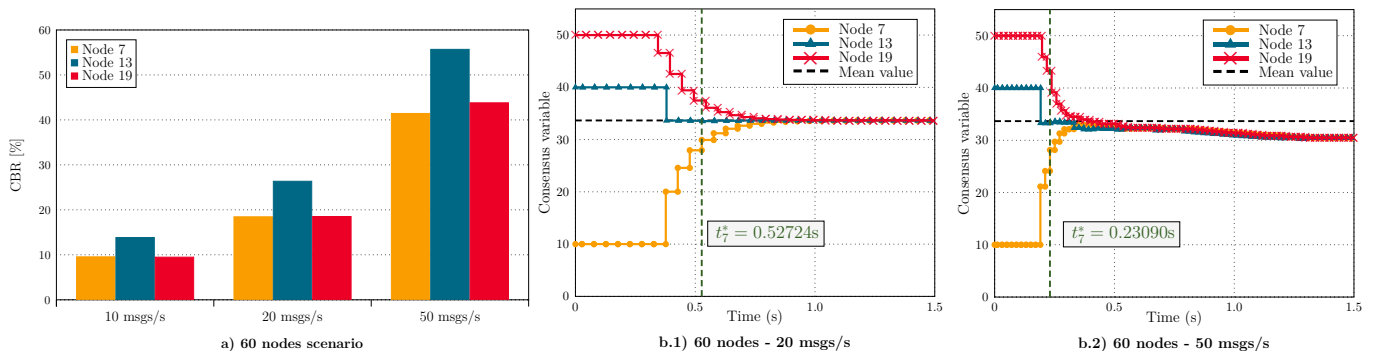


Fig. 4: a) CBR values for three different BSM rates; b) Consensus algorithm output as a function of time. t_7^* marks the convergence time for node 7.

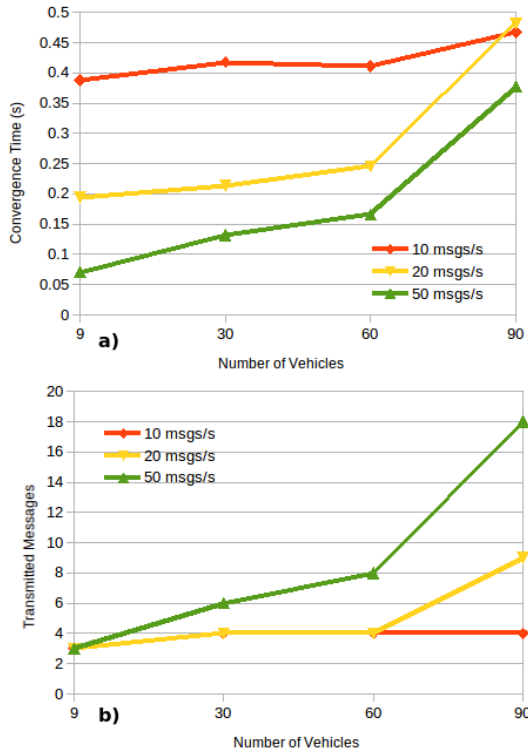


Fig. 5: a) Mean convergence time and b) mean number of transmitted messages necessary to converge. Results are measured on vehicles in area A.

each area includes the same number of vehicles. As in the previous simulations, the scenario is static and mobility is not implemented. Vehicles in area A can communicate with vehicles in area B. Vehicles in area C can also see vehicles in B. However, due to obstacles, vehicles in area A and C cannot see each other. In this test we assume that all vehicles in the same area have the same initial convergence variable $z_i^{(0)}$. Thus, we simulate different vehicle densities (from 3 to 30 vehicles per area) and different BSM sending intervals (from 10 msg/s to 50 msg/s according to the WAVE range of permitted values [12]).

Fig. 4-a shows the measured Channel Busy Ratio (CBR) in a 60-nodes network, for three nodes positioned in the three different areas, for three different BSM rates. Referring to the same scenario and to the same nodes, Fig. 4-b plots the consensus iteration outputs as a function of time. In Fig. 4-a, we note that the increment of the BSM transmit rate in a vehicle-dense scenario leads to a massive impairment of the V2V communications in terms of CBR. Therefore, as shown in Fig. 4-b, whilst a higher rate allows a faster consensus convergence, a higher CBR value (and the relative loss of disseminated messages) prevents convergence to the expected mean value of the consensus variables.

To further study the behavior, Fig. 5-a and Fig. 5-b show, respectively, the mean value of the convergence time and the mean value of the number of messages a node has to transmit to achieve convergence, over all nodes within area A. As expected, in Fig. 5-a we see that in a low-density scenario, better performance of the consensus algorithm can be achieved with higher BSM transmit rates. However, this assumption is no longer valid for a larger number of vehicles. In the worst case of 90 vehicles (i.e., 30 vehicles per area) the increase of the message rate from 10 to 20 msg/s can even lead to a growth of the convergence time. Moreover, the decrease of the convergence time with the highest rate of 50 msg/s is no longer profitable as in the lower-traffic density scenarios. This issue is further confirmed by results in Fig. 5-b. The mean number of BSMs to be transmitted to achieve convergence of the consensus algorithm grows rapidly with the BSM transmit rate. Due to the above-discussed CBR results, with a BSM rate of 50 msg/s we get a four-fold increase over the default setting rate of 10 msg/s. To summarize, for critical congestion conditions, there is an overall impairment of the consensus algorithm performances, in terms of both number of needed iterations to convergence and duration of the *learning period*, during which nodes collect information on the topology. Also, due to the message loss and the high CBR, the algorithm could be unable to converge to the average value.

D. Impact of Mobility in the Simulated Scenario

In this section, our consensus method implementation is run in a scenario with mobility. As shown in Fig. 6a, we have two

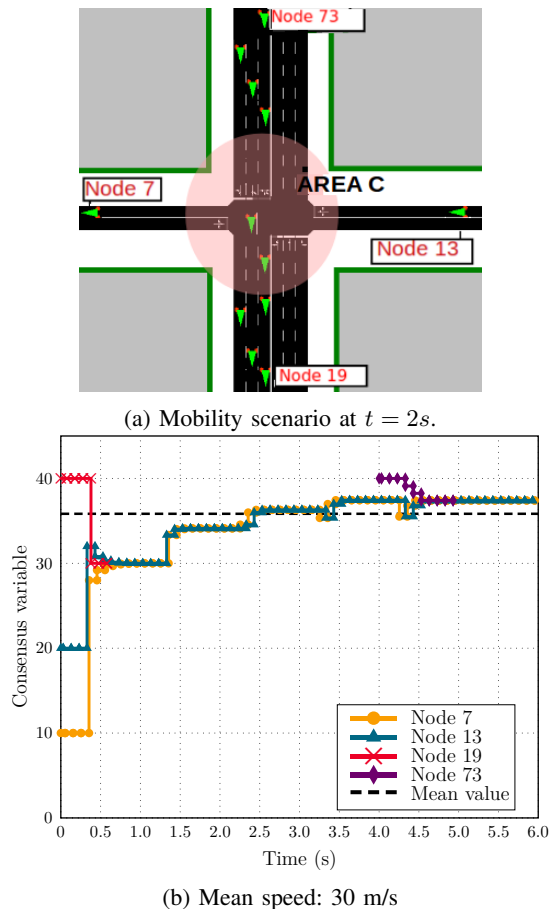


Fig. 6: a) V2V scenario with 10 mobile vehicles and 2 fixed vehicles (node 7 and node 13); b) consensus algorithm output as a function of the time for a mean speed of 30 m/s.

fixed nodes (i.e., node 7 and node 13) in a NLOS condition. Then, we added a mobile flow of 10 vehicles. When vehicles in the flow enter area C, they can communicate with both node 7 and node 13, along a 30 m stretch of road. Several simulations have been performed with different vehicle speeds with the aim to validate the implemented algorithm in a highly dynamic scenario. Simulation results show that when vehicles move at a medium-low speed (e.g., 10–20 m/s), despite the topology changes, the algorithm works correctly and the output values converge to the expected value. On the other hand, with higher speeds we observe that the convergence time starts to increase quickly. In Fig. 6b the algorithm output is depicted as a function of time, when vehicles in the flow move with a mean speed of 30 m/s. For the sake of clarity, we show just the first and the last node of the flow (i.e., node 19 and node 73) when they pass through area C. All mobile vehicles have the same initialization value $z_r^{(0)}$. In the case shown, we implemented in area C a flow with a density of 2 vehicles per second and we assume the first 2 vehicles are positioned within the area at time $t = 0s$. In Fig. 6b it is shown how, because of the rapid topology changes, the consensus converges to a value which is far from the expected average value. In the

figure it is also clear how each vehicle starts a new *learning period* periodically every second: this is related to the SUMO update-period of 1 s used to determine the vehicle position.

IV. CONCLUSION

This paper proposes an ITS-standard compliant framework to examine performance of a consensus algorithm implementation in a realistic vehicular environment. First, we addressed the issue concerning the implementation of distributed consensus algorithms in a generic vehicular environment. Then, we examined the effectiveness of a proposed consensus algorithm implementation under several NLOS conditions. Our consensus-based implementation has been validated in a real urban scenario using both the SUMO traffic simulator and OMNeT++. In the simulated high-density scenario, the network congestion conditions are analyzed and measured in terms of CBR. Numerical results show that the communication overhead significantly impacts the consensus algorithm performance in terms of convergence time and number of messages to be exchanged to achieve convergence. Finally, vehicle mobility has been added, in order to observe the response of the consensus algorithm in a dynamic scenario. As in the other simulated scenarios, also in the dynamic simulation, our results highlight the importance of a high ratio of the time interval between two successive topology changes (due to both vehicles speed and NLOS conditions) to the consensus convergence time.

ACKNOWLEDGMENT

The authors would like to acknowledge the support from EU ECSEL SAFECOP project and Ex-EMERGE project.

REFERENCES

- [1] E. Uhlemann, *The United States and Europe Advances Vehicle-to-Vehicle Deployment*, IEEE Vehicular Technology Magazine, vol.12, pp. 18-22, June 2017.
- [2] R. Olfati-Saber et al., *Consensus and Cooperation in Networked Multi-Agent Systems*, Proc. of the IEEE, vol. 95, no. 1, pp. 2152-233, 2007.
- [3] G. Soatti et al., *Implicit Cooperative Positioning in Vehicular Networks*, IEEE Trans. on Intelligent Transportation Systems, pp. 1-17, Mar. 2018.
- [4] H. Wymeersch et al., *Cooperative Localization in Wireless Networks*, Proc. of the IEEE, vol. 97, no. 2, pp. 427 - 450, Feb. 2009.
- [5] F. Meyer et al., *Distributed Localization and Tracking of Mobile Networks Including Non-Cooperative Objects*, IEEE Trans. on Signal and Inf. Processing over Networks, vol. 2, no. 1, pp. 5761, Mar. 2016.
- [6] G. Soatti et al., *Consensus-based Algorithms for Distributed Network-State Estimation and Localization*, IEEE Trans. on Signal and Information Processing over Networks, Mar. 2017.
- [7] J. Yao et al., *Improving Cooperative Positioning for Vehicular Networks*, IEEE Trans. on Vehicular Technology, vol. 60, no. 6, July 2011.
- [8] A. Nedić et al., *Network Topology and Communication-Computation Tradeoffs in Decentralized Optimization*, Proc. of the IEEE, pp. 1-24, Apr. 2018.
- [9] S. Zemouri et al., *An Altruistic Prediction-Based Congestion Control for Strict Beaconing Requirements in Urban VANETs*, IEEE Transactions on Systems, Man, and Cybernetics, pp. 1-16, Jan. 2018.
- [10] C. Sommer et al., *Bidirectionally Coupled Network and Road Traffic Simulation for Improved IVC Analysis*, IEEE Transactions on Mobile Computing, vol. 10, no. 1, pp. 3-15, Jan. 2010.
- [11] C. Sommer et al., *IVC in Cities: Signal Attenuation by Buildings and How Parked Cars Can Improve the Situation*, IEEE Transactions on Mobile Computing, vol. 13, no. 8, pp. 1733-1745, Aug. 2014.
- [12] U.S. Department of Transportation, *Conformance test specifications for Wireless Access in Vehicular Environments (WAVE) - Networking Services*, WAVENS-TSS&TP, v1.0, March 2016.