A Distributed and Privacy-Aware Speed Advisory System for Optimising Conventional and Electric Vehicles Networks

Mingming Liu, Rodrigo H. Ordóñez-Hurtado, Fabian Wirth, Yingqi Gu, Emanuele Crisostomi and Robert Shorten

***Abstract*—One of the key ideas to make Intelligent Transportation Systems (ITS) work effectively is to deploy advanced communication and cooperative control technologies among the vehicles and road infrastructures. In this spirit, we propose a consensus based distributed speed advisory system that optimally determines a recommended common speed for a given area in order that the group emissions, or group battery consumptions, are minimised. Our algorithms achieve this in a privacy-aware manner; namely, individual vehicles do not reveal in-vehicle information to other vehicles or to infrastructure. Mathematical proofs are given to prove the convergence of the algorithm, SUMO simulations are given to illustrate the efficacy of the algorithm, and hardware-in-the-loop tests involving real vehicles are given to illustrate user acceptability and ease of the deployment.**

***Index Terms*—SUMO, Distributed algorithms, Optimisation.**

I. Introduction

At present, Intelligent Speed Advisory (ISA) systems, as a part of Advanced Driver Assistance Systems (ADASs), have become a fundamental part of Intelligent Transportation Systems (ITS). Such systems offer many potential benefits, including improved vehicle and pedestrian safety, better utilisation of the road network, and reduced emissions. Recently, many papers have appeared on this topic reflecting the problem from the viewpoint of road operators, infrastructure providers, and transportation solution providers [1]–[8].

In this paper, we consider the design of a speed advisory system (SAS) making use of Vehicle-to-vehicle/infrastructure (V2X) technologies. Our starting point is the observation that different vehicle classes are designed to operate optimally at different vehicle speeds and at different loading conditions. Thus, a recommended speed, or speed limit may be optimal for one vehicle and not for others. Given a stretch of road network, the group emissions (CO, CO_2, NO_x, O_3, PM10, PM2.5) may or may not be close to the theoretically minimum possible. This of course depends on the composition of traffic on a given road, and the average speed at which vehicles are travelling.

This paper is a journal version of the idea presented at [9]. Our objective in this paper is to develop a SAS which allows groups of vehicles to collaborate in order to find the optimal speed at which the group should travel. We shall assume that vehicles are equipped with V2X technologies, and can exchange information with their neighbours and can exchange limited information with the infrastructure. We shall show that one can design, using very simple ideas, an effective SAS in a manner that preserves the privacy of individual vehicles. Extensive simulations, including hardware-in-the-loop (HIL) testing using real vehicles, are given to demonstrate the efficacy of our approach. Finally, we explore application of the approach to fleets of electric vehicles. As we shall see, the potential gains in this scenario are even more compelling than for conventional vehicles.

M. Liu and R. Ordóñez-Hurtado, joint first authors, are with the Hamilton Institute, Maynooth University, Maynooth, Ireland (e-mail: mliu@eeng.nuim.ie; Rodrigo.Ordonez@nuim.ie).

F. Wirth is with the Faculty of Computer Science and Mathematics, University of Passau, Germany (e-mail: fabian.lastname@uni-passau.de).

Y. Gu is with the School of Electrical, Electronic and Communications Engineering, University College Dublin, Ireland (e-mail: yingqi.gu@ucd.ie)

E. Crisostomi is with the Department of Energy, Systems, Territory and Constructions Engineering, University of Pisa, Italy (e-mail: emanuele.crisostomi@unipi.it).

R. Shorten is with IBM Research and University College Dublin, Dublin, Ireland (e-mail: ROBSHORT@ie.ibm.com).

II. Related work

In this section, we give a brief review of some related work. First note that a detailed review of this topic is given in [10]. Conventional systems are described in [5]–[7], [11], [12]. These papers describe various aspects of the ISA design process. This includes the design of driver display systems, the incorporation of external environmental information, and the algorithmic aspects of speed and distance recommendations. Recently, there has been a strong trend to also include traffic density information. References [10], [13]–[17] describe work in this direction. In these works density information is included in the procedure via loop detectors or via explicit density estimation via V2V technology. The differentiating feature of the approach followed in this paper is that density and composition of the vehicle fleet is also used, but in an implicit manner as part of the optimisation algorithm. Finally, we note that there is a huge body of work on cooperative control of vehicles and its connection to consensus algorithms [12], [18]–[20]. It is important to note that we are designing a SAS and not a cooperative control system. This distinction is important as it allows us to ignore string stability effects which are a fundamental limitation of many cooperative control architectures [21]–[24].

III. Model and Algorithm

A. Problem Statement

We consider a scenario in which a number of vehicles are driving along a given stretch of a highway with several lanes in the same direction. We wish to find a common recommended speed for these vehicles that takes into account the composition of the vehicles (individual vehicle types) and the number of vehicles. Note that the assumption on different lanes of the highway allows vehicles to overtake whenever it is appropriate. Let N denote the total number of vehicles on a particular section of the highway where the ISA broadcast signal can be received. Each vehicle is equipped with a specific communication device (e.g. a mobile phone with access to WiFi/3G networks) so that it is able to receive/transmit messages from/to either nearby vehicles or available road infrastructure (e.g. a base station). We assume that each vehicle can communicate a limited amount of information with the infrastructure, and that the infrastructure can broadcast information to the entire network of cars, and each vehicle can send a broadcast signal to its neighbours.

For convenience, we assume that all vehicles have access to a common clock (for example, a GPS clock). Let $k \in \{1, 2, 3, ...\}$ be a discrete time instant in which new information from vehicles is collected and new speed recommendations are made. Let $s_i(k)$ be the recommended speed of the vehicle $i \in \underline{\mathrm{N}} := \{1, 2, ..., N\}$ calculated at time instant k. Thus, the vector of recommended speeds for all vehicles is given by $\mathbf{s}(k)^{\mathrm{T}} := [s_1(k), s_2(k), ..., s_N(k)]$, where the superscript T represents the transposition of the vector. Note that between two consecutive time instants $(k, k+1)$, the recommended speeds are constant while the driving speeds are time-varying real-valued variables. We denote by N_k^i the set of neighbours of vehicle i at time instant k, i.e. those vehicles which can successfully broadcast their recommended speeds to vehicle i.

In addition, we assume that each vehicle i can evaluate a function f_i that determines its average emissions, were it to be travelling at the recommended speed $s_i(k)$. Such functions can be found in the literature [25], and are typically convex functions of the vehicle speed. We further assume that these functions are continuously differentiable and with a Lipschitz continuous first derivative f_i' which is assumed with positive bounded growth rate, i.e.

$$0 < d_{\min}^{(i)} \leq \frac{f_i'(a) - f_i'(b)}{a-b} \leq d_{\max}^{(i)}, \tag{1}$$

for all $a, b \in \mathbb{R}$ such that $a \neq b$, and suitable positive constants $d_{\min}^{(i)}$, $d_{\max}^{(i)}$. A schematic diagram of the above is illustrated in Fig. 1. In this context, we consider the following problem.

Problem 1: *Design an ISA system for a network of vehicles, following a common speed such as a speed limit, connected via V2X communication systems, such that the total emission from all vehicles can be minimised by all of them following the same reference speed.*

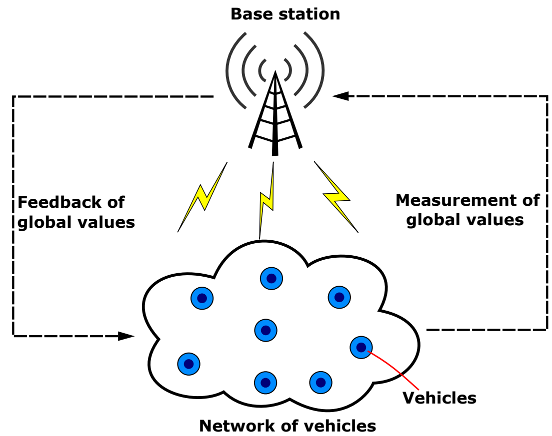

Fig. 1: Schematic diagram of the framework [26].

The optimisation problem that needs to be solved in order to address Problem 1 can be formulated as follows:

$$\begin{aligned} \min_{\mathbf{s} \in \mathbb{R}^N} \quad & \sum_{j \in \underline{\mathrm{N}}} f_j(s_j), \\ \text{s.t. } & s_i = s_j, \ \forall i \neq j \in \underline{\mathrm{N}}. \end{aligned} \tag{2}$$

This problem is an optimised consensus problem and can be solved in a variety of ways (for example using ADMM [27]–[29]). Our focus in this present work is not to construct a fully distributed solution to this problem, but rather to construct a partially distributed solution which allows rapid convergence to the optimum, without requiring the vehicles to exchange information that reveals individual cost functions to other vehicles. This is the privacy preserving component of our problem statement.

Comment: Note that in addressing Problem 1 we are not trying to calculate the recommended speed for all the vehicles in one step. Rather we propose an iterative algorithm that in each step yields individual recommended speeds that will eventually converge to the same value on the consensus constraints. Thus, our objective for the minimisation problem (2) is to seek the optimal solution of the recommended speeds under a consensus constraint. Clearly the constraints for vehicles to travel at roughly the same speed on a section of highway is a reasonable practical assumption.

To solve (2) we use the iterative feedback scheme

$$\mathbf{s}(k+1) = P(k)\,\mathbf{s}(k) + G(\mathbf{s}(k))\,e, \tag{3}$$

where $\{P(k)\} \in \mathbb{R}^{N \times N}$ is a sequence of row-stochastic matrices[1], $e \in \mathbb{R}^N$ is a column vector with all entries equal to 1, and $G : \mathbb{R}^N \mapsto \mathbb{R}$ is a continuous function with some assumptions to satisfy as we shall see in Theorem 1. Algorithms of these type were proposed and studied in [30]–[32]; the principal theoretical contribution here is to extend this framework to a new class of optimisation problems and to give conditions guaranteeing their convergence.

[1] Square matrices with non-negative real entries, and rows summing to 1.

We will require that (2) has a unique solution. Note that, it follows from elementary optimisation theory that if all the f_i's are strictly convex functions, then the optimisation problem (2) has a solution if and only if there exists a $y^* \in \mathbb{R}$ satisfying

$$\sum_{j=1}^{N} f_j'(y^*) = 0. \tag{4}$$

In this case by strict convexity y^* is unique and the unique optimal point of (2) is given by

$$\mathbf{s}^* := y^* e \in \mathbb{R}^N. \tag{5}$$

In order to obtain convergence of (3) we select a feedback signal

$$G(\mathbf{s}(k)) = -\mu \sum_{j=1}^{N} f_j'(s_j(k)). \tag{6}$$

and we obtain the dynamical system

$$\mathbf{s}(k+1) = P(k)\mathbf{s}(k) - \mu \sum_{j=1}^{N} f_j'(s_j(k))e, \quad \mu \in \mathbb{R}. \tag{7}$$

In [26] it is shown that if $\{P(k)\}_{k\in\mathbb{N}}$ is a uniformly strongly ergodic sequence[2] and μ is chosen according to

$$0 < \mu < 2\left(\sum_{j=1}^{N} d_{\max}^{(j)}\right)^{-1}, \tag{8}$$

then (7) is uniformly globally asymptotically stable at the unique optimal point $\mathbf{s}^* = y^* e$ of (2). Such systems were studied in [30] and formally analysed in [26]. For completeness, we formally state the relevant results from these works as a theorem (Theorem 1). An overview of the proof is given in the appendix and the interested reader may refer to [26] for details.

$$\begin{aligned} y(k+1) &= h(y(k)), \\ h(y) &:= y + G(ye). \end{aligned} \tag{9}$$

Theorem 1 ([26]) *Consider the optimisation problem* (2)*, the optimisation algorithm* (3)*, and the associated Lure system* (9)*. If G is defined by* (6) *and the condition* (8) *holds, then the following assertions hold:*

(i) There exists a unique, globally asymptotically stable fixed point $y^* \in \mathbb{R}$ of the Lure system (9)*.*

(ii) The fixed point y^* of (i) satisfies the optimality condition (4) *and thus $y^* e \in \mathbb{R}^N$ is the unique optimal point for the optimisation problem* (2)*.*

(iii) If, in addition, $\{P(k)\}_{k\in\mathbb{N}} \subset \mathbb{R}^{N\times N}$ is a strongly ergodic sequence of row-stochastic matrices, then $y^* e$ is a globally asymptotically stable fixed point for system (3)*.*

An outline of the proof can be found in Appendix A. To apply the above theorem to solve the optimisation problem we proceed as follows. For each k we define the $P(k)$ as

$$P_{i,j}(k) = \begin{cases} 1 - \sum_{j\in N_k^i} \eta_j, & \text{if } j = i, \\ \eta_j, & \text{if } j \in N_k^i, \\ 0, & \text{otherwise.} \end{cases}, \tag{10}$$

where i, j are the indexes of the entries of the matrix $P(k)$, and $\eta_j \in \mathbb{R}$ is a weighting factor. For example, a convenient choice η_j is $\frac{1}{|N_k^i|+1} \in \left(0, \frac{1}{N-1}\right)$, where $|\bullet|$ denotes cardinality, giving rise to an equal weight factor for all elements in the reference speed vector $\mathbf{s}(k)$.

The assumption of uniform strong ergodicity holds if the neighborhood graph associated with the problem has suitable connectedness properties. If sufficiently many cars travel on a given stretch it is reasonable to expect that this graph is strongly connected at most time instances. Weaker assumptions are possible but we do not discuss them here for reasons of space; see [33] for possible assumptions in this context.

Now, we propose the Optimal Decentralised Consensus Algorithm for solving (2) as shown in Algorithm 1. The underlying assumption here is that at all time instants all cars communicate their value $f_j'(s_j(k))$ to the base station, which reports the aggregate sum back to all cars. This is precisely the privacy preserving aspect of the algorithm, as cars do not have to reveal their cost functions to anyone. Also implicit information as derivatives of the cost function at certain speeds is only revealed to the base station but not to any other agent involved in the system.

Algorithm 1 Optimal Decentralised Consensus Algorithm

```
for k = 1, 2, 3, .. do
    for each i ∈ N̲ do
        Get F̃(k) = Σ_{j∈N̲} f'_j(s_j(k)) from the base station.
        Get s_j(k) from all neighbours j ∈ N_k^i.
        Do q_i(k) = η_i · Σ_{j∈N_k^i} (s_j(k) − s_i(k)).
        Do s_i(k+1) = s_i(k) + q_i(k) − μ · F̃(k).
    end for
end for
```

For the purpose of evaluation of the algorithm, we shall adapt the average-speed model proposed in [25] to model each cost function f_i in function of the average speed s as

$$f_i = \mathrm{k}\left(\frac{\mathrm{a} + \mathrm{b}s_i + \mathrm{c}s^2 + \mathrm{d}s^3 + \mathrm{e}s^4 + \mathrm{f}s^5 + \mathrm{g}s^6}{s}\right), \tag{11}$$

where $\mathrm{a}, \mathrm{b}, \mathrm{c}, \mathrm{d}, \mathrm{e}, \mathrm{f}, \mathrm{g}, \mathrm{k} \in \mathbb{R}$ are used to specify different levels of emissions by different classes of vehicles.

Comment: We note that in any real implementation on a given stretch of highway the recommended speed may be bounded above and below by the road operator.

[2] That is, for every $k_0 \in \mathbb{N}$ the sequence $P(k_0)$, $P(k_0+1)P(k_0)$, ..., $P(k_0+\ell)\cdots P(k_0)$, ... converges to a rank one matrix. See [26] for further details.

IV. Evaluation for Conventional Vehicles using SUMO simulations

We now evaluate the algorithm described in the previous section using the SUMO simulator [34], and by implementing it in a real vehicle. We conduct the following experiments:

1) First, we compare, for a given scenario, the optimal speed whit a non-optimum speed limit. In this simulation, for the purpose of the illustration, we force all vehicles to travel at the recommended speeds, subject to implementation constraints coming from the SUMO simulation (e.g. acceleration/deceleration profiles).
2) We then make the first scenario more realistic by allowing vehicles to travel in a small range around the recommended speed.
3) We then give a simulation which is dynamic in nature. Cars enter/leave the simulation dynamically and over a long stretch of road we allow vehicles to travel as they wish with an broad range of level speeds. For example, this may represent a highway situation where the traffic flow is moving better one lane than another lane.
4) We then give a hardware-in-the-loop simulation with a real target vehicle travelling on a road and an emulated network with a fixed number of simulated vehicles.
5) To conclude, we then consider networks with only electric vehicles. Due to the importance of this case, we give a detailed description in its own dedicated section.

The idea in all situations is to show the benefits of Algorithm 1. In the simulations we use the emission profiles from [25] shown in Table I and Fig. 2, corresponding to petrol cars/minibuses with up to 2.5 tons of gross vehicle mass. Besides, we also use the following vehicle types:

- Type 1: accel. 2.15 m/s^2, decel. 5.5 m/s^2, length 4.54 m.
- Type 2: accel. 1.22 m/s^2, decel. 5.0 m/s^2, length 4.51 m.
- Type 3: accel. 1.75 m/s^2, decel. 6.1 m/s^2, length 4.45 m.
- Type 4: accel. 2.45 m/s^2, decel. 6.1 m/s^2, length 4.48 m.

A. SUMO Simulations with a Fixed Number of Vehicles

In this experiment we consider 40 vehicles travelling along a highway. The set-up for this set of experiments is as follows.

- Road: A straight 5 km long highway with 4 lanes.
- Duration of the simulation: 1000 s.
- Algorithm sampling interval: ΔT = 1s.
- Switch-on time: the algorithm is activated at time 500 s.

The results of the experiments are given in Fig. 3 and Fig. 4. As can we see from Fig. 3a, making a small change in the recommended speed yields a significant (about 3%) improvement in CO_2 emissions. We then repeat the experiment and then allow vehicles to travel in a range around the recommended speed, and the results of this is given in Fig. 3b. As can we see from Fig. 3b, also in this case a significant reduction in CO2 emission can be observed. Finally to conclude, we applied the algorithm to a situation where the cars have a different range of emission profiles and with a different range of speeds. The results are given in Fig. 4. In this case, as can be seen in Fig. 4, a reduction of up to 8.07% can be obtained by following the recommended speed. To illustrate this means in terms of grams of carbon, SUMO predicts that these fourty vehicles emit an average of 9591 g/km compared with 8816 g/km when travelling at an optimised speed. This represents a saving of about 765 g/km which integrates over a day into a significant carbon saving.

TABLE I: Emission factors for some CO_2 emission profiles reported in [25].

Profile	a	b	c	d	e,f,g	k
R007	2.2606E+3	3.1583E+1	2.9263E-1	3.0199E-3	0	1
R016	3.7473E+3	1.9576E+2	-8.5270E-1	1.0318E-2	0	1
R017	3.7473E+3	1.8600E+2	-8.5270E-1	1.0318E-2	0	1
R018	3.7473E+3	1.6774E+2	-8.5270E-1	1.0318E-2	0	1
R019	3.7473E+3	1.5599E+2	-8.5270E-1	1.0318E-2	0	1
R021	3.7473E+3	1.0571E+2	-8.5270E-1	1.0318E-2	0	1

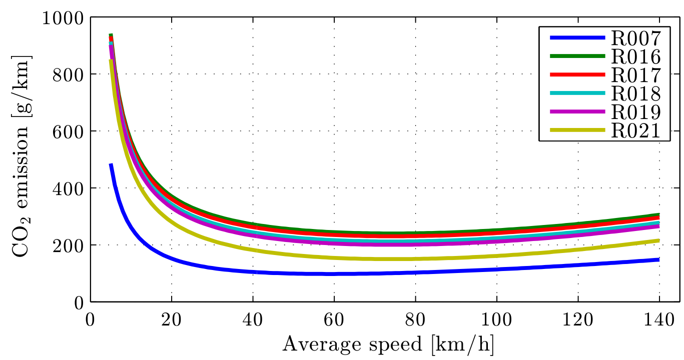

Fig. 2: Curves for the CO_2 emission profiles in Table I.

B. SUMO Simulations with a Dynamic Number of Vehicles

We now consider a dynamic scenario. To do this we partitioned the highway into three consecutive sections L1, L2 and L3. We then proceed as follows. First, vehicles enter the uncontrolled section L1, with constant speed (randomly chosen in a given range); after completing L1, vehicles enter the section L2. On section L2 vehicles calculate and follow a recommended speed. After completing L2 they enter section L3 and on this section they travel freely. The experiments are setup as follows.

- Road: three consecutive straight edges:
  - L1: 5 km long highway with 4 lanes, uncontrolled;
  - L2: 5 km long highway with 4 lanes, ISA controlled;
  - L3: 5 km long highway with 4 lanes, uncontrolled.
- Total number of cars: 650, with uniform distribution of both emission profiles among R016, R017, R018 and R019, and uniform distribution of types of vehicles.
- Vehicular flow entering L1: one new car every 2 seconds until simulation time 1300 s.
- Length of simulation: 3010 s.
- Window size for the calculation of the moving average (MA) of CO_2 emissions for visualisation purposes: 500 time steps.
- Travelling speeds for cars on L1 are randomly chosen with uniform distribution in 3 scenarios:

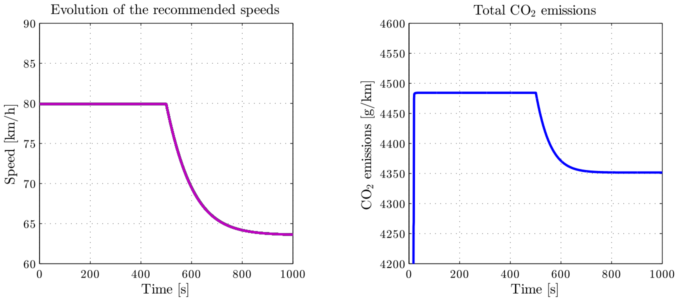

(a) All vehicles with constant speed 80 km/h until time 500 s, and following the recommended speed precisely after time 500 s.

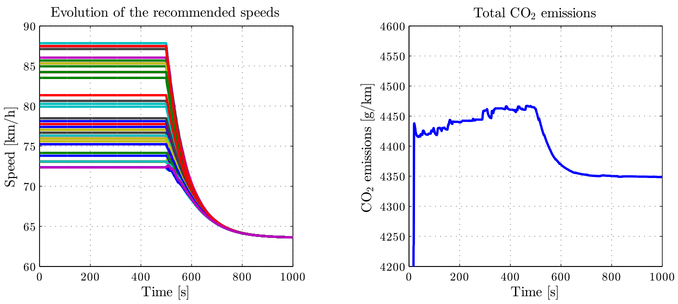

(b) All vehicles with constant speed in the range $(81, 99)$ km/h until time 500 s, and following the recommended speed with a maximum variation tolerance of 10% after time 500 s.

Fig. 3: Results of the SUMO simulation for the static case, before and after the activation of the algorithm at time step 500 s. Setup: 40 vehicles, of which 32 are of emission type R007 and 8 are of emission profile R021, and uniform distribution of types of vehicles.

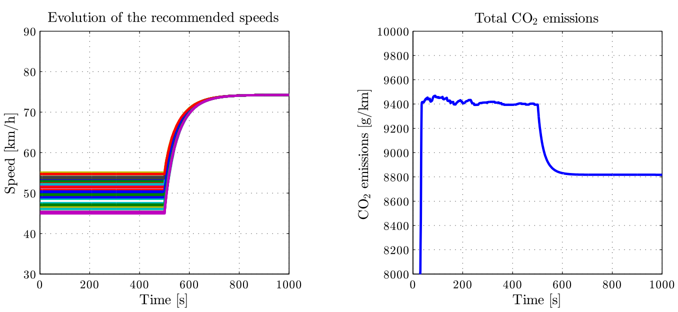

Fig. 4: Results of the SUMO simulation for the static case, before and after the activation of the algorithm at time step 500 s. Setup: 40 vehicles, with uniform distribution of emission profile among R016, R017, R018 and R019, and uniform distribution of types of vehicles.

- Case 1, constant speeds in $(80, 100)$ km/h.
- Case 2, constant speeds in $(60, 80)$ km/h.
- Case 3, constant speeds in $(40, 60)$ km/h.

Note that even though this is a dynamic situation, the vehicle density on each part of the road becomes constant after certain time. A sample of simulation results is given in Fig. 5, which reveals what might be expected from the initial experiments. Namely, the further vehicles are away from the optimal speed, the more is to be gained by deploying the ISA.

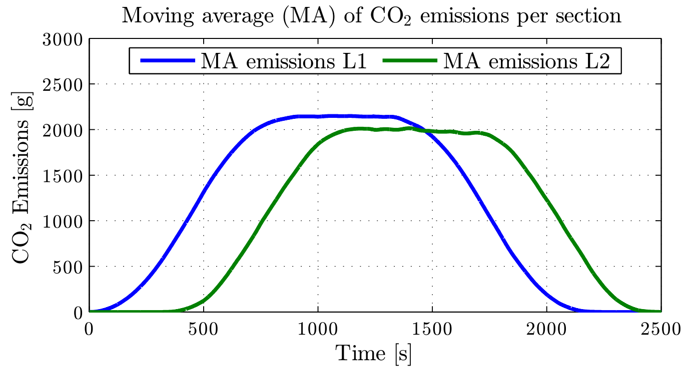

Fig. 5: Example for case 3: 500 time-step MA of CO_2 emissions for a set of initial speeds in the range $(40, 60)$ [km/h].

To complete the section, we conducted one hundred random experiments for each of the four cases described above. In each experiment, we collected the simulation data of the total CO_2 emission generation on each section of the highway from SUMO. Table II summarises the aggregated results of this exercise, and clearly demonstrate the benefits of the ISA.

TABLE II: General results: total emissions per lane.

	Total Emissions∗ [grams]					
	L1 (Uncontrolled)		L2 (Controlled)		Improvement	
Case	Mean	σ	Mean	σ	Mean	σ
1	2639012.7	1498.38	2587629.4	268.2	1.95%	0.05
2	2600710.6	606.87	2583472.8	178.3	0.66%	0.02
3	2787810.6	4200.36	2586943.9	169.9	7.20%	0.14

∗ Sum of emissions at every time step (i.e. time integration).
Mean: average of 100 different measurements.
σ: standard deviation.

Comment: Note that it is clearly the case that parameters of the algorithm have the potential to affect emission savings. For example, the speed of convergence of the algorithm affects the rate of which the emissions are saved.

Comment: Note that the solution that we have obtained is optimal for the environment and for the collective, e.g., in terms of overall reduced emissions. However, the solution might be unfair for some single users who would be recommended to drive at a different speed than originally desired. One way to improve fairness could be to decrease road taxes for virtuous vehicles, to compensate them from the inconvenience caused by the dirty vehicles in terms of recommended average speeds.

C. *Hardware-in-the-loop (HIL) emulation*

Finally, to give a feeling to a driver of how this system might function we now describe a hardware-in-the-loop implementation of the algorithm. Specifically, we use a SUMO-based hardware-in-the-loop (HIL) emulation platform that was developed at the Hamilton Institute [35], [36]. This emulation platform uses the open source road traffic simulator to emulate a real environment and generate virtual cars, along with a dedicated communication architecture supported by TraCI (a Python script implementing a TPC-based client/server architecture) to provide on-line access to SUMO, a smartphone connected to the 3G network and running the plug-in *SumoEmbed* (designed for use with Torque Pro [37], and a OBD-II adaptor [38] to embed a real car into the simulation, as shown in Fig. 6. The idea then is to allow the driver, driving a real vehicle on real streets, to experience being connected to a network of emulated vehicles driving along the same road network. Specifically, we performed this experiment by driving a Toyota Prius on a single-lane street circuit in the North Campus of the Maynooth University, while the Prius is embedded into a HIL emulation and represented by an avatar which interacts with the avatars of 29 other virtual (simulated) vehicles driving along the same stretch of (emulated) road.

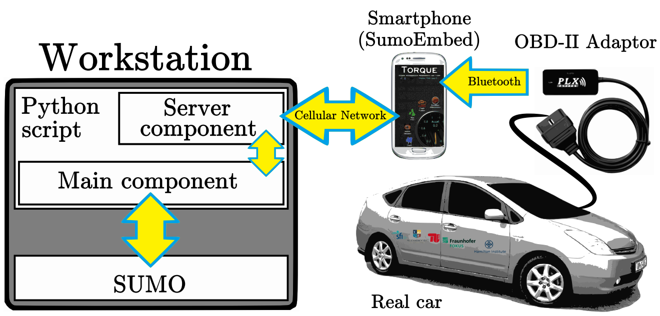

Fig. 6: Schematic of the HIL emulation platform.

The experiment begins when the simulation is started on the workstation and the server component of the Python script waits for a call from the OBD-II connected smartphone in the real vehicle. Since the selected street circuit only has one lane, the vehicles are released sequentially from the same starting point. The avatar representing the Prius departs in the sixth position. Once the connection between the Prius and the workstation is established, the position and speed of the Prius' avatar are updated using real-time information from the Prius via the OBD-II adaptor. From the point of view of the ISA algorithm, the Prius is regarded as a normal agent in the SUMO simulation, i.e. treated just like any other simulated vehicle.

The consensus algorithm for the proposed ISA system is embedded in the main component Python script. Thus, once the respective recommended speeds are calculated, they are sent to the vehicles via the server component and the cellular network to the smartphone in the case of the Prius, and via TraCI commands in the case of the other vehicles in the simulation. Note here that the driver behaviour is different for a simulated car compared to the case of the Prius: while we force each simulated vehicle to follow the recommended speed as far as possible[3], the Prius' driver is allowed to either follow or ignore the speed recommendation (displayed on the smartphone's screen) as desired.

[3]Concerning mainly the interaction between vehicles and the design parameters for the simulated cars such as acceleration, deceleration, car following model or driver information.

The HIL experiment is setup as follows.

- Length of the experiment: 600 s, of which the ISA algorithm is only engaged at around time 300 s;
- Total number of cars: 30, with uniform distribution of emission profiles among R016, R017, R018 and R019, and uniform distribution of types of vehicles, with a maximum speed of 100 km/h.
- The sampling time interval ΔT for collecting new information and updating the recommendations is 1 s.

Results of the experiment are depicted in Fig.7 and Fig.8. Fig. 7 shows that in the turned-off stage of the ISA system (i.e. the first 300 seconds), the overall CO_2 emissions increase almost linearly until all 30 vehicles are added to the emulation (at around 130 s). From this point, it can be observed that the total emissions oscillate around an average peak value of 713 g/km. Again from Fig. 7, we can observe that the overall CO_2 emissions reduce significantly once the the ISA algorithm is switched on, to an average of 475 g/km. In Fig. 8 (top), a comparison between the evolution of the Prius' driving and recommended speeds is presented. As can be observed, the recommended speed can be easily followed by the driver.

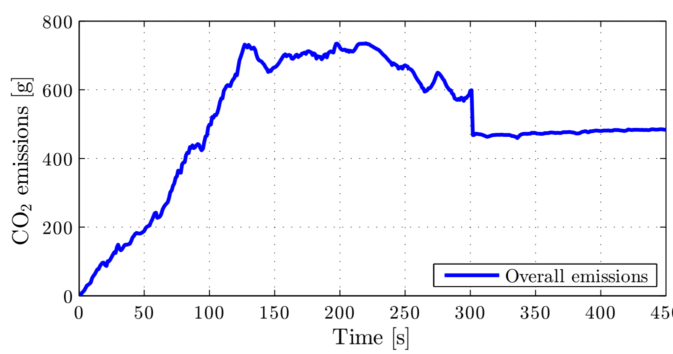

Fig. 7: Evolution of the overall CO_2 emissions. The algorithm was turned on around time 300 s.

V. THE CASE OF ELECTRICAL VEHICLES

In this section, we extend our previous discussion to the special case of a fleet of Electric Vehicles (EVs). The motivation for doing so is the recent interest in EVs as a cleaner alternative to conventional more polluting conventional vehicles. In recent years, some cities have decided to close the city centers to normal traffic, only allowing the transit to specific categories of low (or zero) polluting vehicles; see for instance [39]. Similarly, in many cities all around the world, the urban public transportation fleet has been restricted to electrically driven vehicles, as for the case of fleets of urban electric buses; see for instance the recent cases of São Paulo [40], Louisville in the US [41], or Wien in Europe [42].

Clearly, the previously designed Algorithm I needs to be adapted when applied to a fleet of EVs concerning different aspects. First, in terms of field of application, EVs are typically used for short distances due to their reduced driving ranges,

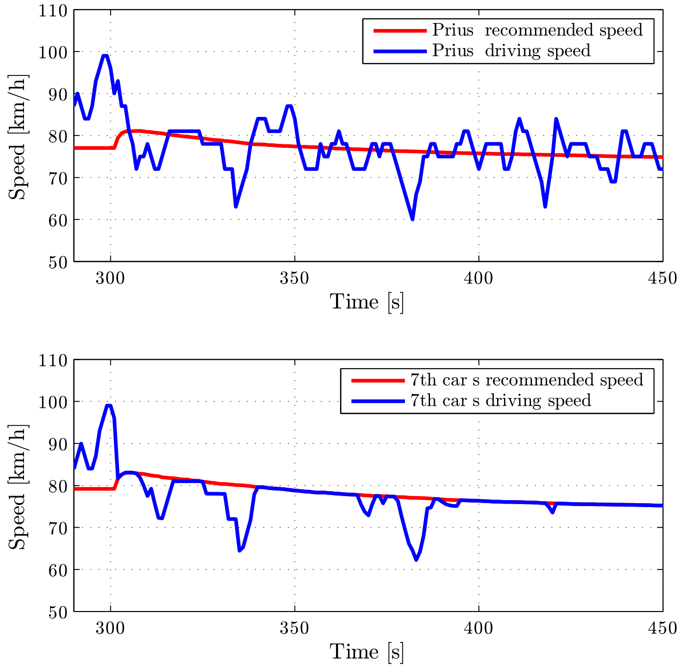

Fig. 8: Evolution of some variables along the whole HIL simulation. Top: related speeds for the Prius (the 6th car), and Bottom: the related speeds for the 7h car (just behind the Prius). The algorithm was turned on around time 300 s.

and thus are most likely deployed in city centers; second, the cost functions should not consider instantaneous emissions that, in the case of EVs, can be considered equal to zero. Accordingly, in the following we show how the previous framework can be adapted to the case of a fleet of EVs, and now the objective is to maximise the energy efficiency of the fleet of cars or, in other words, to extend their driving range.

A. *Cost functions to represent energy consumption in EVs*

Most of the discussion here follows the reference [43], where the ranges of EVs are reported for different brands and under different driving cycles. Power consumption in an EV driving at a steady-state speed (along a flat road) is caused mainly by four sources:

- **Aerodynamics power losses**: they are proportional to the cube of the speed of the EV, and depend on other parameters typical of a single vehicle such as its frontal area and the drag coefficient (which in turn, depends on the shape of the vehicle).
- **Drivetrain losses**: they result from the process of converting energy in the battery into torque at the wheels of the car. Their computation is not simple, as losses might occur at different levels (in the inverter, in the induction motor, gears, etc); in some cases, these power losses have been modelled as a third-order polynomial, whose parameters have been obtained by fitting some experimental data (see [43]).
- **Tires**: the power required to overcome the rolling distance depends on the weight of the vehicle (and thus, on the number of passengers as well), and is proportional to the speed of the vehicle.
- **Ancillary systems**: this category includes all other electrical loads in the vehicle, such as HVAC systems, external lights, audio system, battery cooling systems, etc. Here, the power consumption does not depend on the speed of the vehicle and can be represented by a constant term that depends on external factors (e.g., weather conditions) and personal choices (desired indoor temperature, volume of the radio, etc). According to experimental evaluations [43], the power losses due to ancillary services usually vary between 0.2 and 2.2 kW.

Thus, by summing up all the previous terms, then the power consumption P_{cons} can be represented as a function of the speed v as

$$\frac{P_{cons}}{v} = \frac{\alpha_0}{v} + \alpha_1 + \alpha_2 v + \alpha_3 v^2, \qquad (12)$$

where the left hand side is divided by the speed in order to obtain an indication of energy consumption per km, expressed in kWh/km. Such a unit of measurement is usually employed in energy-efficiency evaluations, and we shall assume that every single EV will use (12) as its personal cost function f_i. Accordingly, Fig. 9a shows a possible relationship between speed and power consumption, obtained using data from Tesla Roadster and assuming a low power consumption for ancillary services of 0.56 kW (i.e., assuming air conditioning switched off). As can be noted from Fig. 9a, there is a large energy consumption at large speeds due to the fact that power increases with the cube of the speed for aerodynamic reasons; however, it is also large for low speeds, due to the fact that travel times increase and, accordingly, constant power required by ancillary services demands more energy than the same services delivered with high speeds.

B. *Experimental results*

According to the previous discussion, we now assume that the objective is to infer the optimal speed that the ISA system should broadcast to a fleet of EVs travelling in a given area of a city (e.g., in the city centre). For this purpose, we assume that a fleet of 100 vehicles travels in the city centre for an hour, and following the next steps:

- In the first 20 minutes, the vehicles travel at the optimal speed calculated from Algorithm I.
- In the second 20 minutes, they travel at a speed below the optimal speed.
- In the last 20 minutes, they travel at a speed above the optimal speed.

In the first stage we assume that the communication graph among the EVs changes in a random way, i.e., at each time step an EV receives information from a subset of vehicles belonging to the fleet. This is a simplifying assumption that can be justified by assuming that in principle all vehicles might communicate to all the other vehicles (i.e., they are relatively close), but some communications might fail due to obstacles, shadowing effects, external noise, or other. Besides, in the two last stages we assume that the change of speed occurs instantaneously, since there is no requirement to iteratively compute an optimal speed.

We tuned our parameters in Algorithm 1 as $\eta = \mu = 0.001$, and we simulate different cost functions for each EVs by assuming a random number of people inside each car (between 1 and 5 people) with an average weight of 80 kg, and by assuming a different consumption from ancillary services within the typical range of $[0.2, 2.2]$ kW. The curves of the cost functions used in our experiment are shown in Fig. 9b. The evolution of the speeds of the EVs are shown in Fig. 10a, while the average energy consumption is shown in Fig. 10b.

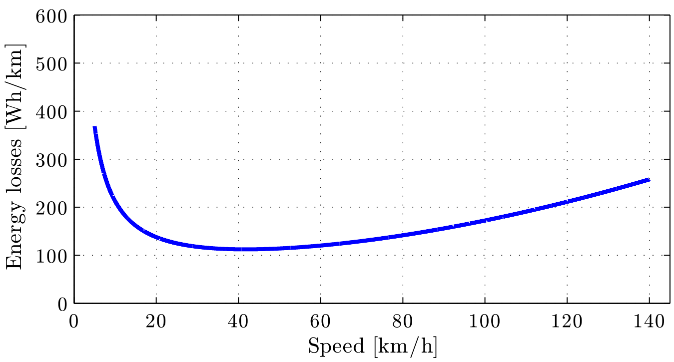

(a) An individual cost function.

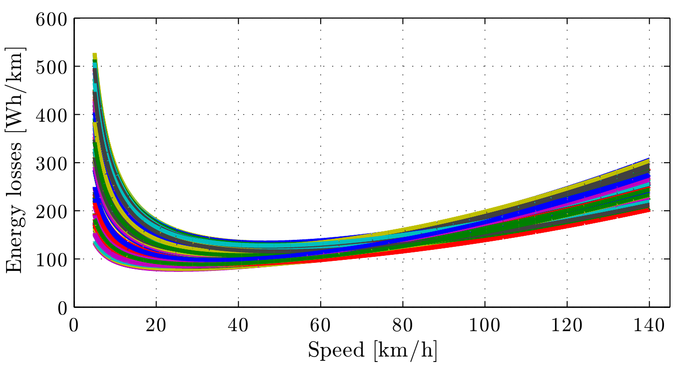

(b) All the cost functions overlapped.

Fig. 9: Curves for the cost functions used in the experiment. All of them were chosen convex.

As is shown in Fig. 10b, a simple ISA can be used for the case of a fleet of EVs just as in the case of conventional cars (for what regards the mathematical background) by simply adapting the cost functions to the new application of interest.

VI. Conclusion

In this paper we present a new ISA system. The system is based on a solution to an optimised consensus problem. We show that the ISA can be implemented in a privacy preserving manner, in a manner that accounts for vehicle density and composition, and in a manner that is provably convergent. Simulations and HIL emulations are given to illustrate the efficacy and the acceptability of the algorithm. Finally, the algorithm has been implemented in a real production vehicle using nothing more than a smartphone and a commercially available OBD-II plug-in.

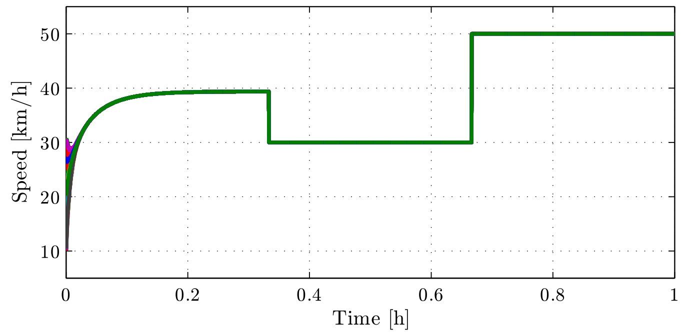

(a) Evolution of the vehicles' speeds.

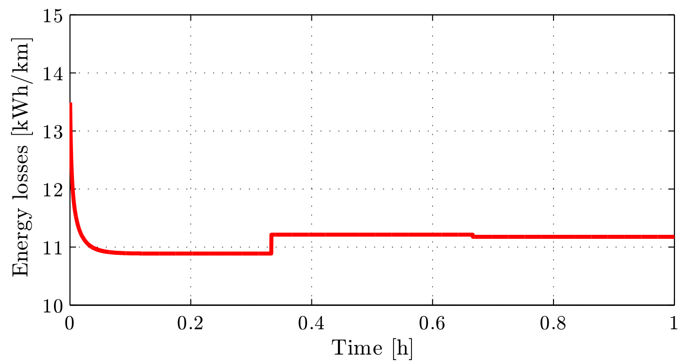

(b) Evolution of the overall energy loss.

Fig. 10: Simulation results for the network of EVs: Algorithm 1 is applied until time 0.33 h, and then two different speeds (below and above the optimal one) are suggested in $[0.33, 0.66]$ h and $[0.66, 1]$ h, respectively.

Appendix
Outline of Proof of Theorem 1

In this section, we give an outline of the proof for the claims in Subsection III-A in which we largely rely on the results obtained in [26].

The statement of Theorem 1 (i) is a consequence of the Banach contraction theorem. It is a straightforward calculation to show that the bound (8) ensures that the function h defining the Lure system (9) is in fact a global strict contraction on $\mathbb{R}$. Statement (ii) then follows directly from the definition of h: if $h(y^*) = y^*$ then $G(y^* e) = 0$ and by (6) this is equivalent to the optimality condition (4). The global optimality of $y^* e$ for the optimisation problem (2) of this optimal point then follows as (4) is the standard first order necessary condition for optimality and because strict convexity of the cost functions implies that this condition is also sufficient. Uniqueness is a further consequence of strict optimality.

It therefore remains to show that Theorem 1 (iii) holds. To this end we recall the following two lemmas from [26].

Lemma 2 ([26]) *Let* $\{P(k)\}_{k\in\mathbb{N}}$ *be a sequence of row-stochastic matrices. If* $\{y(k)\}_{k\in\mathbb{N}}$ *is a solution of the Lure system* (9) *then* $\{y(k)\, e\}_{k\in\mathbb{N}}$ *is a solution of* (3).

Lemma 3 ([26]) *Let $\{P(k)\}_{k\in\mathbb{N}}$ be a strongly ergodic sequence of row-stochastic matrices, and suppose that $G:\mathbb{R}^n\to\mathbb{R}$ is continuous and satisfies the following conditions:*

(i) there exists an $\varepsilon>0$ such that G satisfies a Lipschitz condition with constant $L>0$ on the set

$$B_\varepsilon(E) := \{x\in\mathbb{R}^n : \operatorname{dist}(x,E)\le\varepsilon\},$$

where $\operatorname{dist}(x,E) := \inf\{\|x-z\| : z\in E\}$ is the distance of a vector $x\in\mathbb{R}^n$ to the consensus set $E := \operatorname{span}\{e\}$; and

(ii) there exists constants $\beta,\gamma>0$ such that

$$|h(y)|\le|y|-\gamma \quad when \quad |y|\ge\beta,$$

where $h(y)=y+G(ye)$.

Then, every trajectory of (3) is bounded.

It is easy to see that G as defined in (6) satisifies the conditions of Lemma 3. Indeed, G is even globally Lipschitz continuous because of the Lipschitz continuity assumption (1). Furthermore, as h is a strict contraction on $\mathbb{R}$ with fixed point y^* we may denote the contraction constant of h by $0<c<1$ and obtain for any $y\in\mathbb{R}$ that

$$\begin{aligned}|h(y)| \le |h(y)-y^*|+|y^*| \quad &\le c\,|y-y^*|+|y^*| \\ \le c\,|y|+(1+c)\,|y^*| \quad &= |y|-(1-c)\,|y|+(1+c)\,|y^*|,\end{aligned}$$

from which it is easy to derive constants β and γ.

Finally, if every trajectory of (3) is bounded, then every trajectory has a nonempty bounded ω-limit set. Because of the averaging property of stochastic matrices and the assumption of uniform strong ergodicity, this ω-limit set is a subset of the span of e. By part (i) of the theorem, the Lure system has a globally asymptotically stable fixed point. Lemma 2 on the other hand ensures that on $\operatorname{span}\{e\}$ the trajectories of (3) and (9) (multiplied by e) coincide. It follows that restricted to $\operatorname{span}\{e\}$, the optimisation algorithm (3) has only one ω-limit set, namely y^*e. It then follows from a continuity argument that y^*e is a globally asymptotically stable fixed point of (3).

References

[1] E. Giordano, R. Frank, G. Pau, and M. Gerla, "CORNER: A realistic urban propagation model for VANET," in *7th Int. Conf. on Wireless On-demand Network Systems and Services*. IEEE, 2010, pp. 57–60.

[2] G. Yan, S. Olariu, and M. C. Weigle, "Providing VANET security through active position detection," *Computer Communications*, vol. 31, no. 12, pp. 2883–2897, 2008.

[3] A. Buchenscheit, F. Schaub, F. Kargl, and M. Weber, "A VANET-based emergency vehicle warning system," in *IEEE Vehicular Networking Conference (VNC2009)*, 2009, pp. 1–8.

[4] O. Carsten, F. Lai, K. Chorlton, P. Goodman, D. Carslaw, and S. Hess, "Speed limit adherence and its effect on road safety and climate change-final report," *tech. rep., University of Leeds - Institute for Transport Studies*, Oct 2008.

[5] E. Adell, A. Varhelyi, M. Alonso, and J. Plaza, "Developing human–machine interaction components for a driver assistance system for safe speed and safe distance," *IET Intelligent Transport Systems*, vol. 2, no. 1, pp. 1–14, 2008.

[6] N. Hounsell, B. Shrestha, J. Piao, and M. McDonald, "Review of urban traffic management and the impacts of new vehicle technologies," *Intelligent Transport Systems, IET*, vol. 3, no. 4, pp. 419–428, 2009.

[7] N. Tradisauskas, J. Juhl, H. Lahrmann, and C. S. Jensen, "Map matching for intelligent speed adaptation," *Intelligent Transport Systems, IET*, vol. 3, no. 1, pp. 57–66, 2009.

[8] M. Barth and K. Boriboonsomsin. (2010) "Intelligent" Ways to Cut Transportation's CO2 Emissions. University of California Transportation Center. [Online]. Available: http://www.uctc.net/research/briefs/PB-2010-03.pdf

[9] Y. Gu, M. Liu, E. Crisostomi, and R. Shorten, "Highway speed limits, optimised consensus, and intelligent speed advisory systems," in *3rd Conference on Connected Vehicles and Expo (ICCVE 2014)*, 2014.

[10] R. H. Ordóñez-Hurtado, W. M. Griggs, K. Massow, and R. N. Shorten, "Intelligent speed advising based on cooperative traffic scenario determination," in *Optimization and Optimal Control in Automotive Systems*. Springer, 2014, pp. 77–92.

[11] R. Gallen, N. Hautiere, A. Cord, and S. Glaser, "Supporting drivers in keeping safe speed in adverse weather conditions by mitigating the risk level," *IEEE Transactions on Intelligent Transportation Systems*, vol. 14, no. 4, pp. 1558–1571, 2013.

[12] R. M. Murray, "Recent research in cooperative control of multivehicle systems," *Journal of Dynamic Systems, Measurement, and Control*, vol. 129, no. 5, pp. 571–583, 2007.

[13] F.-Y. Wang, "Parallel control and management for intelligent transportation systems: Concepts, architectures, and applications," *IEEE Transactions on Intelligent Transportation Systems*, vol. 11, no. 3, pp. 630–638, 2010.

[14] L. Garelli, C. Casetti, C. Chiasserini, and M. Fiore, "MobSampling: V2V communications for traffic density estimation," in *73rd IEEE Vehicular Technology Conference (VTC Spring)*, 2011, pp. 1–5.

[15] V. Tyagi, S. Kalyanaraman, and R. Krishnapuram, "Vehicular traffic density state estimation based on cumulative road acoustics," *IEEE Transactions on Intelligent Transportation Systems*, vol. 13, no. 3, pp. 1156–1166, 2012.

[16] W. J. Schakel and B. van Arem, "Improving traffic flow efficiency by in-car advice on lane, speed, and headway," *IEEE Transactions on Intelligent Transportation Systems*, vol. 15, no. 4, pp. 1597–1606, Aug 2014.

[17] R. Ordonez-Hurtado, W. Griggs, K. Massow, and R. Shorten, "Evaluation of a new intelligent speed advisory system using hardware-in-the-loop simulation," in *Connected Vehicles and Expo (ICCVE), 2013 International Conference on*, Dec 2013, pp. 945–946.

[18] S. Kato, S. Tsugawa, K. Tokuda, T. Matsui, and H. Fujii, "Vehicle control algorithms for cooperative driving with automated vehicles and intervehicle communications," *IEEE Transactions on Intelligent Transportation Systems*, vol. 3, no. 3, pp. 155–161, 2002.

[19] M. C. Choy, D. Srinivasan, and R. L. Cheu, "Cooperative, hybrid agent architecture for real-time traffic signal control," *IEEE Transactions on Systems, Man and Cybernetics, Part A: Systems and Humans*, vol. 33, no. 5, pp. 597–607, 2003.

[20] F. Flemisch, J. Kelsch, C. Löper, A. Schieben, J. Schindler, and M. Heesen, "Cooperative control and active interfaces for vehicle assistance and automation," in *FISITA World automotive congress*, 2008.

[21] J. Hedrick, M. Tomizuka, and P. Varaiya, "Control issues in automated highway systems," *IEEE Control Systems*, vol. 14, no. 6, pp. 21–32, Dec 1994.

[22] D. Swaroop and J. Hedrick, "String stability of interconnected systems," *IEEE Transactions on Automatic Control*, vol. 41, no. 3, pp. 349–357, March 1996.

[23] S. Klinge and R. Middleton, "Time headway requirements for string stability of homogeneous linear unidirectionally connected systems," in *Proc. 48th IEEE Conference on Decision and Control, 2009 held jointly with the 2009 28th Chinese Control Conference. CDC/CCC 2009*, Dec 2009, pp. 1992–1997.

[24] S. Knorn and R. Middleton, "Stability of two-dimensional linear systems with singularities on the stability boundary using LMIs," *IEEE Transactions on Automatic Control*, vol. 58, no. 10, pp. 2579–2590, Oct 2013.

[25] P. Boulter, T. Barlow, and I. McCrae, "Emission factors 2009: Report 3-exhaust emission factors for road vehicles in the United Kingdom," *TRL Report PPR356. TRL Limited, Wokingham*, 2009.

[26] M. Liu, F. Wirth, M. Corless, and R. Shorten, "On global convergence of consensus with nonlinear feedback, the Lure problem, and some applications," *submitted to IEEE Transaction on Automatic Control*, 2014.

[27] S. Boyd, N. Parikh, E. Chu, B. Peleato, and J. Eckstein, "Distributed optimization and statistical learning via the alternating direction method of multipliers," *Foundations and Trends® in Machine Learning*, vol. 3, no. 1, pp. 1–122, 2011.

[28] W. Shi, Q. Ling, K. Yuan, G. Wu, and W. Yin, "On the linear convergence of the admm in decentralized consensus optimization," *IEEE Transactions on Signal Processing*, vol. 62, no. 7, 2014.

[29] T. Erseghe, D. Zennaro, E. Dall'Anese, and L. Vangelista, "Fast consensus by the alternating direction multipliers method," *IEEE Transactions on Signal Processing*, vol. 59, no. 11, pp. 5523–5537, 2011.

[30] F. Knorn, "Topics in cooperative control," Ph.D. dissertation, National University of Ireland Maynooth, Maynooth, Ireland, 2011.

[31] F. Knorn, R. Stanojevic, M. Corless, and R. Shorten, "A framework for decentralised feedback connectivity control with application to sensor networks," *International Journal of Control*, vol. 82, no. 11, pp. 2095–2114, 2009.

[32] F. Knorn, M. Corless, and R. Shorten, "Results in cooperative control and implicit consensus," *International Journal of Control*, vol. 84, no. 3, pp. 476–495, 2011.

[33] L. Moreau, "Stability of multiagent systems with time-dependent communication links," *IEEE Transactions on Automatic Control*, vol. 50, no. 2, pp. 169–182, 2005.

[34] SUMO, Simulation of Urban MObility. Last accessed: April 1, 2015. [Online]. Available: http://www.sumo-sim.org

[35] W. M. Griggs and R. N. Shorten, "Embedding real vehicles in SUMO for large-scale ITS scenario emulation," in *International Conference on Connected Vehicles and Expo (ICCVE), 2013*, 2013, pp. 962–963.

[36] W. M. Griggs, R. H. Ordóñez-Hurtado, E. Crisostomi, F. Häusler, K. Massow, and R. N. Shorten, "A Large-Scale SUMO-Based Emulation Platform," 2014, accepted in the IEEE Transactions on Intelligent Transportation Systems.

[37] I. Hawkins. Torque pro. Available from Google Play. [Online]. Available: https://play.google.com/store/apps/details?id=org.prowl.torque(Lastaccessedwebsiteon23July,2013)

[38] Plx devices inc. 440 Oakmead Parkway, Sunnyvale, CA 94085, USA. Phone: +1 (408) 7457591. [Online]. Available: http://www.plxdevices.com(Lastaccessedwebsiteon23July,2013)

[39] http://gis.uba.de/website/umweltzonen/umweltzonen.php.

[40] http://cities-today.com/2014/01/sao-paulo-to-introduce-its-first-fleet-of-fully-electric-buses/.

[41] http://cleantechnica.com/2015/01/16/louisville-gets-first-fully-electric-zero-emissions-city-buses/.

[42] http://www.siemens.com/innovation/en/home/pictures-of-the-future/mobility-and-motors/electric-mobility-electric-buses.html.

[43] R. Van Haaren, "Assessment of electric cars' range requirements and usage patterns based on driving behavior recorded in the national household travel survey of 2009," in *A Study conducted as part of the Solar Journey USA project*, 2011.

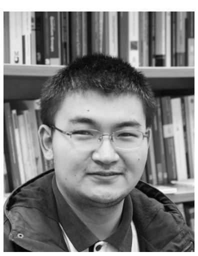

Mingming Liu received his B.E. degree in Electronic Engineering (first class honours) from National University of Ireland Maynooth (now Maynooth University) in 2011. He is currently working towards a Ph.D. degree in Applied Mathematics at the Hamilton Institute, Maynooth University. His current research interests are nonlinear system dynamics, distributed control techniques, modelling and optimisation in the context of smart grid and smart transportation systems.

Rodrigo H. Ordóñez-Hurtado received his degree in Engineering in Industrial Automatica from the University of Cauca, Colombia, in 2005. He commenced his Ph.D. in Electrical Engineering at the University of Chile in March 2008, and between April-November 2012, he was an intern at the Hamilton Institute, a multi-disciplinary research centre established at the National University of Ireland in Maynooth. In November 2012, he received his Ph.D. degree from the University of Chile, and since December 2012 he has held a postdoctoral position at the Hamilton Institute, working with Professor R. Shorten and his research group. Rodrigo's interests include robust adaptive systems (control and identification), stability of switched systems, swarm intelligence, large-scale systems and intelligent transportation systems. His focus is on applications to the mining industry and transportation systems.

Fabian Wirth received his PhD from the Institute of Dynamical Systems at the University of Bremen in 1995. He has since held positions in Bremen, at the Centre Automatique et Systèmes of Ecole des Mines, the Hamilton Institute at NUI Maynooth, Ireland, the University of Würzburg and IBM Research Ireland. He now holds the chair for Dynamical Systems at the University of Passau in Germany. His current interests include stability theory, queueing networks, switched systems and large scale networks with applications to networked systems and in the domain of smart cities.

Yingqi Gu received her B.E. degree in Electronic Engineering (first class honours) from Maynooth University in 2013. She obtained her M.Sc degree in Signal Processing and Communications at School of Engineering, University of Edinburgh in 2014. From February 2015, she commenced her Ph.D. degree in the University College Dublin with Prof. Robert Shorten. Her current research interests are modelling, simulation and optimisation of intelligent transportation systems.

Emanuele Crisostomi received the B.S. degree in computer science engineering, the M.S. degree in automatic control, and the Ph.D. degree in automatics, robotics, and bioengineering, from the University of Pisa, Italy, in 2002, 2005, and 2009, respectively. He is currently an Assistant Professor of electrical engineering with the Department of Energy, Systems, Territory and Constructions Engineering, University of Pisa. His research interests include control and optimization of large-scale systems, with applications to smart grids and green mobility networks.

Robert Shorten graduated from University College Dublin (UCD) in 1990 with a First Class Honours B.E. degree in Electronic Engineering. He was awarded a Ph.D. in 1996, also from UCD, while based at Daimler-Benz Research in Berlin, Germany. From 1993 to 1996, Prof. Shorten was the holder of a Marie Curie Fellowship at Daimler-Benz Research to conduct research in the area of smart gearbox systems. Following a brief spell at the Center for Systems Science, Yale University, working with Professor K. S. Narendra, Prof. Shorten returned to Ireland as the holder of a European Presidency Fellowship in 1997. Prof. Shorten is a co-founder of the Hamilton Institute at NUI Maynooth, where he was a full Professor until March 2013. He was also a Visiting Professor at TU Berlin from 2011-2012. Professor Shorten is currently a senior research manager at IBM Research Ireland. Prof. Shorten's research spans a number of areas. He has been active in computer networking, automotive research, collaborative mobility (including smart transportation and electric vehicles), as well as basic control theory and linear algebra. His main field of theoretical research has been the study of hybrid dynamical systems.