A Distributed and Privacy-Aware Speed Advisory System for Minimisation of Emissions in an ITS Scenario

Mingming Liu, Rodrigo H. Ordóñez-Hurtado, Fabian Wirth, Yingqi Gu, 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 highway in order that the group emissions are minimised. Our algorithms achieve this in a privacy-aware manual; 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 point of view of road operators, infrastructure providers, and transportation solution providers [1]–[7].

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.

M. Liu and R. Ordóñez-Hurtado 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)

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

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.

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 [8]. Conventional systems are described in [5]–[7], [9], [10]. These papers describe various aspects of the ISA design process. This include 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 [8], [11]–[15] 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 [10], [16]–[18]. 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 [19]–[22].

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. 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 [23], 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.*

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

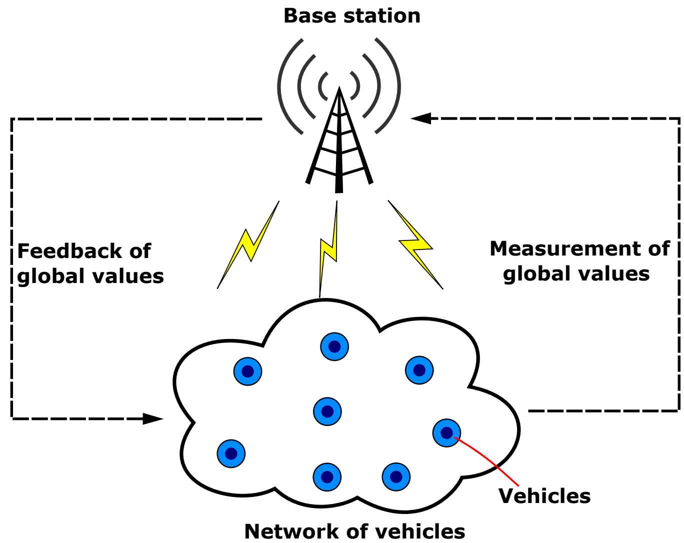

Fig. 1. Schematic diagram of the framework [24].

[25]–[27]). 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 an iterative feedback scheme of the form

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

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 [28]–[30]; the principal theoretical contribution here is to extend this framework to a new class of optimisation problems and to give conditions guaranteeing their convergence.

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)

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

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

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

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 [24] 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). For completeness, we formally state this as a theorem. In Theorem 1, we refer to a one-dimensional Lure system associated to (3). This new system is used to demonstrate the stability of system (3), for which the form of the Lure system is given by

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

Theorem 1 ([24]) *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, $\{P\left(k\right)\}_{k\in\mathbb{N}} \subset \mathbb{R}^{n\times n}$ 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 $P\left(k\right)$ as

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

[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 [24] for further details.

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

The assumption of uniform strong ergodicity holds if the neighborhood graph associated to 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 [31] 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 $f_j'\left(s_j\left(k\right)\right)$ 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 $i \in \underline{\mathrm{N}}$ **do**
    Get $\tilde{F}\left(k\right) = \sum_{j\in\underline{\mathrm{N}}} f_j'\left(s_j\left(k\right)\right)$ from the base station.
    Get $s_j\left(k\right)$ from all neighbours $j \in N_k^i$.
    Do $q_i\left(k\right) = \eta_i \cdot \sum_{j\in N_k^i} \left(s_j\left(k\right) - s_i\left(k\right)\right)$.
    Do $s_i\left(k+1\right) = s_i\left(k\right) + q_i\left(k\right) - \mu \cdot \tilde{F}\left(k\right)$.
  end for
end for

For the purpose of evaluation of the algorithm, we shall adapt the average-speed model proposed in [23] to model each cost function f_i:

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

where $\mathrm{a}, \mathrm{b}, \mathrm{c}, \mathrm{d}, \mathrm{e}, \mathrm{f}, \mathrm{g} \in \mathbb{R}$ are used to specify different levels of emissions by different classes of vehicles as shown in Table I. The corresponding cost functions for the CO_2 emission are given in Table I.

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

IV. Evaluation using SUMO simulations

We now evaluate the algorithm described in the previous section using SUMO [32]. We consider in the following three scenarios.

1) A network with fixed number of simulated vehicles (static case).

TABLE I
DIFFERENT STANDARD CO_2 EMISSION FACTORS [23].

	CO_2 emission type			
	Euro I	Euro II	Euro III	Euro IV
a	3.7473×10^3			
b	1.9576×10^2	1.8600×10^2	1.6774×10^2	1.5599×10^2
c	-8.527×10^{-1}			
d	1.0318×10^{-2}			
e	0			
f	0			
g	0			

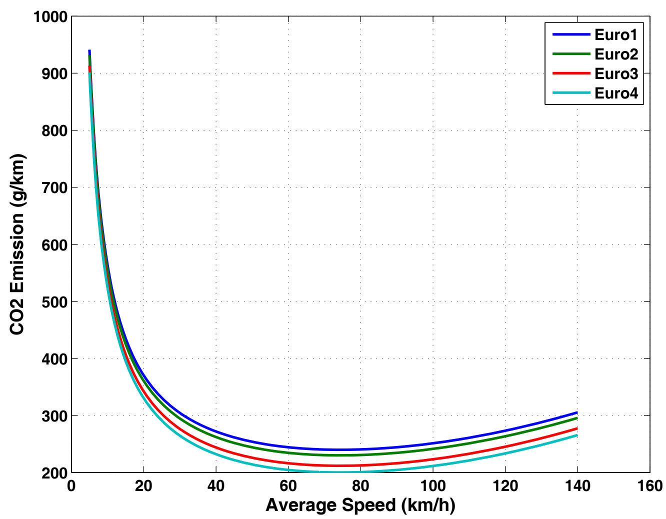

Fig. 2. Curves for the CO_2 emission types presented in Table I.

2) A network with a dynamic number of simulated vehicles (dynamic case).
3) 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.

The idea in all situations is to find the speed that minimises the emissions using Algorithm 1, and to document the benefits of all vehicles following the recommended speed. In all simulations we assume that there are four classes of vehicles. These are described below.

- type 1: Euro 1 emission class, acceleration 2.15 m/s^2, deceleration 5.5 m/s^2, length 4.54 m;
- type 2: Euro 2 emission class, acceleration 1.22 m/s^2, deceleration 5.0 m/s^2, length 4.51 m;
- type 3: Euro 3 emission class, acceleration 1.75 m/s^2, deceleration 6.1 m/s^2, length 4.45 m;
- type 4: Euro 4 emission class, acceleration 2.45 m/s^2, deceleration 6.1 m/s^2, length 4.48 m.

A. SUMO Simulations with a Fixed Number of Vehicles

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 has converged).
- Algorithm sampling interval: ΔT = 1s.
- Switch-on time: the algorithm is activated at simulation time 500 s.
- Number of vehicles: 10 vehicles of each available type.
- Initial speeds: randomly chosen with uniform and independently distribution in $(40, 60)$ [Km/h].

Comment: For comparison purposes, the assumption on the number and the initial speed distribution of vehicles are consistent with our previous work in [33]. Results of the experiments are presented in Fig. 3. As can be seen, an instantaneous reduction of up to 6.13% 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 9582 g/Km compared with 8817 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.

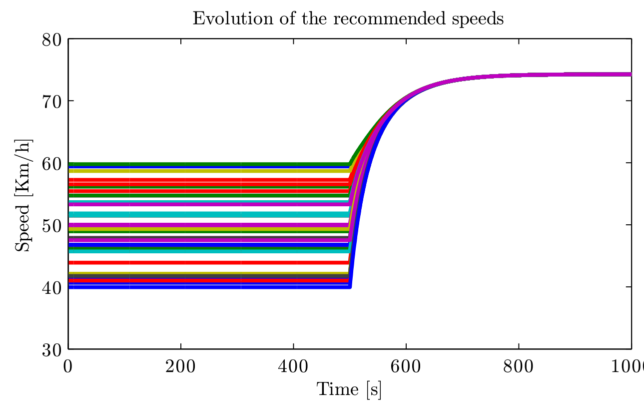

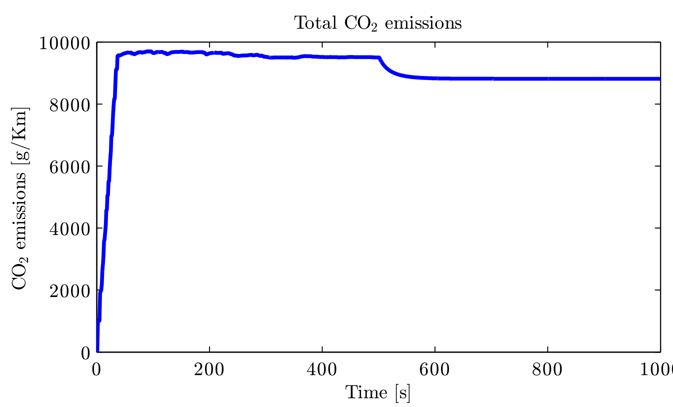

Fig. 3. Results of the SUMO simulation for the static case, before and after the activation of the algorithm at time step 500 s. Top: dynamic of the recommended speed for all the 40 vehicles; Bottom: total CO_2 emissions.

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.
- Vehicular flow at the beginning of L1: one new car every 2 seconds until simulation time 1300 s.
- Length of simulation: 2010 s (no more cars on L1 or L2).
- Window size for the calculation of the moving average (MA) of CO_2 emissions for visualisation purposes: 500 time steps.
- Initial speeds for L1 are randomly chosen with uniform distribution in 4 scenarios:
  - case 1, initial speeds in $(80, 100)$ [Km/h];
  - case 2, initial speeds in $(60, 80)$ [Km/h];
  - case 3, initial speeds in $(40, 60)$ [Km/h].

A sample of simulation results for case 3 are shown in Fig. 4 and Fig. 5.

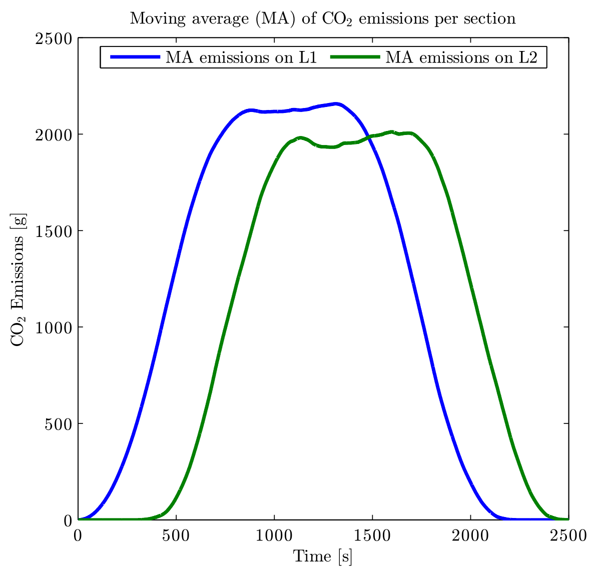

Fig. 4. Example for case 3: 500 time-step MA of CO_2 emissions for a set of initial speeds in $(40, 60)$ [km/h].

The figures reveal 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.

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.

Table II clearly demonstrate the benefits of the ISA.

Comment: Table II is consistent with our previous discussion. In particular, the proposed ISA system could reduce emissions most effectively in a scenario where most vehicles

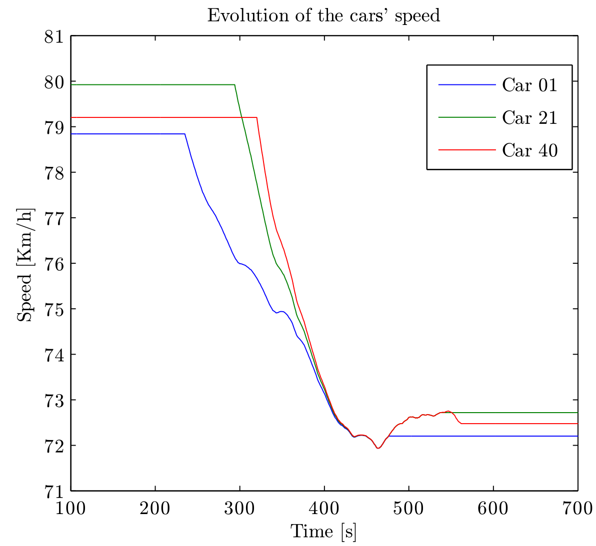

Fig. 5. Speed evolution of some vehicles from case 3, using the dynamic version of the algorithm.

TABLE II
GENERAL RESULTS: TOTAL EMISSIONS PER LANE.

	Total Emissions∗ [grams]					
	L1 (Uncontrolled)		L2 (Controlled)		Improvement	
Case	Mean	σ	Mean	σ	Mean	σ
1	2638711.8	1334.9	2586171.7	206.9	1.99%	0.05
2	2600747.9	628.3	2584162.7	167.5	0.64%	0.03
3	2787509.4	3934.1	2586801.9	156.2	7.20%	0.13

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

are required to speed up for speed coordination. In Ireland, for instance, this ISA system can be facilitated on a road with speed limit of 80 Km/h just after a road with a speed limit of 60 Km/h.

Finally, it is clearly the case that parameters of the algorithm have the potential to affect emission savings. For example, clearly the speed of convergence of the algorithm affects the rate of which the emissions are saved. To conclude this section, we now give a brief empirical exploration of the manner in which the algorithm parameters affect emissions. To do this, we adjust η and μ in the Algorithm 1 in their proper stable range of operation. Then, we evaluate the mean and standard deviation of the emission reduction in case 2 above, using a fixed set of speed distributions for five random sets of experiments. The result are depicted in Fig. 6 and Fig. 7. As expected, the choice of η and μ affect the practical performance of the algorithm.

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 [34], [35]. This emulation platform uses the open source road traffic simulator to emulate

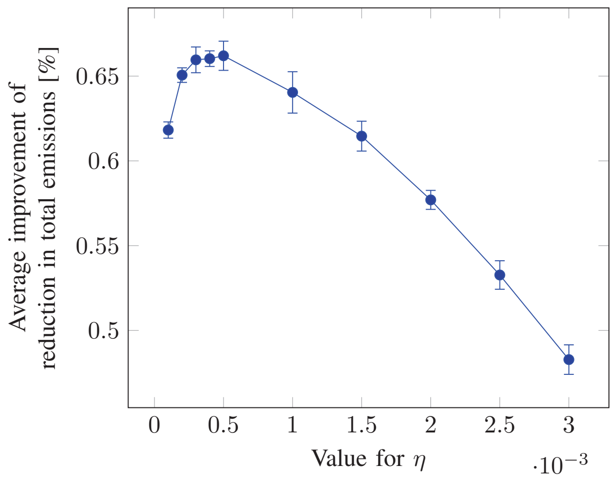

Fig. 6. Average improvement of reduction in total emissions for Case 2, in terms of η: mean and standard deviation for 5 different sets of experiments per set of values of η. Value of μ is fixed in 0.01.

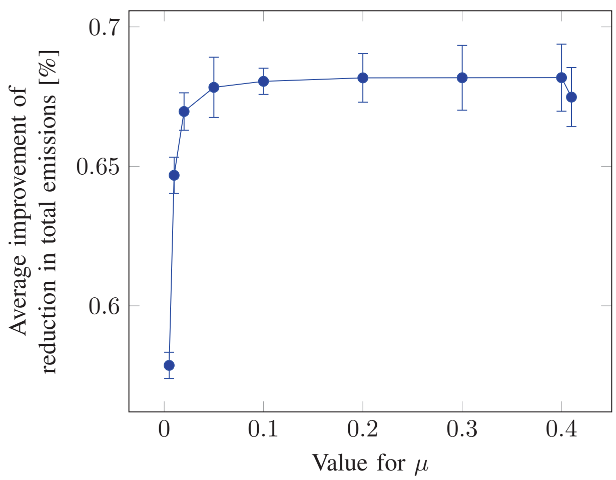

Fig. 7. Average improvement of reduction in total emissions for Case 2, in terms of μ: mean and standard deviation for 5 different sets of experiments per set of values of μ. Value of η is fixed in 0.001.

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[3]), and a OBD-II adaptor[4] to embed a real car into the simulation, as shown in Fig. 8. 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.

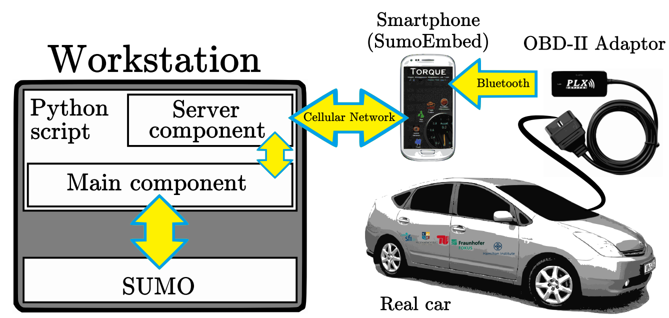

Fig. 8. Schematic of the HIL emulation platform.

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[5], the Prius' driver is allowed to either follow or ignore the speed recommendation (displayed on the smartphone's screen) as desired.

The HIL experiment is setup as follows.

- The total duration of the experiment is 600 s, of which the ISA algorithm is only engaged after 300 seconds;
- To emulate a disturbance to the algorithm we stop the Prius from the time instant 450 s to 470 s.
- 7 vehicles of each type Euro 1/2/3 and 9 vehicles of type Euro 4 are used, all of them with a maximum speed of 100 Km/h.
- The sampling time interval ΔT for collecting new information and updating the recommendations is 1 s.

[3]Torque Pro by Ian Hawkins. Available from Google Play: https://play.google.com/store/apps/details?id=org.prowl.torque. (Last accessed website on 23 July, 2013.)

[4]PLX Devices Inc., 440 Oakmead Parkway, Sunnyvale, CA 94085, USA. Phone: +1 (408) 7457591. Website: http://www.plxdevices.com. (Last accessed website on 23 July, 2013.)

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

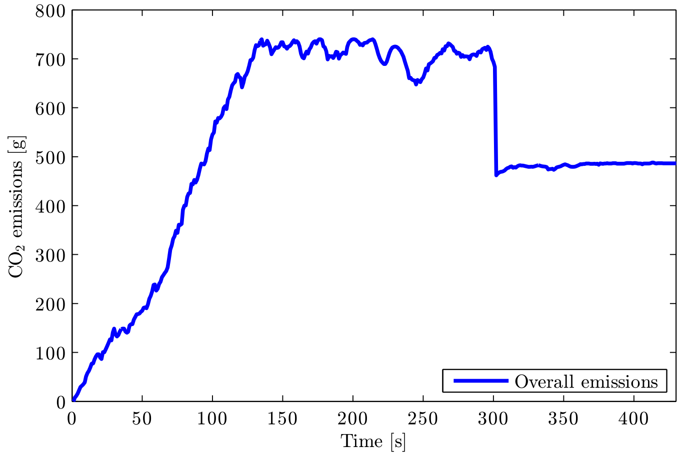

Fig. 9. Evolution of the overall CO_2 emissions.

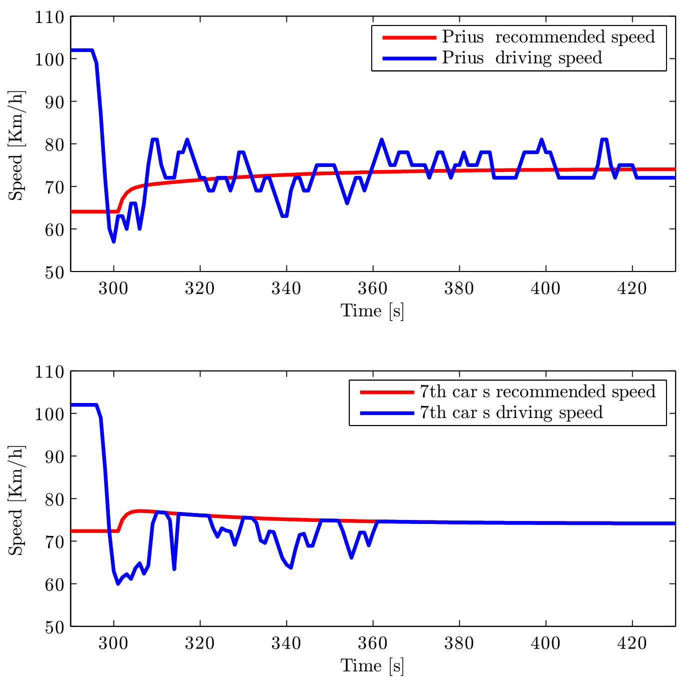

Fig. 10. Evolution of some variables along the whole HIL simulation. Top: related speeds for the Prius (the 6th car), and Bottom: the related speeds for the 7h car (just behind the Prius).

Results of the experiment are depicted in Fig.9 and Fig.10. First, Fig. 9 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. 9, we can observe that the overall CO_2 emissions reduce significantly once the the ISA algorithm is switched on, to an average of 484 g/Km.

In Fig. 10 (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.

V. 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.

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 [24].

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 [24].

Lemma 2 ([24]) *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 ([24]) *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

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

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

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

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

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] 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.

[9] 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.

[10] 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.

[11] 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.

[12] 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.

[13] 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.

[14] 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.

[15] 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.

[16] 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.

[17] 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.

[18] 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.

[19] 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.

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

[21] 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.

[22] 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.

[23] 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.

[24] 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.

[25] 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.

[26] 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.

[27] 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.

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

[29] 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.

[30] 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.

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

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

[33] 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.

[34] 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.

[35] 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.

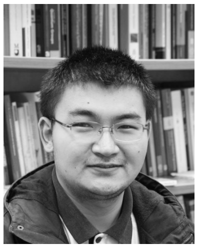

Mingming Liu received his B.E. degree in Electronic Engineering (first class honours) from National University of Ireland Maynooth (now Maynooth University) in 2011. He is currently working towards a Ph.D. degree in Applied Mathematics at the Hamilton Institute, Maynooth University. His current research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

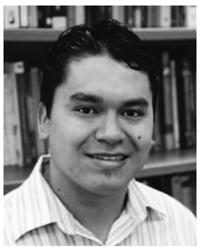

Rodrigo H. Ordóñez-Hurtado received his degree in Engineering in Industrial Automatica from the University of Cauca, Colombia, in 2005. He commenced his Ph.D. in Electrical Engineering at the University of Chile in March 2008, and between April-November 2012, he was an intern at the Hamilton Institute, 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.

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.