

Author Posting.

© The Institute of Electrical and Electronics Engineers, Inc. (IEEE), 2011.

This is the author's version of the work.
It is posted here by permission of the IEEE for personal use, not for redistribution.

The definitive version is published in the proceedings of the
50TH IEEE CONFERENCE ON DECISION AND CONTROL & EUROPEAN CONTROL CONFERENCE
12–15 December 2011, Orlando, FL, USA.

A Result On Implicit Consensus with Application to Emissions Control

Florian KNORN, Martin J. CORLESS, and Robert N. SHORTEN

Abstract—This paper is concerned with a class of decentralised control problems that arise in contemporary applications where agents cooperate to control and regulate a global quantity, are limited in the manner in which they communicate with each other, and are required to reach consensus on some implicit variable (for instance, CO₂ emissions). An algorithm is presented for achieving this goal. A simplified application of the algorithm to emissions control for a fleet of Plug-in Hybrid Electric Vehicles (PHEVs) is given.

I. INTRODUCTION

Recent survey papers such as [1, 2] give a good overview of the large interest in coordination, cooperation and consensus in multi-agent systems. While the field appears mature, most of the work focuses on so-called *unconstrained* and *explicit* consensus problems. The former refers to situations where the network is to agree on a common value / state, but the actual value of the consensus is not important. The consensus problem is explicit if the agents can directly manipulate the value of interest that consensus is to be reached on. However, in many applications, neither assumption can be made. Oftentimes an aggregate measure of system behaviour (assumed to depend on *all* the network's states) is of interest as well, and it is not always possible to directly change the quantity of interest in each agent. There is thus not only a desire for cooperation among the agents in order for the network to achieve a common goal (“cooperative control”), but, whatever protocol is considered, it must not assume that the agents can directly manipulate the consensus value.

In the present paper we review recent cooperative control algorithms that allow a network to reach consensus on one quantity of interest, while additionally meeting a separate, global goal. The main contributions of this paper is a theorem coupled with an application which consists of a cooperative control scheme designed for city-wide CO₂ emissions control with Plug-in Hybrid Electric Vehicles (PHEVs). It is based on results of a recently accepted paper [3], which we reproduce here in a condensed form.

The remainder of the paper is structured as follows: Next, we will review related work in the area. We then introduce the problem setting more concretely and define some necessary notation. This is followed by a presentation of an

algorithm that can achieve the desired goal. An application of our results to emissions control for a fleet of hybrid vehicles is given, followed by a closing discussion in the last section.

Previous work

For a good introduction into the field and examples of its many diverse applications see for instance the surveys [1, 2]. Consensus problems have been formulated in many different ways, often dictated by particular applications. Just to illustrate a few: Variations include whether or not the topology of the communication network changes over time; whether or not it is undirected or directed; whether or not the agents can manipulate the state on which to reach consensus directly or not (and instantly or only with finitely fast dynamics); whether the consensus value is scalar or multidimensional; whether or not there are delays in the information exchange; and whether states are updated in a synchronous or asynchronous fashion; *etc.* Initial work in the field [4–8] assumed bi-directional information exchange between neighbouring nodes (which leads to undirected communication graphs). Extensions to directed communication graphs can be found in [9, 10]. Further generalisations allowed asynchronous consensus protocols [9, 11], changing graph topologies [8, 12, 13], linear agent dynamics [12, 13], or implicit consensus [14]. This latter paper inspired much of the present work.

For unconstrained consensus problems, three basic approaches appear to exist [15]: leader-following [16, 17], virtual structure based [18, 19] or behaviour based [20, 21] approaches. The first concept is a common technique used in formations of autonomous mobile agents that are to follow a prescribed trajectory. The second treats the entire network of agents as a single entity with the desired behaviour, relative to which each member then adjusts its own behaviour. The behavioural approach makes use of weighted sums of several desired behaviours (such as goal seeking, formation keeping and obstacle avoidance). A different approach is taken in [22] where constrained consensus is formulated as an optimisation problem.

It is in this third class that our work should be placed, as the desired behaviour of the nodes in our network is also a combination of a localised constraint (consensus on the utility values) and a global constraint.

II. OVERALL SETTING

A. Problem setting

To facilitate the discussion, let us introduce our general set-up and some necessary terminology. The overall multi-agent system is assumed to consist of individual agents

This work was supported in part by the Irish Higher Education Authority (HEA) Programme HEA/PRTL15 TGI and Science Foundation Ireland Principal Investigator (PI) Award 07/IN.1/1901.

Florian Knorn (florian@knorn.org) and Robert N. Shorten (robert.shorten@nuim.ie) are with the Hamilton Institute, National University of Ireland, Maynooth, Ireland.

Martin J. Corless (corless@purdue.edu) is with the School of Aeronautics & Astronautics, Purdue University, West Lafayette, IN, USA.

or nodes that are linked together via some form of local, broadcast based communication. In other words, we assume, loosely speaking, that only communication among “nearby” nodes is possible, and that this communication is not necessarily bi-directional (i.e. node A may send information to node B , but may not necessarily receive any information back). Each agent in the network has a *physical state* (or just “state”) that describes an actual property of the agent (e.g. a car’s driving speed) that the agent can change. Additionally, to each node is also associated a *utility value* which depends directly on the physical state (e.g. CO₂ emission, which, in a first approximation, can be assumed to depend on the car’s driving speed). This value represents the quantity of interest on which consensus is to be reached. The dependency between physical state and utility value is given by the *utility function*, and these functions can vary between agents. This implies that when consensus is reached on the utility values, the agents’ states may not necessarily be equal.

Additionally, we define a *global value* that depends directly on *all* the nodes’ physical states through the so-called *global function*. By suitable means of communication (or decentralised estimation) either all or just some nodes in the network have access to this global value.¹ Finally, we assume that the nodes use the inter-agent communication system to share their current utility value with neighbouring nodes. This set-up is illustrated in Figure 1.

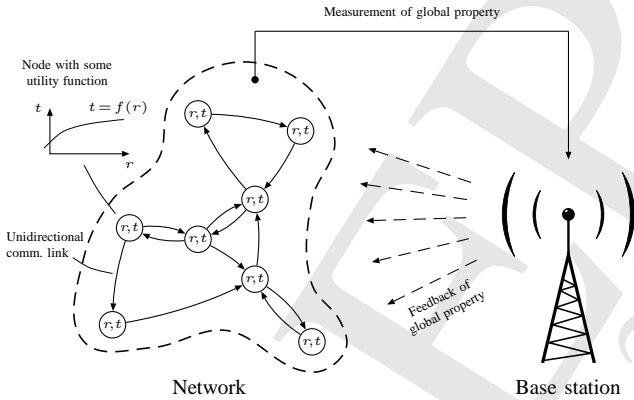


Fig. 1. Illustration of our basic setting.

B. Problem statement

The objective is for all agents in the system to reach consensus on their utility values, while also driving a global value to a target value. This should be achieved in a decentralised way, using simple algorithms that will operate in a variety of settings, including time-varying topologies of the communication network, non-linear utility functions that are only known approximately, when not all nodes have access to the global value and when the state updates are not necessarily performed synchronously.

¹That is, either it can be measured / estimated locally by the nodes, or it will be communicated to them by some “external” broadcast (for instance sent from a base station that itself can estimate or measure that value).

In the following, we present a recursive update rule for each node’s physical state that solves this problem. The basic idea behind the algorithm is to use classic consensus techniques together with an additional term that takes into account the global objective.

C. Notation

Throughout, we use superscripts in parentheses to denote individual elements in a vector, and subscripts to denote the time index. Vectors and matrices are typeset in bold letters to improve readability.

We represent the communication network topology at each time step $k = 0, 1, \dots$ as a directed graph where each agent forms a node, and the (directed) communication links the arcs. The set of (in-)neighbours of node i at time k is called $\mathcal{N}_k^{(i)}$; it contains all the nodes j that can send information to node i (including itself). The network is called *connected* (in the literature also referred to as *strongly connected*) if there exists a path from every node to every other node in the network, respecting the orientation of the arcs.

A matrix $\mathbf{P} \in \mathbb{R}^{n \times n}$ is called *row-stochastic* if all its entries are non-negative and all its row-sums equal one, in other words $p^{(ij)} \geq 0$ and $\mathbf{P}\mathbf{1} = \mathbf{1}$, where $\mathbf{1}$ denotes the all-ones vector of appropriate dimensions. Similarly, *row sub-stochastic* matrices are defined to be real valued, non-negative matrices whose row-sums are less than or equal to one (but with at least one row-sum strictly less than one). A *strictly row sub-stochastic* matrix is a row sub-stochastic matrix where *all* row-sums are strictly less than one.

Let $r_k^{(i)} \in \mathbb{R}$ be the *physical state* of node i at time k , so that \mathbf{r}_k forms the state vector of the network. Node i ’s *utility value* $t^{(i)} \in \mathbb{R}$ depends on the physical state via the *utility function* $f^{(i)} : \mathbb{R} \rightarrow \mathbb{R}$, that is $t_k^{(i)} = f^{(i)}(r_k^{(i)})$. Furthermore, let $g : \mathbb{R}^n \rightarrow \mathbb{R}$ be the *global function* that depends on all the states $r^{(i)}$. Desired values are denoted with subscript asterisks, so that, for example, the desired value for the global function is denoted by g_* . Based on this desired value, the solution to the problem thus consists of a vector \mathbf{r}_* for which $f^{(i)}(r_*^{(i)}) = t_*$ for all $i = 1, \dots, n$ and $g(\mathbf{r}_*) = g_*$.

D. Growth conditions and feasibility

We require that the utility functions be continuous and strictly increasing with a rate that is bounded away from zero and upper bounded; the global function must also be continuous and strictly increasing coordinate-wise with non-zero but also upper bounded rates. Formally (using \mathbf{e}_i to denote the i th unit vector of appropriate dimension) there must be positive constants $\underline{d}^{(i)}, \bar{d}^{(i)}, \underline{h}^{(i)}, \bar{h}^{(i)}$ such that for all $i = 1, \dots, n$

$$\underline{d}^{(i)} \leq \frac{f^{(i)}(r_a) - f^{(i)}(r_b)}{r_a - r_b} \leq \bar{d}^{(i)}$$

for all $r_a, r_b \in \mathbb{R}$ with $r_a \neq r_b$, and

$$\underline{h}^{(i)} \leq \frac{g(\mathbf{r} + \Delta r \mathbf{e}_i) - g(\mathbf{r})}{\Delta r} \leq \bar{h}^{(i)}$$

for all $\mathbf{r} \in \mathbb{R}^n$ and all $\Delta r \in \mathbb{R}$ with $\Delta r \neq 0$.

These conditions guarantee that the continuous utility functions are strictly monotone increasing and unbounded and thus have continuous inverse functions $\phi^{(i)}$ so that $\phi^{(i)}(f^{(i)}(r)) = r$ and $f^{(i)}(\phi^{(i)}(t)) = t$ for all $t, r \in \mathbb{R}$. Given these conditions, it can easily be shown that the problem always has a unique feasible solution, [3].

III. AN ALGORITHM TO ACHIEVE IMPLICIT, CONSTRAINED CONSENSUS

In this section we now present an algorithm that allows a network to achieve implicit consensus and to cooperatively control a global goal. Its implementation only requires knowledge of the *bounds* on the growth rates of the utility functions as well as the global function, but not the functions themselves.

Comment Different assumptions to the ones presented in the following theorem are possible. For instance, there may be situations where the utility functions are perfectly known. In that case, they can be inverted and the algorithm is essentially reduced to the problem discussed in [23]. Conversely, the utility functions may not be of static but dynamical nature. In this case, the utility values may be filtered values of the physical state, a set-up that is also considered in [3]. Due to space limitations, we cannot discuss these results here.

Theorem [3] Consider the standard situation as described in the Notation section and assume that the utility functions $f^{(i)}$ and the global function g are continuous and satisfy the growth condition. For any initial condition $\mathbf{r}_0 \in \mathbb{R}^n$ and any sequence of strongly connected communication graphs, suppose that the nodes iteratively update their physical states based on

$$r_{k+1}^{(i)} = r_k^{(i)} + \sum_{j \in \mathcal{N}_k^{(i)}} \eta_k^{(ij)} (t_k^{(j)} - t_k^{(i)}) + \mu_k^{(i)} \sigma_k \quad (1)$$

where $\sigma_k = g_* - g(\mathbf{r}_{k+1-M})$ if $k+1$ is a multiple of $n-1$, or $\sigma_k = 0$ otherwise. Furthermore, suppose there exist constants $\varepsilon_1, \varepsilon_2, \underline{\mu}, \bar{\mu} > 0$ such that for each $i = 1, \dots, n$, $j \in \mathcal{N}_k^{(i)}$ and $k \geq 0$ all the gains $\eta_k^{(ij)}$ satisfy

$$\sum_{j \in \mathcal{N}_k^{(i)}} \eta_k^{(ij)} \leq 1/\bar{d}^{(i)} - \varepsilon_2 \quad \text{with} \quad \eta_k^{(ij)} \geq \varepsilon_1 \quad (2)$$

and $0 < \underline{\mu} \leq \mu_k^{(i)} \leq \bar{\mu}$.

Then, if $\bar{\mu}$ is sufficiently small, the state vector \mathbf{r}_k converges asymptotically to the vector \mathbf{r}_* for which $f^{(i)}(r_*^{(i)}) = t_*$ for all i and $g(\mathbf{r}_*) = g_*$. •

Comment The update law (1) can roughly be interpreted as a controller that contains two terms: One that produces the consensus on the utility values and one that adjusts the state so as to drive the global value to the desired level. In each time step, each node i only incorporates information from neighbouring nodes $\mathcal{N}_k^{(i)}$ to reach consensus on the utility values; access to the global value is only required every $n-1$ time steps.

Extensions

The algorithm can easily be extended to cases where not all nodes have access to the global value (it can be shown that it is sufficient for only one node to have access to this value) as well cases of asynchronous communications and state updates, see [3] for more details.

We would now like to demonstrate the use of the above algorithm by applying it to a proposed emissions control scheme for a fleet of plug-in hybrid electric vehicles.

IV. APPLICATION TO EMISSIONS CONTROL WITH PLUG-IN HYBRID ELECTRIC VEHICLES

A. Introduction

Reducing greenhouse gas emissions as well as emissions of directly harmful gasses and particulates are one of the major challenges of the future. In the European Union for instance, see [24], attempts to reduce emissions include schemes to encourage optimum driver behaviour (emissions reducing driving style for instance), more efficient use of the transport network (traffic management and smart navigation systems to reduce congestion, dedicated lanes for specific vehicle types, real-time information systems for locations of available parking spaces, etc.), or to modify the transport demand (improved logistics to reduce commercial traffic, better public transport, more low-polluting vehicles, etc.).

Attempts by large cities like London [25] or Berlin [26] to reduce emissions have received much public attention, particularly due to the direct impact they have on the public's mobility. They try to either discourage drivers from taking their cars into the city centre by charging a significant fee for doing so, or by strictly only allowing (certified) low-polluting vehicles to enter. While these attempts indeed succeed in somewhat diminishing the number of vehicles in the typically congested city centres, they basically are open-loop schemes that do not use feedback to respond to the actual situation. Factors like the weather, the time of day, day of the week, or public holidays all have a significant impact on air quality and green house gas emissions. Another problem is that although cars become greener and greener, there are more and more cars in circulation so that the effects of more efficient and less polluting engines is offset by the ever growing number of cars, [25].

Research and development in the field of electric vehicles has progressed significantly in recent years. *Hybrid electric vehicles* (HEV), which combine a conventional internal combustion engine (ICE) based propulsion system with an electric engine, were introduced to the mass market around the early 2000s, and, apart from their economic advantages in terms of fuel economy and their "green appeal", a number of additional factors have led to fast growing sales, [27]. Just to name a few, strong tax incentives in most country make a compelling argument for these low-emission vehicles; social preferences and awareness for environmental quality or energy security have increased; fuel prices can rise and already have risen sharply in the past, with a consistent

upward trend over time; etc. However, consumer adopting rates could still be improved upon, [28].

A new generation of hybrid vehicles are so-called *Plug-in Hybrid Vehicles* (PHEV). These cars have a much larger battery than traditional hybrids and are designed to be charged not only while driving (through regenerative braking for instance), but more importantly by means of “plugging” into an external power supply when the car is parked, [29]. At the current state of the art, this allows the car to drive several tens of kilometres purely on electric power, hence producing zero *local* emissions. The electrical energy still needs to be produced somewhere, either in a “clean” way (such as wind, solar, water or nuclear power based) or “dirty” way (traditional coal or oil based power plants). But while the latter also pollute the air and produce greenhouse gases, the overall emissions and harmful particulates may be filtered more effectively and, since power plants are usually located far away from urbanised zones, their pollution does not accumulate in the cities as is the case with traditional, fossil fuel based transport. Thus, the air quality in densely populated areas — which pose major health concerns, [30] — will be improved either way.

Unfortunately, market adoption of PHEVs is still somewhat slow, mainly due to economical and technical reasons with the batteries. It appears that battery technology still needs to improve in order for this class of vehicle to be economically viable, and very few vehicles currently can drive farther than 100km in purely electrical mode (a figure drastically reduced in cold weather conditions). For that reason, the combustion engine serves mainly as a range extender, allowing the car to run several *hundreds* of kilometres — but at the expense of local air pollution.

B. Controlling emissions, maximising driving distance

Hybrid electric vehicles offer new possibilities for urban mobility. For the first time, cars can be truly context-aware. In principle, it is possible to combine GPS and engine management unit to enable vehicles to choose where to be most polluting. For example, it makes eminent sense for a hybrid vehicle to switch to full electric mode in the neighbourhood of schools or hospitals. In the following application we explore, at a very high level, a fleet-wide notion of such context awareness. We wish to, in a manner that is fair, adjust the behaviour of the hybrid electric vehicles such that city-wide pollution and/or emissions are regulated. Before proceeding, we give a few words on hybrid electric vehicle fundamentals.

Hybrid vehicles come in several power-train configurations. In the *parallel* power-train configuration, a combustion engine works in conjunction with a small electric motor to provide extra torque, or, particularly in the case of plug-in hybrids, to extend the driving range. The two methods of propulsion can either run exclusively or in conjunction (“blended mode”). In other words, it is possible to “mix” the power sources and vary between emission-free, all-electric mode (short driving range) or emission-producing combustion-based mode (large range).

In the following, we propose a scheme to manage this trade-off in order to cooperatively regulate CO₂ emissions² in a fleet of n vehicles, while maximising their overall driving range. For that, we make a number of assumptions: (1) The participating PHEVs have a parallel power-train configuration that allows arbitrary blending between the power output of the combustion engine and the electric motor. (2) This power mixing can be described by a convex combination, for which we define the blending-parameter $r^{(i)} \in [0, 1]$ for each car i . By convention, let $r^{(i)} = 0$ if the car is in all-electric mode, and $r^{(i)} = 1$ if the car is only propelled by the combustion engine. (3) The vehicles are equipped with some broadcast-based vehicle-to-vehicle communication system (such as the proposed 802.11p protocol for *Co-operative Awareness Messages*, [31]) that allows each car to broadcast its current emission level to other cars in the area. The emissions need not be measured in real-time but could be derived from offline measurements, taking into consideration the currently used power blend. (4) Information about the aggregate CO₂ emissions are available to each car. They could either be measured externally and broadcast to the fleet (through the *Traffic Management Channel* for instance, [32]), or the cars could collectively estimate them through some distributed averaging scheme, [23]. (5) The emissions control scheme should be fair in the sense that no car should be allowed to have higher emissions than others.

C. Simulations

Given these assumptions, this set-up can easily be cast into the framework presented earlier and the algorithms proposed in Section III can be applied: The utility functions here are linear functions mapping the interval $[0, 1]$ of the blending-parameter $r^{(i)}$ (the “physical state”) into the corresponding range of emissions $t^{(i)}$ that vary between 0 (when in emissions-free all-electric mode) and $\bar{t}^{(i)}$, the nominal CO₂ emissions of the combustion engine. In order to satisfy the fairness requirement, emissions between the different agents must be equalised. The global function is simply the sum of all the CO₂ emissions. Given the linear / multi-linear nature of these functions, it is then trivial to derive the required growth-rates in order to calculate suitable gains $\eta_k^{(ij)}$ and $\mu_k^{(i)}$.

We now present three simulation runs of this set-up. For that, we generated fleets of $n = 4$ as well as $n = 50$ cars whose emissions are realistically distributed between the different emissions classes. In each time step, the topology of the communication graph was changed randomly (but so as to always guarantee strong connectedness). For each simulation run, the agents were initialised to use a 50/50 power mix, that is $r_{k=0}^{(i)} = 0.5$ for each $i = 1, \dots, n$. From then on, the blending-parameter was modified iteratively based on the update law presented earlier. In each case, the desired aggregate emissions level g_* was set to be 25% lower than that at time $k = 0$, thus requiring all the cars in the

² Note that we use CO₂ emissions here only as an example — our scheme can easily be applied to any other type of emissions.

network to adjust their energy mix in order to reduce overall emissions by 25%.

The gains $\mu_k^{(i)}$ required in the update law were set both according to the (only sufficient) stability conditions above, as well as manually: As commented on earlier and discussed in [3], the bounds on $\mu_k^{(i)}$ required by Theorem 1 are very conservative and may result in slow convergence to the desired global value. However, they may be increased significantly without compromising stability, which allowed us to achieve significantly faster convergence rates. (The gains $\eta_k^{(ij)}$ were always set according to the theorem).

In all the following figures, the first sub-plot shows the evolution over time of the overall emissions $g(\mathbf{r}_k)$ (with the desired level g_* indicated by the dashed line), the next sub-plot displays the corresponding evolution of the blending-parameters $r^{(i)}$, and the last sub-plot gives the evolution of the emissions $t^{(i)}$. In Figure 2, we used a small network of $n = 4$ cars and set the gains $\eta_k^{(ij)}$ strictly as per Theorem 1. It can be seen that the global emissions converge to the desired level and that all cars indeed equalise their local emissions, which is achieved by using different energy blends depending on the emissions characteristics of the car.

However, convergence is somewhat slow compared to Figure 3 where we manually set the $\mu_k^{(i)}$ about 20 times higher than in the previous case — resulting in about 10 times faster convergence. Figure 4 shows the results for a fleet of $n = 50$ cars (again with adjusted gains). The “jumps” in all the sub-plots at times k that are multiples of $n - 1 = 49$ is of course due to the inclusion of the global term in the update rule. Note that in all simulations some agents use a larger blending-parameter than others (these would be cars with overall less polluting engines), which means they rely more on their combustion engines. This in turn means that these cars would typically be able to drive farther than others, so that their eco-friendliness is rewarded with extended range.

Comments In the simulations here only the update law in its basic form was used here, but in a real-life setting, the application may also require the two extensions of the algorithms to be used.

Also, the algorithm presented earlier requires the states (and utility values) to be defined for the entire field of real numbers. In the application presented here, however, both variables are only defined on compact intervals. We thus assume that, with the blending-parameters all initialised properly, the solution is feasible and does not drive the parameters beyond their domain of definition. If, however, this was the case, the blending-parameters would simply saturate at either fully electric or combustion mode.

Lastly, the CO₂ emissions of cars are typically strongly dependent on the driving speed as well as the individual driver’s behaviour — both of which is not taken into account here. We rather focus on the *average* emissions that would be produced in typical city traffic. Furthermore, the frequency at which new aggregate emissions measurements are provided determines the rate at which the discrete updates occur.

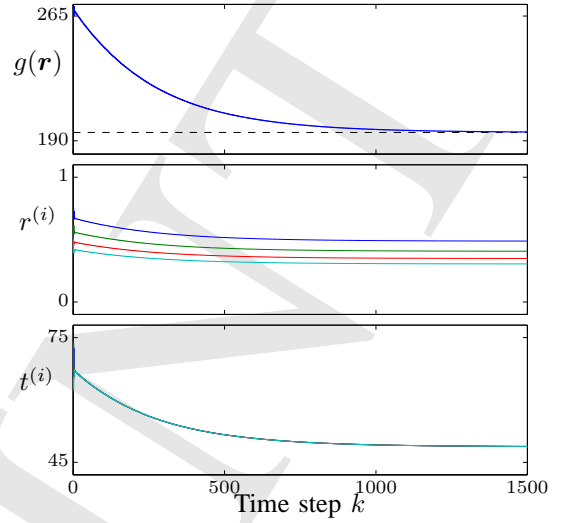


Fig. 2. Plots of the evolution over time of the aggregate CO₂ emissions, the blending parameters used and the resulting local emissions, respectively. Fleet of $n = 4$ cars, gains $\mu_k^{(i)}$ set according to Theorem 1.




Fig. 3. Same set-up as in Figure 2, but the gains $\mu_k^{(i)}$ were set 20 times larger than required by Theorem 1.




Fig. 4. Fleet of $n = 50$ cars, gains $\mu_k^{(i)}$ set manually.

V. CONCLUSION

In this paper, we have briefly reviewed previous work in the area of distributed consensus finding, and presented an algorithm that allows a multi-agent network to solve a very general implicit and constrained consensus problem. We then suggested an application for this algorithm that aimed at cooperatively regulating CO₂ emissions in a fleet of plug-in hybrid electric vehicles, and provided some simulation results that demonstrated its effectiveness. We would like to stress again, that we used CO₂ emissions purely for illustration purposes, any kind of emission (such as the directly harmful respirable dust produced by combustion engines) or combinations of different emissions may indeed be considered.

Future work in this area should consider the effect of nodes joining and leaving the network, how effects like saturation of the states could be incorporated directly into the mathematical framework, and ideally derive tighter bounds on the maximum permissible gain for the global term (as the bounds presented here are only sufficient for stability, and we have shown in the simulations that they can be increased significantly without compromising stability). Also, it would be interesting to attempt a real-life implementation of our suggested application.

REFERENCES

- [1] R. Olfati-Saber, J. A. Fax, and R. M. Murray, "Consensus and cooperation in networked multi-agent systems," *Proc. IEEE*, vol. 95, no. 1, pp. 215–233, Jan. 2007.
- [2] W. Ren, R. W. Beard, and E. M. Atkins, "A survey of consensus problems in multi-agent coordination," in *Proc. ACC'05*, vol. 3, Jun. 2005, pp. 1859–1864.
- [3] F. Knorn, M. J. Corless, and R. N. Shorten, "Results in cooperative control and implicit consensus," *Int. J. Control*, vol. 84, no. 3, pp. 476–495, Mar. 2011.
- [4] V. Borkar and P. P. Varaiya, "Asymptotic agreement in distributed estimation," *IEEE Trans. Autom. Control*, vol. 27, no. 3, pp. 650–655, Jun. 1982.
- [5] J. N. Tsitsiklis, D. P. Bertsekas, and M. Athans, "Distributed asynchronous deterministic and stochastic gradient optimization algorithms," *IEEE Trans. Autom. Control*, vol. 31, no. 9, pp. 803–812, Sep. 1986.
- [6] C. W. Reynolds, "Flocks, herds, and schools: A distributed behavioral model," in *Proc. SIGGRAPH'87*, Jul. 1987, pp. 25–34.
- [7] T. Vicsek, A. Czirók, E. Ben-Jacob, I. Cohen, and O. Shochet, "Novel type of phase transition in a system of self-driven particles," *Phys. Rev. Lett.*, vol. 75, no. 6, pp. 1226–1229, Aug. 1995.
- [8] A. Jadbabaie, J. Lin, and A. S. Morse, "Coordination of groups of mobile autonomous agents using nearest neighbor rules," *IEEE Trans. Autom. Control*, vol. 48, no. 6, pp. 988–1001, Jun. 2003.
- [9] R. Olfati-Saber and R. M. Murray, "Consensus problems in networks of agents with switching topology and time-delays," *IEEE Trans. Autom. Control*, vol. 49, no. 9, pp. 1520–1533, Sep. 2004.
- [10] L. Fang, P. J. Antsaklis, and A. Tzimas, "Asynchronous consensus protocols: Preliminary results, simulations and open questions," in *Proc. CDC-ECC'05*, Dec. 2005, pp. 2194–2199.
- [11] V. D. Blondel, J. M. Hendrickx, A. Olshevsky, and J. N. Tsitsiklis, "Convergence in multiagent coordination, consensus, and flocking," in *Proc. CDC-ECC'05*, Dec. 2005, pp. 2996–3000.
- [12] H. G. Tanner, A. Jadbabaie, and G. J. Pappas, "Stable flocking of mobile agents part II: Dynamic topology," in *Proc. CDC'03*, vol. 2, Dec. 2003, pp. 2016–2021.
- [13] R. Olfati-Saber, "Flocking for multi-agent dynamic systems: Algorithms and theory," *IEEE Trans. Autom. Control*, vol. 51, no. 3, pp. 401–420, Mar. 2006.
- [14] R. Stanojević and R. Shorten, "Fully decentralized emulation of best-effort and processor sharing queues," in *Proc. SIGMETRICS'08*. ACM, 2008, pp. 383–394.
- [15] J. Lawton, R. W. Beard, and B. J. Young, "A decentralized approach to formation maneuvers," *IEEE J. Robot. Autom.*, vol. 19, no. 6, pp. 933–941, Dec. 2003.
- [16] P. K. C. Wang, "Navigation strategies for multiple autonomous mobile robots moving in formation," in *Proc. IROS'89*, Sep. 1989, pp. 486–493.
- [17] M. Ji, A. Muhammad, and M. Egerstedt, "Leader-based multi-agent coordination: Controllability and optimal control," in *Proc. ACC'06*, Jun. 2006, pp. 1358–1363.
- [18] M. A. Lewis and K.-H. Tan, "High precision formation control of mobile robots using virtual structures," *Auton. Robot.*, vol. 4, no. 4, pp. 387–403, 1997.
- [19] R. W. Beard, J. R. T. Lawton, and F. Y. Hadaegh, "A feedback architecture for formation control," in *Proc. ACC'00*, vol. 6, Jun. 2000, pp. 4087–4091.
- [20] T. Balch and R. C. Arkin, "Behavior-based formation control for multirobot teams," *IEEE J. Robot. Autom.*, vol. 14, no. 6, pp. 926–939, Dec. 1998.
- [21] Q. Chen and J. Luh, "Coordination and control of a group of small mobile robots," in *Proc. ICRA'94*, vol. 3, May 1994, pp. 2315–2320.
- [22] A. Nedić, A. Ozdaglar, and P. A. Parrilo, "Constrained consensus and optimization in multi-agent networks," *IEEE Trans. Autom. Control*, vol. 55, no. 4, pp. 922–938, Apr. 2010.
- [23] F. Knorn, R. Stanojevic, M. Corless, and R. Shorten, "A framework for decentralised feedback connectivity control with application to sensor networks," *Int. J. Control*, vol. 82, no. 11, pp. 2095–2114, Nov. 2009.
- [24] A. Spence, S. T. A. Schelling, T. Benz, J.-P. Medevielle, I. McCrae, J. Jaaskelainen, and E. Boethius, "Methodologies for assessing the impact of ITS applications on CO₂ emissions," EC-METI Task Force, Tech. Rep., Mar. 2009.
- [25] "Central London Congestion Charging, Impacts Monitoring — Sixth Annual Report," Transport for London, Congestion Charging Division, Jul. 2008.
- [26] J. Schoenburg, "Germany introduces environmental zones," *Front. Ecol. Environ.*, vol. 6, p. 5, Feb. 2008.
- [27] K. S. Gallagher and E. J. Muehlegger, "Giving green to get green: Incentives and consumer adoption of hybrid vehicle technology," John F. Kennedy School of Government, Harvard University, Cambridge, MA, USA, Tech. Rep. RWP08-009, Feb. 2008.
- [28] B. Lane and S. Potter, "The adoption of cleaner vehicles in the UK: Exploring the consumer attitude-action gap," *J. Cleaner Prod.*, vol. 15, no. 11, pp. 1085–1092, 2007.
- [29] S. G. Wirasingha, N. Schofield, and A. Emadi, "Plug-in hybrid electric vehicle developments in the US: Trends, barriers, and economic feasibility," in *Proc. VPPC'08*, Harbin, China, Sep. 2008, pp. 1–8.
- [30] "Road transport, air pollution and health," Briefing. Friends of the Earth Trust, London, UK, Mar. 1999.
- [31] K. Bilstrup, E. Uhlemann, E. G. Ström, and U. Bilstrup, "Evaluation of the IEEE 802.11p MAC method for vehicle-to-vehicle communication," in *Proc. VTC'08-Fall*, Calgary, Canada, Sep. 2008, pp. 1–5.
- [32] TMC Forum, "What is Traffic Message Channel (TMC)?" Website: <http://goo.gl/Xfnvx>, 2004.