A framework for real-time emissions trading in large scale vehicle fleets

Florian Häusler[1,6], Mahsa Faizrahnemoon[2,6], Emanuele Crisostomi[3], Arieh Schlote[2], Ilja Radusch[4], Robert Shorten[5]

1. Fraunhofer FOKUS, Germany, Kaiserin-Augusta-Allee 31, 10589 Berlin, +49 3463 7117, florian.haeusler@fokus.fraunhofer.de
2. The Hamilton Institute, NUI Maynooth, Ireland
3. DESTEC, University of Pisa, Italy
4. Technical University of Berlin, Germany
5. IBM Research, Ireland

In this paper a framework for the real-time trading of budgeted emission rights between a fleet of participating vehicles is presented. The trading problem is formulated as a utility maximization or as a utility fairness problem, which can be solved in real-time either in a centralized or in a distributed manner. In both cases, we illustrate the basic issues that arise when such a framework is realized in practice, and we show the efficacy of the approaches by providing several simulation examples and a realistic case study.

1. Introduction

Within the context of smart-cities research [1-2] collaborative mobility is rapidly becoming one of the main growth areas in fields such as: computer science; electrical engineering; applied mathematics; operations research; networking, as well as transportation engineering. Motivated by the fact that road networks in our cities are often hugely inefficient, and by recent advances in technologies such as communication networks, vehicles, optimization, the provision of products related to smart mobility is expected to develop into a multi-billion dollar industry over the next decade [3-4].

Cooperative mobility is not a new topic in transportation research. The notion of vehicles communicating with other vehicles and with infrastructure (V2X) to improve the quality of service (QoS) to the vehicle owner is closely related to the area of transport telematics [5]. Transport telematics has been an active area of research for several decades but its impact has always been limited by available technologies. However, due to the aforementioned advances, and strict new guidelines from regulatory bodies, the interest in collaborative mobility has never been greater. Roughly speaking, transport related regulation is driven by four main factors; congestion on our roads (and related inefficiencies); vehicular safety; greenhouse gas emissions; and the quality of air in our cities. Each of these issues makes a compelling case for investing in smarter transportation systems. Congestion is not only unpleasant to experience, but it is also a major inhibitor of economic growth [6]. Similarly, car-related safety issues, both for occupants and other road users, are also a major issue [7]. Also, in the EU, transportation accounts for 20% of greenhouse gas emissions. However, perhaps the most compelling reason to revisit our concept of personal mobility, is due to the fact that the internal combustion engine (ICE) is extremely damaging to human health [8]. By-products of the ICE include: CO; NOx; SO; Ozone; Benzene; PM10; and PM25 [9].

Before proceeding, it is worth noting that governments and municipal authorities have already started to respond to the air-quality issue. Cities in some countries ban certain vehicles from densely populated areas (low emission zones), or adapt speed limits to respond to pollution peaks. Going further in this direction, strict measures are being planned to ban the ICE from

[6] Florian Häusler and Mahsa Faizrahnemoon are joint first authors.

A short version of this paper was presented at the 9th ITS European Congress, Dublin, Ireland, 2013.

our cities in the near future. Degradation of air quality results from the aggregate effect of vehicles. But today's instruments as the ones introduced before are per-vehicle measures. Enforcing per-vehicle measures (unless we ban vehicles all together), takes no account of this effect. Second, contemporary measures are open-loop measures. The regulation is the same irrespective if there is one vehicle in a spatial area (in the middle of the night), or if there are millions of vehicles in the same area~~. In this context it makes little sense restricting vehicle access when there is little pollution (one vehicle), but it makes great sense when there are many vehicles as critical levels of pollution are due to aggregated effect.~~ Much of the current literature addressing traffic related air pollution focuses on shifting mobility away from dirty modes of transportation. In this context an effort is made to improve the perceived and actual QoS of electric vehicles and bike sharing systems [42-43]. In a similar manner intermodality is advertised as a way to simultaneously decrease inefficiencies in different transport modes [44]. Another focus is on developing automated driving systems that take the sometimes erratic and wasteful human behavior out of the picture [45].

The aggregate effect of vehicles in a geographic region has been recently addressed in [41]. We extend the approach by introducing a method to implement a perfect fairness, such that all vehicles – considering individual utility functions – experience the same utility from the system. But we also broaden the perspective and include the optimization of a vehicle fleet's emission, disregarding their geographic positioning. Here, the basic assumption is that a fleet operator wishes to control the aggregate effect of a group of vehicles (e.g., pollution), so that some (other) utility function of interest for the fleet operator is maximized. Note that many problems in Intelligent Transportation Systems (ITS) and logistics research can be formulated in this manner. As a practical example, dispatching companies such as DHL, Schenker, or UPS, are usually requested to monitor and control their fleets in regard to their ecological impact. Such companies, do already provide information about emissions associated with single trips [34]. Moreover, there are products from such companies that offer customers the possibility to choose whether they are interested in shipping their good either minimizing delivery time or total emissions. Another example arises when policy makers wish to determine a speed limit for a road segment, such that pollution levels are below or at a certain level, and subject to this constraint we wish to minimize fuel consumption. Advances in V2V technologies, communications, and in distributed optimization, allow us to implement such a system in real time without the need for significant centralized processing, and to respond to changing vehicle density levels. We may expect that the demand for such algorithms to increase as OEMs (and cities) will be asked to provide information about the environmental impact of products throughout their lifecycle [35] (and in the case of municipalities, air quality). Furthermore, customers of such services often pay for different levels of service. Such a degree of freedom can be usefully adopted to make packets requiring a high level of service "borrow" emissions from the other packets, with the objective of maintaining the overall emissions per fleet constant [36]. In this paper, we illustrate some algorithms that can be utilized for this purpose.

To summarize, our main objective in this paper is to present a method to regulate the aggregate effect (e.g., pollution) of a fleet of vehicles, so that the utility (benefit) to the vehicle owners is at all times maximized. In our particular application examples, we show how a fleet of hybrid vehicles has a total budget of CO_2 they can emit. Thus, in this example we control the aggregate emission of CO_2. Clearly, there are many ways to achieve this, and this non-uniqueness creates an opportunity to regulate and optimize simultaneously [9-12]. The present setting offers the opportunity of not only regulating, say, aggregate CO_2, but of doing this in a manner that another objective (e.g., maximization of QoS, minimization of CO or PM production) is achieved at the same time.

Specifically, in this manuscript: (i) we formulate the emissions problem, (ii) we discuss communication constraints, (iii) we present a hybrid algorithm, and (iv) we show some simulations. Our paper follows closely [9] but goes beyond this work by developing algorithms that have much faster convergence rates. The paper is structured as follows. In Section II we shall discuss the typical constraints that arise in cooperative mobility settings. These constraints guide the type of algorithms that can be deployed in such environments.

In Section III we formally state our problem, and in Section IV we will discuss algorithms for solving utility maximization problems in vehicular environments. Simulation results are presented in Section V, where we also develop the case study given in [41] to include utility fairness constraints.

2. Background

We consider the following problem:

$$max \sum_{i=1}^{N} f_i(D_i) \tag{1}$$

subject to the constraint that

$$\sum_{i=1}^{N} D_i = C \tag{2}$$

where the functions $f_i(D_i)$ map the scalar D_i to some utility scalar. As we have discussed, many problems in an ITS context map to this scenario. In our context C is a pollution budget assigned to a group of vehicles, and D_i is the pollution budget assigned to vehicle i, and $f_i(D_i)$ is the benefit of the i'th vehicle using the allocation D_i. A rich repertoire of techniques exist to deal with such problems and many of these have been successfully applied in other application domains [23]. Indeed, given reasonable assumptions on the $f_i(D_i)$ a centralized optimization is always possible. Nevertheless, despite the availability of such techniques, solving problems that arise in an ITS context are very challenging due to a number of factors. We now summarise some of the basic issues considering the specific ITS context.

1. **Scale**: ITS problems tend to be characterized by a large scale level of participation. Optimization problems over very large graphs are considered to be some of the most challenging in the theory of Operations Research [13].
2. **Non-homogeneous levels of vehicle participation**: ITS problems are also characterized by varying level of vehicle participation as vehicles leave and join the optimization problem.

Algorithms deployed to solve the maximization problem must also respect constraints imposed by the communication capabilities of the ITS system.

3. **I2V broadcast:** Here vehicles listen to a broadcast from a centralized node. This scenario, the Infrastructure-to-vehicle broadcast, is probably the most common and well-established communication link in our context. The required technologies to implement I2V broadcasts are already available in most vehicles and a feature of broadcasting is, that the total communication effort is scalable and does not depend on the number of vehicles in the fleet.
4. **I2V unicast:** Direct (IP) communication from the central infrastructure (e.g., a fleet management back-office) to individual users using cellular networks is a feature, which could not be implemented easily today, due to network address translation (NAT) traversal problems. Available technologies for NAT traversal require extra infrastructure and communication protocols or control over network infrastructure.

5. **V2I unicast:** In contrast to unicast I2V, communicating into the other direction is technically easy to implement. A standard cellular communication network is required and the communication effort increases linearly with the number of vehicles.
6. **Asymmetric and symmetric V2V:** Vehicle-to-vehicle communication using IEEE 802.11p technology is currently under standardization but not yet deployed. V2V communication may be classified as being symmetric and asymmetric, where symmetry refers to a situation where vehicle *i* communicates with vehicle *j* if and only if *j* can communicate with *i*.

A further issue that arises in ITS applications is the ability to solve large scale optimization algorithms. For example, in some situations it may not be possible or desirable to solve optimization problems in a centralized manner.

7. **Computation:** A basic question concerns as to whether the central infrastructure is capable of solving large scale optimization problems and communicating bidirectionally to mobile vehicles at relatively high sampling rates.

Finally, a number of very important performance constraints also arise in ITS applications. These include the following:

8. **Robustness to failure and uncertainty:** A basic requirement is that the optimization algorithm should be robust with respect to the effect of failures in the system. Also, in many situations, utility functions are not known exactly.
9. **Speed of convergence:** A further requirement, in the context of decentralized algorithms, is that algorithms should converge to the optimal solution reasonably quickly.
10. **Scalability:** Finally, deployed algorithms should function satisfactorily irrespective of level of vehicle participation.

Given the above considerations, the question now arises as to the type of solution that is suitable for ITS applications. Roughly speaking, three choices are available to us:

1. **Fully centralized solutions:** involve solving the optimization problem defined in Equation (1) and Equation (2) centrally. Clearly, the requirements to implement such a solution involve perfect knowledge of the $f_i(D_i)$, bi-directional communication between infrastructure and all vehicles, and significant computational ability of the centralized node. Advantages of the centralized approach include speed of convergence (instantaneous) and that there is no V2V communication requirement. Disadvantages of this approach include sensitivity to node failure, sensitivity of solutions to noise and uncertainty, and the feasibility of centralized approaches in large scale environments.
2. **Fully decentralized solutions:** involve solving the optimization problem defined in Equations (1-2) using local computation and communication only. Clearly, such an approach involves a V2V communication capability requirement, and advantages include that there is no need for a centralized optimization step, the ability to deal with uncertainty in the $f_i(D_i)$, and the robustness with respect to node failures. Disadvantages include the very slow convergence rates - especially in large scale applications.
3. **Hybrid solutions:** involve using a small amount of centralized computation/communication to speed up the convergence of decentralized solutions and to mitigate some of the constraints of ITS environments. In the next section we shall present two such algorithms.

2.1 Related work

Before proceeding it is worth noting that similar problems have been addressed elsewhere in the literature, as well as in a cooperative mobility context; see [9, 10-12, 14-15] where some related ideas are realized in the context of hybrid vehicles. Cooperative sensing, control and optimization, are also closely related to many of the most established techniques in the control systems community that have been applied to similar problems. Amongst these related areas, distributed control and optimization is perhaps the most relevant one; see [16-18] and the references therein. Other related work includes consensus, flocking and multi-agent systems [19-20]. Distributed resource allocation is also very relevant with the work [21-22] particularly interesting, as [15, 23-26], on implicit consensus. Our proposed work will build heavily on this last one, and the more recent [12, 41]. Specific examples of other related work include [27] where a fully centralized system was implemented to allocate a (limited) number of parking spaces to users respecting their individual (multidimensional) preferences. Drivers shared their utilities with the central management center, which maximized the sum of individual utilities by allocating well-suited parking spaces to them. Similarly, in [23] a fully distributed approach was introduced to distribute load on CPU clusters with the objective to lower aggregate power consumption and thus to lower costs. CPUs are able to adjust speed based on demand and performance constraints.

The above discussion presents only a very brief overview of related work. Nevertheless, the need for new algorithmic developments is highlighted by noting the constraints under which each of the studies was conducted. For example, the classical distributed control techniques found in the literature are unable to deal with the massive scale encountered in a cooperative mobility setting. Similarly, the work by Stanojevic assumes symmetric communication channels, and this is clearly very restrictive in a cooperative mobility setting. This assumption is later removed by Knorn, but at the cost of very slow algorithmic convergence. Finally, the work by Geng [27] assumes massive centralized computational power and perfect knowledge of the agents utility when computing the utility maximization step. Our objective is to now develop algorithms which alleviate all or most of these assumptions.

3. Problem Statement and Assumptions

Given a network of N nodes (vehicles) with individual utility functions $f_i(D_i)$, and a constraint $\sum_{i=1}^{N} D_i = C$, find an allocation D_i that solves $max \sum_{i=1}^{N} f_i\left(D_i\right)$ subject to $\sum_{i=1}^{N} D_i = C$.

Comment: In some applications we may be interested in solving a minimization problem.

3.1 Assumptions

The following assumptions appear to be a reasonable reflection of current and near future state-of-the-art in ITS systems, and suggest a hybrid solution whereby the optimization is solved using decentralized computation with aid of limited signaling from the centre.

- Utility functions $f_i(D_i)$ are concave. If we are interested in minimizing the sum of utility functions (e.g., minimum CO), then we assume that the utility functions are convex. In addition, the precise nature of the $f_i(D_i)$ may be unknown, but the vehicles have access to a measurement or a-posteriori calculation based on observable variables of $f_i(D_i)$ at each time step. In some cases, e.g. the utility fairness case, we will also assume that the utility functions are increasing (decreasing) in the maximization (minimization) case.

- We assume that an optimal and feasible solution exists. For instance, this corresponds to assuming that vehicles have the ability to achieve the optimal QoS or to produce the optimal quantity of CO emissions.

- We assume that all vehicles can communicate directly (or indirectly through a measurement) with a centralized infrastructure.

- We assume that the centralized infrastructure may broadcast information to the vehicles, but not communicate directly with each vehicle.

3.2 Solution

Equations (1) and (2) define a classical optimization problem whose solution may be found with the aid of Lagrange multipliers. Let $H(D_1,...,D_n,\lambda)$ be the Lagrangian associated with the problem, i.e.

$$H(D_1,...,D_n,\lambda) = \sum_{i=1}^{N} f_i(D_i) - \lambda(\sum_{i=1}^{N} D_i - C). \tag{3}$$

The optimal solution of the maximization problem is given by finding the allocation D_i^* such that $\lambda^* = \partial f_i\left(D_i^*\right)/\partial D_i, \forall i=1,...,N$, subject to the linear constraint being satisfied. The form of this equation suggests two (hybrid) approaches to solve the maximization problem. First, it suggests that the maximization problem can be solved using feedback in combination with local on-vehicle computational power. Namely, by regulating the outputs of the vehicle functions in such a way that they follow the λ, which is broadcasted from a centralized infrastructure. The task of the central infrastructure then becomes to find the appropriate multiplier λ so that the utility maximization problem can be solved. As we shall see, this can be solved by embedding the multiplier as part of a feedback loop as is done in subgradient methods. This approach requires no V2V or I2V unicast capabilities, can deal with uncertainty in the utility functions, and requires no centralized computation. A disadvantage of the approach is that it is not robust to failure of the centralized node and selecting the control gains can be an issue. A second approach is based on the fact that optimality is also achieved when each of the vehicles achieve consensus of the values of the functions $g_i(D_i) = \partial f_i(D_i)/\partial D_i$. This observation was first exploited in [23] in the context of symmetric communication graphs, and later extended to the asymmetric case in [15], and to the case where the outputs of the vehicle are adjusted as part of a feedback loop. Advantages of these approaches are that no I2V unicast capabilities are required, and the techniques can deal with uncertainty in the utility functions. In addition no centralized computation is required, and the technique is robust with respect to node failure.

4 Algorithms

4.1 Algorithm I. Feedback control and local computation

The basic idea in Algorithm 1 is to use an integral control to asymptotically find the Lagrange multiplier λ*. For this purpose, the centralized infrastructure updates iteratively according to:

$$\bar{\lambda}(k+1) = \bar{\lambda}(k) - \mu\left(C - \sum_{i=1}^{N} D_i(k)\right), \tag{4}$$

where we have assumed that the centralized infrastructure either receives communication of the $D_i\left(k\right)$ from each vehicle, or measures it. Each node (vehicle) receives a broadcast of $\bar{\lambda}(k)$ and updates its utility (implicitly) according to:

$$D_i(k+1) = D_i(k) - \beta_i\left(\bar{\lambda}(k) - \partial f_i(D_i(k))/\partial D_i(k)\right). \tag{5}$$

Roughly speaking, this corresponds to a classic subgradient method from optimization theory, see [28-29] and the gains can be selected using methods from nonlinear control theory. The convergence of the algorithm, and the choice of the model parameters is beyond the scope of this paper. Inappropriate choice of gains will lead to oscillations which is a problem for some applications (leading to breakage of mechanical systems) **[GIVE CITATION]**. In otherwords, the system converges to a ball in the neighbourhood of the optimal point. In our case oscillations are not a problem; rather they lead to a reduction in performance [28]. The convergence argument and choice of gains follows from standard systems theory ideas [16, 23, 33] and appropriate assumptions on the nonlinearities. Typically, one assumes sector bounds for the nonlinear functions. Note that Algorithm 1 is a variant on a standard optimisation algorithm, the dual subgradient descent algorithm, where we replaced the primal optimisation step with a second ascend type update. This kind of algorithms is sometimes denoted as saddle point algorithms, see [29].

4.2 Algorithm II. Implicit consensus with input

As previously mentioned, it was observed in [23] that utility maximization problems can be formulated as a consensus problem where all agents achieve the same value of the derivative $g_i(D_i) = \partial f_i(D_i)/\partial D_i$ of their own utility function. Note that this solution is particularly attractive as it can be realized in a fully decentralized manner. However, consensus based solutions usually require symmetric communication graphs, to guarantee that $\sum_{i=1}^{N} D_i(k) = C$ for all *k*, and might suffer from other convergence related issues. Some of these issues are addressed in [15] in which a basic feedback error signal, $C - \sum_{i=1}^{N} D_i(k)$, from the infrastructure is used to yield a modified integral control of the following form:

$$D_i(k+1) = D_i(k) + \varepsilon \sum_{j\in\aleph_i} \left(g_i(D_i(k)) - g_j(D_j(k)) \right) + \gamma \left(C - \sum_{q=1}^{N} D_q(k) \right) \tag{6}$$

in which $\aleph_i$ is the set of nodes that send information to node *i*. Note that if node *i* has N_i neighbors that can send information to node *i*, then Equation (6) can be rewritten as

$$D_i(k+1) = D_i(k) + \varepsilon \left(N_i g_i(D_i(k)) - \sum_{j\in\aleph_i} g_j(D_j(k)) \right) + \gamma \left(C - \sum_{q=1}^{N} D_q(k) \right) \tag{7}$$

Thus, by denoting $\varepsilon_i = \varepsilon \cdot N_i$,the problem can be stated in a consensus-like form, as:

$$D_i(k+1) = D_i(k) + \varepsilon_i \left(g_i(D_i(k) - \frac{1}{N} \sum_{j\in\aleph_i} g_j(D_j(k)) \right) + \gamma \left(C - \sum_{q=1}^{N} g_j(D_j(k)) \right) \tag{8}$$

where *j* ranges over all neighbours of *i*. Roughly speaking, the first term of the update equation (8) is used to achieve a consensus around the optimal value, while the second term ensures that the global equality constraint is satisfied.

As before, the stability and convergence properties of Algorithm II, and the choice of γ and ε can again be determined using ideas from systems theory (see [25]). While a proper discussion of this issue is again beyond the scope of the present paper we note that Algorithm 2 requires V2V communication, and has slow convergence properties [9]. As we outline extensions or different version of this approach, we refer to this algorithm in the following as algorithm 2.a.

Comment: The sign of the parameter ε in Equations (6-8) depends on the choice of the functions *g*. In particular, the sign will be positive if functions *g* are decreasing, and negative if functions *g* are increasing. Note that different policies can be implemented by replacing functions $g_i(D_i) = \partial f_i(D_i)/\partial D_i$ in Equation (8) with other appropriate functions. For example, if we use $g_i(D_i)=f_i(D_i)$ instead of the partial derivatives, a utility fairness policy will be implemented. A discussion on the properties of the candidate functions g_i is beyond the purposes of this paper, but conditions on the functions g_i and a related discussion can be found in [15] and [23].

Modified algorithm: One of our contributions in this paper is to present a modification to Algorithm II (Algorithm II.a) to speed up convergence, which can be outlined as: "At each time step the vehicles communicate the maximum and minimum values of $\partial f_i\left(D_i(k)\right)/\partial D_i(k)$ to the central infrastructure. These values are then broadcasted to all nodes who then augment their neighborhood information to encorporate the new values". Simulation results show that speed of convergence is indeed dramatically increased, and in the next paragraph we give a plausibility argument of why this occurs. We refer to this modified algorithm as algorithm II.b The completely distributed version without any feedback from the infrastructure with an update function shown in Equation 9 is referred as Algorithm II.c.

$$D_i(k+1) = D_i(k) + \varepsilon_i \left(g_i(D_i(k) - \frac{1}{N} \sum_{j \in \aleph_i} g_j(D_j(k)) \right) \quad (9)$$

4.3 Some details of Algorithm 2

As simulation examples show, the implicit consensus algorithm appears to reach convergence in a faster manner when the maximum and the minimum value of the second term in Equation (8) are also revealed to the other nodes. This result is not totally unexpected. It is well known that in undirected graphs, the second smallest eigenvalue of the graph Laplacian, also called *algebraic connectivity*, quantifies the speed of convergence of consensus algorithms, see for instance [37]. Also, it is well known that in this case the algebraic connectivity of a new graph G_2, $a(G_2)$, obtained from adding a new edge to graph G_1 is characterised by:

$$a(G_1) \le a(G_2) \le a(G_1) + 2 \quad (10)$$

thus establishing that adding a new edge can potentially increase the speed of convergence, or at least does not decrease it [38]. Note that having two more nodes broadcasting their own value, as suggested from the improvement to Equation (8), corresponds in practice to adding up to $2 \cdot (n-1)$ new edges.

The case of directed graphs is notoriously more difficult to handle, and there is not even a unique definition of the graph Laplacian. As a consequence, there are no theorems similar in spirit to Equation (8), at least to the knowledge of the authors, however the following result from [39]

$$a(G) \le min_{v \ne w} \{ d_0(v) + d_0(w) \} \quad (11)$$

where v and w are two nodes of the graph G, can be used to state that increasing the outdegree d_o of some nodes of the graph (as in practice occurs when the algorithm is modified to circulate the minimum and the maximum values) does also increase (or at least, does not decrease) an upper bound of the algebraic connectivity of the graph.

Finally, we also mention that similar strategies to improve the convergence of network consensus algorithms in directed graphs were also devised in [40], where the authors claim that having as many as possible vertices with the maximum out-degree of $n-1$, and having the in-degree of each vertex around m/n, where m is the number of edges and n that of nodes, can improve the convergence speed.

5 Simulations

5.1 Utility minimization

In the first simulation, we assume that a fleet of 20 hybrid vehicles has a total budget of CO_2 they can emit, and the objective is to minimize the total quantity of emitted *PM*. Furthermore, as in [15], we assume that vehicles can regulate the quantity of emissions by deciding which fraction of the desired speed is supplied by the ICE or by the electric motor. This is consistent with a power split hybrid vehicle. In our simulations, we compute emissions using a simple "average speed model", and data taken from [30] set in an urban setting where vehicle speeds vary between *10* and *50 Km/h*. Then, according to the data from [30] (vehicle code R019, corresponding to petrol LDVs with an ICE equivalent of class EURO4, data can be found on page 164 and 106 for CO_2 and *PM* respectively), one can approximate the relationship between *PM* and CO_2 as

$$PM\left(CO_2\right) = 50CO_2^{\,2} - 19CO_2 + 6.4\,, \qquad (12)$$

where *PM* is expressed in mg/km and CO_2 in kg/km, which is a good approximation for speeds included between 10 and 50 Km/h. Our objective is thus $\min \sum_{i=1}^{N} f_i\left(CO_2^{(i)}\right)$ subject to $\sum_{i=1}^{N} CO_2^{(i)} = C$, where C is the overall CO_2 budget allocated to the fleet of vehicles, and $CO_2^{(i)}$ is the amount of CO_2 allocated to the i'th vehicle. Note that this problem is equivalent to the one described by Equations (1-2) which has already been discussed in the previous sections. We assume that the vehicles have slightly different parameters, due to characteristics specific of the particular vehicle (e.g., brand, age), and we use the parameter values $\gamma = 0.1$ and $\varepsilon = 0.01$. Figure 1 shows how the CO_2 budget is shared among the fleet of vehicles in order to minimize *PM* emissions using Algorithm II.a.

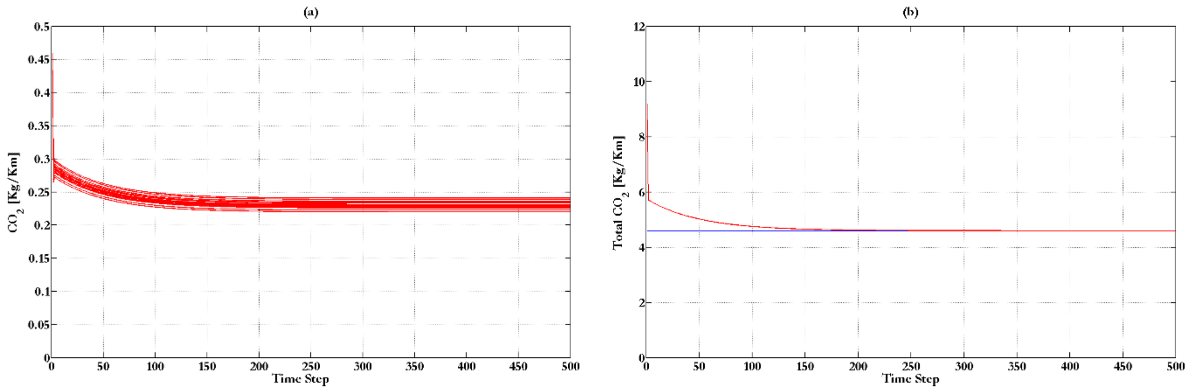

Fig. 1. Figure 1.a shows how the CO_2 budget is shared among the fleet of vehicles in order to minimize PM emissions. At the same time, the CO_2 budget is asymptotically respected, Figure 1.b.

5.2 Utility fairness

The objective of the previous simulation was to minimize the production of *PM* while keeping the total amount of CO_2 equal to a desired threshold. However, such an optimal solution might be unfair in some circumstances. In fact, in some situations it may be of interest to make all vehicles produce an equal amount of pollution. In many countries road users pay a flat fee for road useage irrespective of the damage that they are doing to the environment (road tax). The idea here is to make users share the available budget fairly, and to implicitly make those cars that are dirty pay more by forcing them to use more of their electrical energy budget. In this situation dirty vehicles pay more by using more electrical vehicle energy (thus they are forced to recharge earlier). This idea is easily accommodated in our framework. In the next simulation we are therefore interested in making the vehicles produce the same quantity of *PM*, while still respecting the total CO_2 budget. While such a solution is clearly not overall optimal for the environment (i.e., more *PM* than before is overall produced), it provides a solution that is fairer for the hybrid vehicle owners.

We now compare the three algorithms discussed in Section IV (algorithm 2.a., 2.b and 2.c), and for simplicity, we assume that the utility function of node *i* is given by a function of the form $f_i = \alpha_i \log D_i$, which is increasing and concave. The choice of this function is motivated by dispatch type problems where economic utility is an increasing function of allocated emissions. Also, we initially consider $g_i\left(D_i(k)\right) = f_i\left(D_i(k)\right)$.

Figure 2 compares the performance of Algorithm I (Figure 2.b), Algorithm 2.c (Figure 2.b), and Algorithm 2.b (Figure 2.d) in a simulation example involving 41 overall vehicles, with a communication graph shown in Figure 2.a. In the described example, parameters ε and γ of Equation (6) in the distributed algorithms are chosen as *0.002* and *0.0024* respectively. In the case of the centralised algorithm, μ in Equation (4) is chosen as *0.01*, while β_i of Equation (5) is chosen equal to *0.02*. We further assume that the budget C is initially equally divided between the vehicles. In the shown example, we assume that vehicles are able to communicate along the roads, while buildings do not allow vehicles travelling along different roads to communicate with each other. Note the increased convergence speed in Figure 2.d over that in Figure 2.b.

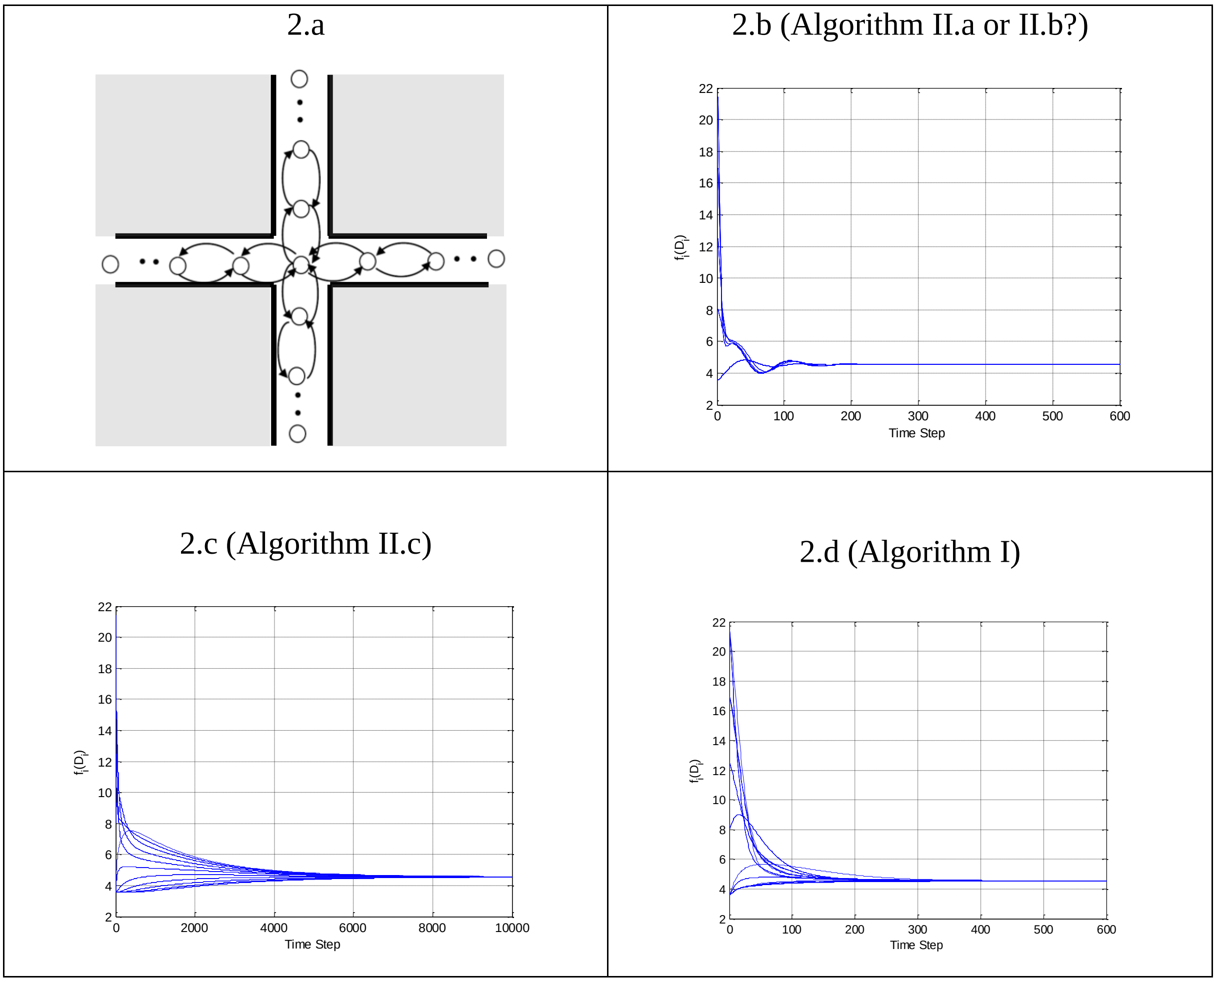

Fig. 2: The CO_2 budget is shared in such a way to equalize the production of *PM*. In Fig. 2.d the infrastructure communicates the value of λ(k) as explained in Algorithm I, and the vehicles do not communicate among themselves; in Fig. 2.c there is no I2V communication (Algorithm II.c), and vehicles can communicate only with the neighbouring vehicles, see equation (9); in 2.b vehicles further receive some more information from the infrastructure as explained in Algorithm II.b (or is it Algorithm II.a???). Note that this last opportunity algorithm improves the time of convergence of the hybrid algorithm with respect to II.d.

However, making all vehicles produce the same amount of pollution might not be itself a fair choice in some circumstances. For instance, vehicles that travel longer distances should be allowed to pollute more. In this case, we take a different fairness approach, and choose $g_i(D_i) = \partial f_i(D_i)/\partial D_i$, as discussed in Section 3.2 and 4.2. Accordingly, Figure 3 compares again Algorithms I, II.a, II.b and II.c in Figures 3.a, 3.b and 3.c respectively. The multipliers ε and γ of Equations (4), (5) and (6) are the same of Figure 2.

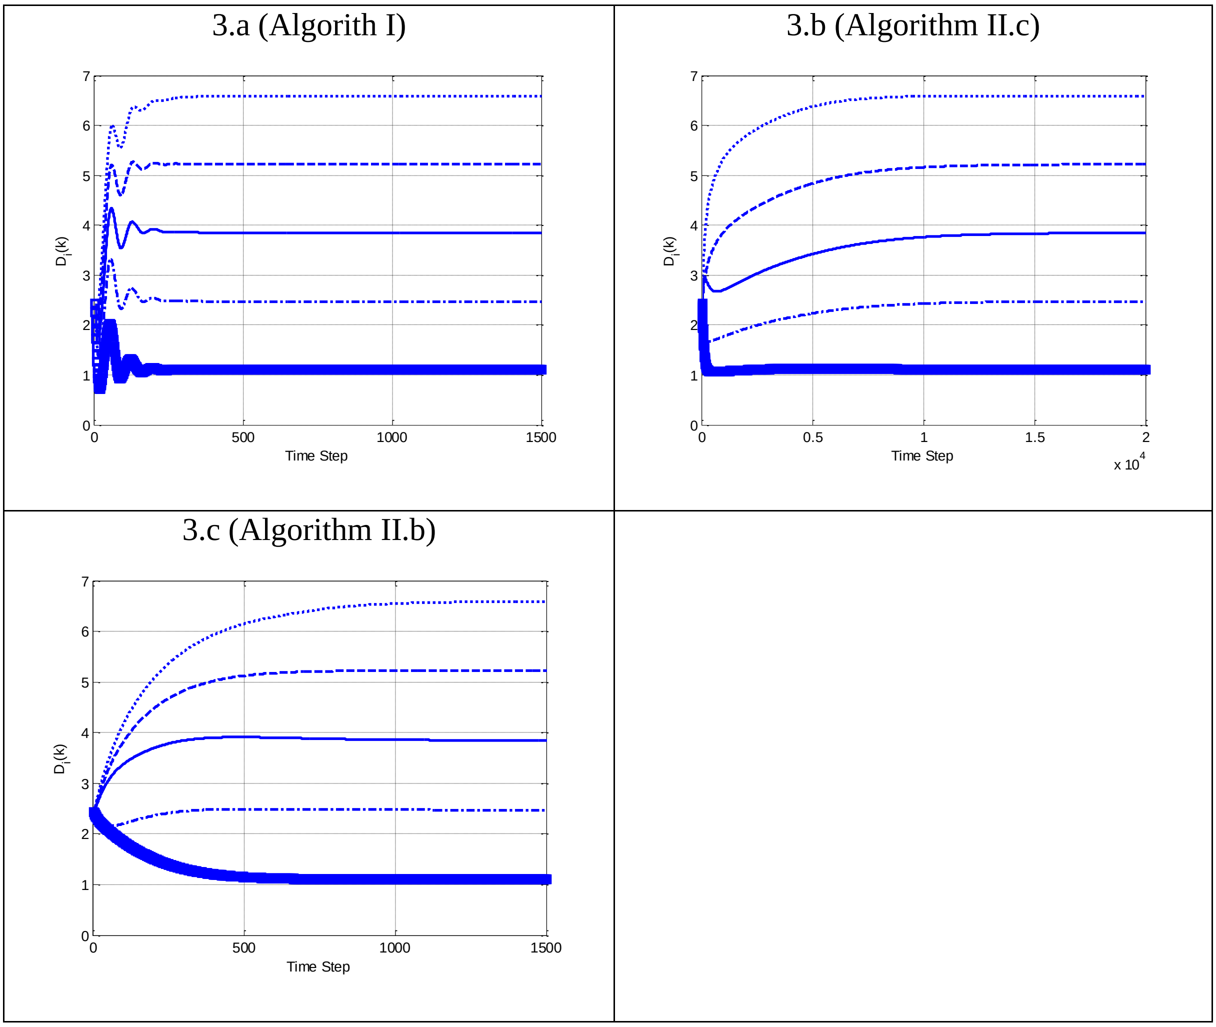

Fig. 3: The time evolution of the pollution of some of the vehicles is depicted in the case we are interested in equalizing the derivatives of the utility functions. In 3.a the infrastructure communicates the value of λ(k), and the vehicles do not communicate among themselves (Algorithm I); in 3.b there is no I2V communication, and vehicles can communicate only with the neighbouring vehicles (Algorithm II.c); in Fig. 3.c vehicles further receive some more information from the infrastructure as explained in the Algorithm II.b, and convergence is faster than in 3.b.

5.3 Utility fairness: a case study

In this simulation we present a more realistic application of the utility fairness paradigm through a case study, which is an extension of a previous work within the so-called TwinLIN initiative [41]. In order to control the *CO* pollution in urban areas, the TwinLIN work implemented a central control approach, which forced hybrid vehicles to travel either in electric or ICE mode in order to maintain the control variable below, or close to, a desired level. The objective was achieved by assuming that the infrastructure could broadcast a probability p_{ev} to all vehicles. This probability was then used locally by the vehicles, and possibly modified to achieve other objectives at the same time, to decide whether they should travel in electric or ICE mode. To support the TwinLIN idea, a proof-of-concept vehicle had been constructed, by customizing a 2008 Toyota Prius, so that it is possible to influence the drive mode choice, by accessing the Toyota CAN-bus via Smartphone.

We now extend the TwinLIN framework to further implement the utility fairness strategy, and test our ideas using the mobility simulator SUMO, with its HBEFA-emission model and TRACI interface [32]. We simulate traffic in an arbitrary area of Berlin, shown in Figure 4, which we divided into nine equivalent rectangular zones of the same size (in a 3x3 grid).

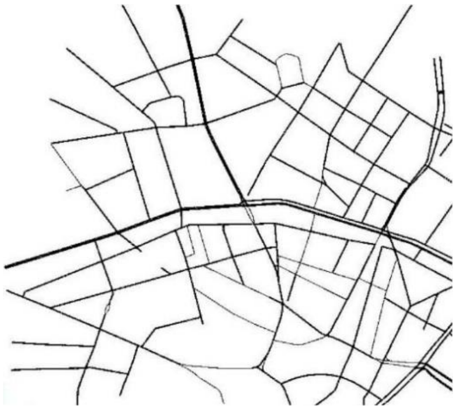

Fig. 4. Road network imported from OpenStreetMap [31], area around Rosenthaler Platz in Berlin, Germany

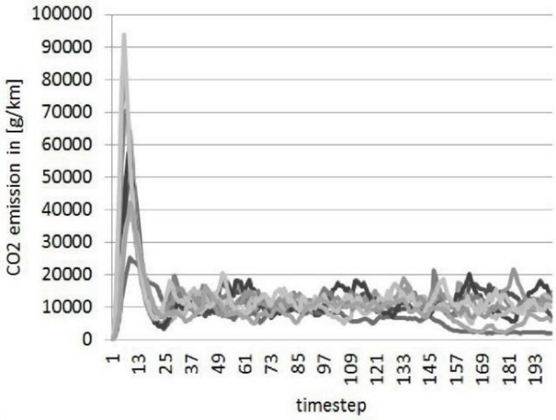

Fig. 5. The level of CO_2 is maintained equal to the CO_2 budget in each subzone of the Berlin area of interest.

Then, we apply a simplified version of Algorithm 2 to every single sub-area of the Berlin region of interest. In practice, we assume that there is a central infrastructure in each sub-area that broadcasts the feedback error signal $C - \sum_{i=1}^{N} D_i(k)$, in order to maintain the level of CO_2 equal to the desired one. At the same time, each vehicle in the sub-area uses the error signal to further adjust the proportion of speed generated from the ICE in order to minimise the overall quantity of emitted CO, according to the algorithm the Equation (8). SUMO simulation results are given in Figure 5, which shows the level of CO_2 in each sub-zone of the Berlin area of interest, and in Figure 6 which shows the normalised CO produced by vehicles belonging to 5 different classes of vehicles (with ICEs equivalent to EURO 1 to 5 vehicles). Overall, we have approximately the same level of CO_2 in each sub-zone of Berlin, and we have that each car produces the same quantity of CO, independently from its specific vehicle class.

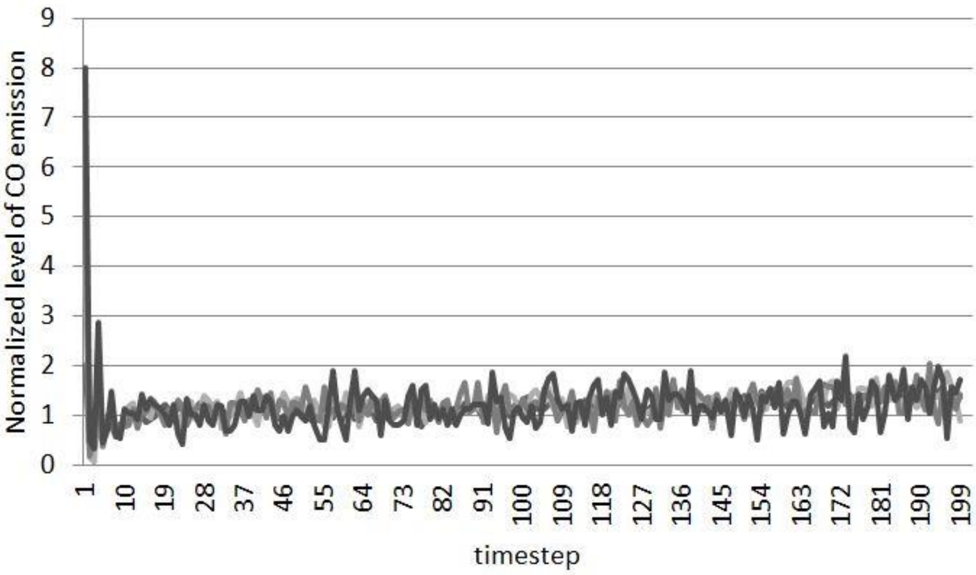

Fig. 6. Vehicles belonging to different classes produce the same quantity of CO.

6 Conclusions

Following a detailed discussion about typical constraints in the design of agorithms to address topics in intelligent transportation systems, a framework for real-time trading of emission rights between a fleet of vehicles has been presented in this paper. Building on a recently proposed cooperative emission control approach, a number of algorihms are presented and discussed for distributed emissions trading in a fleet of vehciles. It is worth noting that these algorithms operate directly on the engine of the vehicles and do not require behavioural change of the driver. Future work will involve hardware in the loop testing to validate the concepts and to test the proposed algorithms.

Acknowledgement
The authors would like to thank the (anonymous) reviewers whose remarks contributed to improve the quality and the clarity of the final version of this article. This work was in part supported by Science Foundation Ireland grant 11/PI/1177.

References

1. IBM (2007). *Delivering smarter transportation systems,* [online]. Available from URL: http://www-935.ibm.com/services/us/igs/pdf/transport-systems-white-paper.pdf [Accessed: 18 April 2013].
2. Hodgkinson, S. (2011). *Is your city smart enough*, [online]. Available from URL: http://www.cisco.com/web/strategy/docs/Is_your_city_smart_enough-Ovum_Analyst_Insights.pdf [Accessed: 20 August 2013].
3. Ezell, S. (2010). *Intelligent Transportation Systems,* [online]. Available from URL: http://www.itif.org/files/2010-1-27-ITS_Leadership.pdf [Accessed: 20 August 2013].
4. Woodburn, A., Allen, J., Leonardi, J. (2010). “International Road and Rail Freight Transport: The Impact of Globalisation on Activity Levels”, in *OECD, Globalisation, Transport and the Environment*, OECD Publishing, Paris.
5. Chatterjee, K., Mcdonald M.; Paulley N.; Taylor N. B. (1999). “Modelling the impacts of transport telematics: Current limitations and future developments”, *Transport Reviews*, Vol. 19, Issue 1, pp. 57-80.
6. Dirks, S., Constantin Gurdgiev, C. and Keeling, M. (2010). *Smarter Cities for Smarter Growth*, IBM Institute for business value, White Paper.
7. Peden, M., McGee, K. and Sharma, G. (2002). *The Injury Chart Book: A Graphical Overview of the Global Burden of Injuries*, World Health Organization, Geneva.
8. Guerreiro, C., de Leeuw, F. and Foltescu, V. (2013). “Air quality in Europe - 2013 report”, European Environment Agency, Copenhagen, Denmark.
9. Knorn, F., Corless, M.J., Shorten, R. (2011). “A result on implicit consensus with application to emissions control”, *Proceedings of 50th IEEE Conference on Decision and Control and European Control Conference (CDC-ECC).*
10. Stüdli, S., Crisostomi, E., Middleton, R. and Shorten, R. (2012). “A flexible distributed framework for realising electric and plug-in hybrid vehicle charging policies”, *International Journal of Control*, vol. 85, no. 8, pp. 1130–1145.
11. Stüdli, S., Crisostomi, E., Middleton, R. and Shorten, R. (2012). “AIMD-like algorithms for charging electric and plug-in hybrid vehicles”, *IEEE International Conference on Electric Vehicles*, Greenville.

12. Schlote, A., Häusler, F., Hecker, T., Bergmann, A., Crisostomi, E., Radusch, I., Shorten R. (2012). "Cooperative regulation of emissions using plug-in hybrid vehicles", *ACM/IEEE/IFAC/TRB International Conference on Connected Vehicles and Expo*, Beijing, China.

13. Hager, W., Hearn, D., Pardolas, P., (1994). *Large Scale Optimization: State of the Art*, Kluwer Academic.

14. Shorten, R., Wirth, F., and Leith, D. (2006). "A positive system model of TCP like congestion control: asymptotic results", *IEEE/ACM Transactions on networking*, vol. 14, no. 3.

15. Knorn, F. (2011). *Topics in Cooperative Control*, PhD Thesis, Hamilton Institute, NUI Maynooth.

16. Siljak, D. (2007). *Large-Scale Dynamic Systems: stability and structure*, North-Holland.

17. Lestas, I.,Vinnicombe, G. (2010). "Heterogeneity and scalability in group agreement protocols: beyond small gain and passivity approaches", *Automatica*, 46 (7). pp. 1141-1151.

18. Moylan, P.J., Hill, D.J. (1978). "Stability criteria for large-scale systems", *IEEE Trans Automat Control*, AC-23 (2) (Special Issue on Large Scale Systems), pp. 143-149.

19. Borkar, V., Varaiya, P.P. (1982). "Asymptotic Agreement in Distributed Estimation", *IEEE Transactions on Automatic Control*, vol. 27, pp. 650-655.

20. Jadbabaie, A., Lin, J., and Morse, A.S. (2003). "Coordination of groups of mobile autonomous agents using nearest neighbor rules", *IEEE Transactions on Automatic Control*, vol. 48, pp. 988-1001.

21. Srikant, R. (2003). *The mathematics of Internet Congestion Control,* Systems & Control: Foundations & Applications, Birkhaueser.

22. Kelly, F.P. (1999). "Mathematical modelling of the Internet", in *Proceedings of ICIAM 99, 4th International Congress of Industrial Applied Mathematics*, Edinburgh, UK.

23. Stanojevic, R. and Shorten, R. (2009). "Load balancing vs. distributed rate limiting: an unifying framework for cloud control", in *Proceedings of ICC*.

24. Stanojevic, R., Shorten, R. (2008). "Fully decentralised emulation of best-effort and processor sharing queues", in *Proceedings of ACM Sigmetrics*.

25. Knorn, F., Corless, M. and Shorten, R. (2011). "Results in cooperative control and implicit consensus", *International Journal of Control*.

26. Knorn, F., Stanojevic, R., Corless, M., and Shorten, R. (2009). "A Framework for Decentralised Feedback Connectivity Control with Application to Sensor Networks", *International Journal of Control*, 82, Issue 11, 2095-2114.

27. Geng, Y., Cassandras C.G. (2011). "Dynamic resource allocation in urban settings: A "smart parking" approach", *Proceedings of CACSD.*

28. Boyd, S., Vandenberghe, L. (2004), *Convex Optimization,* Cambridge University Press.

29. Nedić, A., Ozdaglar, A. (2009), "Subgradient methods for saddle-point problems", *Journal of Optimization Theory and Applications*, vol. 142, no. 1, pp. 205–228.

30. Boulter, P.G., Barlow, T.J. and McCrae, I.S. *Emission factors 2009: Report 3 – exhaust emission factors for road vehicles in the United Kingdom*. Project report PPR356, TRL Limited, 2009.

31. Open Street Map, [online]. Available from URL: http://www.openstreetmap.org/ [Accessed: 20 August 2013].

32. Krajzewicz, D., Bonert, M. and Wagner, P. (2006). “The open source traffic simulation package SUMO”, *RoboCup 2006 Infrastructure Simulation Competition*, Germany.

33. Khalil, H. (2001). *Nonlinear Systems*, Prentice Hall.

34. D. P. DHL. “Go green”, available online at url: http://www.dhl.com/en/logistics/freight_transportation/go_green.html, [Accessed: 20 August 2013].

35. Golden, J.S., Dooley, K.J., Anderies, J.M., Thompson, B.H., Gereffi, G., and Pratson, L. (2011). “Sustainable product indexing: navigating the challenge of ecolabeling,” *Ecology and Society*, vol. 15, no. 3.

36. Haase, H. (2011). “Logistics emission trading system for green optimization,” final report, available online at url [in German]: http://www.emissionshandel-letsgo.ovgu.de/letsgo_media/Downloads/Schlussbericht_LETSGO_ILM2012.pdf, [Accessed: 20 August 2013].

37. Olfati-Saber, R., Fax, J.A., and Murray, R.M. (2007). “Consensus and Cooperation in Networked Multi-Agent Systems,” *Proceedings of the IEEE*, vol. 95, no. 1.

38. de Abreu, N.M.M. (2007). “Old and new results on algebraic connectivity of graphs,” *Linear Algebra and its Applications*, vol. 423, pp. 53-73.

39. Wu, C.W. (2005). “Algebraic connectivity of directed graphs,” *Linear and Multilinear Algebra*, vol. 53, no. 3, pp. 203-223.

40. Cao, M., and Wu, C.W. (2007). “Topology design for fast convergence of network consensus algorithms,” *IEEE International Symposium on Circuits and Systems*.

41. Schlote, A., Häusler, F., Hecker, T., Bergmann, A., Crisostomi, E., Radusch, I., and Shorten R. (2013). “Cooperative Regulation and Trading of Emissions Using Plug-in Hybrid Vehicles,” *IEEE Transactions on Intelligent Transportation Systems*, vol. 14, no. 4, pp. 1572-1585, 2013.

42. de Weerdt, M., Gerding, E., Stein, S., Robu, V. and Jennings, N. (2013) “Intention-aware routing to minimise delays at electric vehicle charging stations”, in Proceedings of the 23rd International Joint Conference on Artificial Intelligence, pp. 83-89.

43. Chen, B., Pinelli, F., Sinn, M., Botea, A. and Calabrese, F. (2013). “Uncertainty in urban mobility: Predicting waiting times for shared bicycles and parking lots”, in Proceedings of the 16th International IEEE Annual Conference on Intelligent Transportation Systems, The Hague, Netherlands.

44. Botea, A., Berlingerio, M. and Nikolova, E. (2013). “Multi-modal journey planning in the presence of uncertainty”, in Proceedings of the International Conference on Automated Planning and Scheduling, pp.20-28.

45. Göhring, D., and Latotzky, D., Wang, Miao. and Rojas, R. (2013). “Semi-autonomous Car Control Using Brain Computer Interfaces”, Intelligent Autonomous Systems 12, Advances in Intelligent Systems and Computing, Springer Berlin Heidelberg.