

Smart Procurement of Naturally Generated Energy (SPONGE) for Plug-in Hybrid Electric Buses

Joe Naoum-Sawaya, Emanuele Crisostomi, Mingming Liu, Yingqi Gu, and Robert Shorten,

Abstract—We discuss a recently introduced ECO-driving concept known as SPONGE in the context of Plug-in Hybrid Electric Buses (PHEB)’s. Examples are given to illustrate the benefits of this approach to ECO-driving. Finally, distributed algorithms to realise SPONGE are discussed, paying attention to the privacy implications of the underlying optimisation problems.

Note to Practitioners: **Abstract**—In this paper we present a new idea for ECO-driving for buses. It is an IoT concept - that instead of connecting devices in space, connects devices in time via forecasting engines. Basically, a bus uses knowledge of the available energy at the next charging step, to optimise its performance beforehand. The system can be implemented using available (free) forecasting engines, and existing distributed optimisation tools. A sample implementation is described using a Toyota plug-in Prius (as a proxy for a hybrid bus). Apart from the forecasting and optimisation analytics, the only additional work needed was the development of an interface unit to control EV mode of the vehicle, and the development of a smart-phone app. Future work will investigate impacts of our approach on the grid, integration of the ideas into the hybrid drive cycle, and using driver behaviour as an input into the design of the utility functions.

Index Terms—Primary Topics: Number 8, Number 9; Secondary Topic Keywords: Distributed Systems, Control Theory

I. INTRODUCTION

WE discuss a recently introduced holistic ECO-driving concept known as SPONGE (Smart Procurement of Naturally Generated Energy) in the context of Plug-in Hybrid Electric Buses (PHEB)’s. PHEB’s are increasingly seen as an effective tool in combating air pollution in our cities, and as a tool for reducing our cities reliance on fossil fuels (thereby reducing greenhouse gas emissions) [1], [2]¹. Consequently, the design and operation of such buses has been the subject of much research interest. Hitherto, significant research effort has focused on improving the fuel economy while guaranteeing that both the engine and the electric machine work in the high-efficiency area; typically, by taking into account knowledge of both bus routes and passenger loadings in a predictive manner.

J. N. Sawaya is with the Ivey Business School, University of Western Ontario, 1255 Western Rd, London, Canada (e-mail: jnaoum-sawaya@ivey.ca).

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

Y. Gu and M. Liu and R. Shorten are with the School of Electrical, Electronic and Communications Engineering, University College Dublin, Ireland (e-mail: yingqi.gu@ucdconnect.ie; mingming.liu@ucd.ie).

R. Shorten is also with IBM Research, Dublin, Ireland (e-mail: robshort@ie.ibm.com).

¹See <https://chargedevs.com/newswire/ultramodern-plug-in-buses-go-into-service-in-gothenburg/> for further examples.

Selected examples of work in this direction can be found in [3]–[5].

Our objective in this paper is to extend this line of inquiry further. Our basic setting is to consider a bus operator that has access to a fixed amount of cheap renewable energy on a daily basis. For example, some operators may own solar farms or have access to wind generation. It makes sense to use this (inexpensive) *free energy* before consuming electrical energy that is bought from the grid, and in situations where there is an oversubscription for this free energy, the operator then has a choice as to how this energy is distributed to each bus. For example, some drivers may be more efficient than others. Thus, it makes eminent economic sense, to distribute this free energy to reduce the impact of less efficient drivers in optimising the hybrid engine cycle, while at the same time ensuring that sufficient energy is consumed to make room for every unit of free energy that arrives the next time the buses recharge. Specifically, SPONGE for buses operates as follows.

- A. We introduce the forecast of generation of energy from renewable resources on a day ahead basis as a further variable to influence the energy management system for the bus operator.
- B. We use this forecast to prioritise the manner in which individual buses choose their driving mode.
- C. We do this by prioritising the utilisation of energy from renewable sources over other resources, and by taking account of the fact that some drivers/routes are more energy efficient than others.

Prioritising energy from renewable sources in this manner introduces a number of benefits for the bus operator and society.

- The use of energy from renewable sources (e.g., wind turbines, dynamic water power, or solar power) achieves environmental health benefits with respect to the use of the “power grid average” electricity [6].
- Financial benefits for the bus operator.
- Depleting PHEBs’ batteries of a pre-specified quantity of energy allows better grid-demand balancing. That is, the energy provider knows in advance how much energy will be required by PHEBs, when connected for charging. This makes the electrical load of PHEBs to be fully predictable and dispatchable, thus mitigating the burden of the power grid to accommodate a not-known-in-advance electrical load.

Note that the proposed energy management approach closely resembles the widely discussed practice of demand side management, where electricity customers shift their elec-

trical loads taking into account the expected availability of energy from renewable sources (e.g., solar panels on the roofs of their houses). In fact, in this paper we are considering the possibility that buses orchestrate the consumption of their batteries by considering the amount of energy that will be available from renewable sources when recharging.

A. Contribution

This paper extends previous work of some of the authors in [7] for the case of electric cars. While the seminal idea of matching energy from renewable sources with space in the battery of the EVs remains the same, the case of PHEBs substantially differs from the case of PHEVs for a number of reasons: (i) in the case of buses it is possible to know the route in advance, thus, in contrast to [7] such knowledge is considered in the optimisation formulation; (ii) in the case of buses, it makes perfect sense to assume that buses of the same company will collaborate to achieve a common goal (e.g., the minimisation of the electric energy bought by the company to supply the electric public transport services); (iii) another difference is however that the optimisation problem is here solved off-line in a batch fashion, taking advantage of the available information (i.e., power generation forecasts and the knowledge of the daily routes). On the other hand, the optimisation problem must be solved in real-time in the case of single cars, given that the time of use and the daily routes are not known in advance. Accordingly, speed of convergence is of paramount importance when choosing an algorithm to be applied in real-time, while here we are more interested in other aspects that include communication requirements, agent actuation, and privacy preservation. With this latter aspect in mind, our final contribution is to give a brief comparison of two competitive optimisation algorithms.

The paper is structured as follows. The SPONGE problem formulation is presented in Section II. The discussion of the proposed Additive Increase Multiplicative decrease (AIMD) optimisation algorithm is presented in Section III. The experimental results are presented in Section IV. The practical implementation of the proposed SPONGE system is briefly discussed in Section V. Finally, a brief conclusion is presented in Section VI.

II. SPONGE PROBLEM FORMULATION

A. Assumptions

The starting point in our work is to consider the actuation possibilities offered by a hybrid powertrain, namely the ability to switch in and out of EV mode, as a means not only to improve the efficiency of an individual vehicle, but also to serve the needs of other stakeholders. This view is consistent with other recent works where the engine management logic is used to help other stakeholders - such as pedestrians by keeping local air quality clean, and energy suppliers by helping to balance the needs of the grid and the transportation network [8]. In particular, it is with this latter view in mind that the paper is written.

Let $\mathcal{N} = \{1, 2, \dots, N\}$ denote the set of N PHEBs participating to the SPONGE programme. We shall make the following assumptions:

- We assume that after a number of trips along their (different) routes, the N PHEBs stop for charging at the bus station. For instance, we can assume that the PHEBs will not drive from 11pm to 6am, and they will be charged in this time frame;
- We also assume that a 24-hour ahead forecast of energy from the renewable energy sources available to the operator will be available as well (e.g., a forecast of how much energy will be generated by the wind plants connected with the charging station at night time). We denote this amount of available energy by E_{av} ;
- Early in the morning, before being dispatched along their routes, the buses will compute how the energy E_{av} should be optimally shared among themselves during the day (i.e., in terms of energy consumption of their own batteries);
- In order to compute the optimal allocations of energy, we shall assume that each PHEB is equipped with a device to transmit messages to the central infrastructure via Vehicle-to-Infrastructure (V2I) technology;
- The central infrastructure has the ability to broadcast messages to the whole network of PHEBs using some Infrastructure-to-Vehicles (I2V) technology.

Note that in our set-up we shall not require vehicles to exchange information among themselves, and thus, we shall not require PHEBs to be equipped with Vehicle-to-Vehicle (V2V) communication devices. A schematic diagram of the above SPONGE scenario is illustrated in Fig. 1.




Fig. 1: A schematic diagram of the SPONGE programme.

B. Optimisation problem

In this context, we denote by d_i the energy consumption by the i 'th bus during the day. Then we are interested in

computing the solution of the following optimisation:

$$\begin{cases} \max_{d_1, d_2, \dots, d_N} & \sum_{i \in \mathcal{N}} f_i(d_i) \\ \text{s.t.} & \sum_{i \in \mathcal{N}} d_i = E_{av} \end{cases} \quad (1)$$

In the optimisation problem (1), the terms d_i can be interpreted as a “budget” of energy that is allocated to the i 'th bus in order to maximise a utility function of interest, such that the sum of the energy budgets allocated to all the buses matches E_{av} as in the SPONGE spirit. Although in principle the utility function $f_i(d_i)$ may be chosen in an arbitrary fashion, to represent any *utility*, in this work we shall explore the particular case where one is interested in the utility of CO_2 emissions savings $f_i(d_i)$ as achieved by each bus. Clearly, each $f_i(d_i)$ is an increasing function of d_i as no CO_2 emissions are saved when the bus travels all the time in ICE mode, while no pollution occurs when all the vehicles travel in electric mode all the time.

Remark 1 - Switching mode: The assumption that a hybrid bus can travel in pure ICE or pure electric mode is realistic (e.g., for parallel hybrids) but it is not strictly required. In fact, our work can be generalised to include switching between two (or even more) arbitrary driving modes (e.g., ECO-drive and sportive mode) that give rise to different energy consumption patterns when undergoing the same driving cycle. More specifically, for ease of exposition, the next section describes how the utility functions f_i are constructed based on the assumption of two driving modes only.

Remark 2 - Exploiting hybrid vehicles to provide ancillary services: Note that traditionally, the principal concern of the hybrid architecture is the fuel efficiency of the vehicle. However, the hybrid architecture allows cities to move from a sole focus of optimising the performance of the bus (with the driver, or the bus company, as the stakeholder) to optimising the performance of the vehicle with respect to other stakeholders (pedestrians) etc. See [8] for examples of work in this direction. Our present strategy can be viewed as a mix of these two approaches, where energy budgets are used to optimise bus performance, but where the vehicle can choose where and when to deploy the pure EV mode with a view to maximising some social utility.

C. Construction of the utility functions

Our main assumption in constructing the utility functions is that a forecast of the expected energy consumption and CO_2 emissions (perhaps with some other higher order statistics) is available for each route, for each of the two driving modes. Such a forecast is itself a function of the time of the day and the specific day of the week. The forecasts can be easily made by measuring such quantities directly on-board for each trip of each bus-line in order to build a data-base of data, and post-processing the recorded data (e.g., averaging measurements to remove stochastic effects). In fact, while instantaneous energy consumption or emissions can not be accurately predicted in advance, it is reasonable to assume that the consumption

patterns associated with a given bus trip at a given time on a given day is predictable to some degree; i.e., similar bus trips require on average a similar aggregate amount of energy (or generate comparable quantities of CO_2 emissions).



Fig. 2: Road network of Dublin City, Ireland, imported from OpenStreetMap, used in our simulations. The trajectory in the map illustrates one bus trip starting from a bus-stop located in the south-west of Dublin city to another situated in Dublin city centre.

By “trip” we intend a single journey from a bus terminus to the other terminus, as illustrated in Fig. 2. Then it is assumed that the next trip will correspond to the return journey, and so on for the remainder of the day. For the purpose of this paper, we employ simulation measurements rather than real measurements to obtain the energy consumption and emission averages, by adopting the popular mobility simulator SUMO and simulating the routes of 16 PHEBs in Dublin city, Ireland. More details about the mobility platform and the simulation set-up will be provided in Section IV-A. The final utility functions are shown in Fig. 3 as a function of the percentage of the use of the electrical engine for each bus.

The utility functions $f_i(d_i)$ depicted in Fig. 3 show how much CO_2 has been saved by the i 'th bus, provided that the bus is allowed to travel in EV mode a given percentage of its route. In particular, it can be noted that the utility functions are non-decreasing functions (i.e., the more one PHEB is allowed to travel in EV mode the more CO_2 is saved) and that some bus-lines generate more CO_2 than others (this information can be retrieved by observing how much pollution can be saved by each bus-line if the bus travels all the time in electric mode). In order to derive the mathematical formulation of the utility functions, let us denote by e_{ij} and p_{ij} the expected energy consumption by the i 'th bus during its j 'th trip when travelling in EV mode, and the expected pollution by the i 'th bus during its j 'th trip when travelling in ICE mode, respectively. Then




Fig. 3: Utility functions of 16 PHEBs in Dublin city (considered outside the greenzone). Note that some buses pollute more than others (and thus, have a greater potential in terms of CO_2 savings) depending on the characteristics of their routes (e.g., speed limits, traffic conditions, topography of the route). The “.” points are fitted using cubic splines. For the trips passing by the city centre, we here show the utility functions corresponding to the part of the trip outside the green zone (for which the driving mode has to be decided).

we have that

$$\left\{ \begin{array}{l} f_i(d_i) = \max_{\gamma_{ij}} \sum_{j \in \mathcal{T}_i} \gamma_{ij} \cdot p_{ij} \\ \text{s.t.} \sum_{j \in \mathcal{T}_i} \gamma_{ij} \cdot e_{ij} = d_i \\ 0 \leq \gamma_{ij} \leq 1, \forall j \in \mathcal{T}_i \\ \gamma_{ij} \cdot e_{ij} \geq \underline{e}_{ij} \end{array} \right. \quad (2)$$

where γ_{ij} is the fraction of time that the i 'th bus spends in EV mode during the j 'th trip and \mathcal{T}_i is the set of daily trips made by the i 'th bus. The last equation of (2) takes into account that some trips need to be partly travelled in EV mode due to possible strict laws. In fact, in some cities it is mandatory to travel in EV mode to access some areas denoted as “green zones” (e.g., in the proximity of the city centre). See for instance the case of Germany². Mathematically, this corresponds to assuming that the energy allocated for each trip has to exceed a lower bound \underline{e}_{ij} that corresponds to the (expected) energy required to travel in EV mode in the green zone. In our work, we assumed that the green zone can be exemplified with a circular area in the city centre, as shown in Fig. 2.

Due to the fact that all bus routes are fixed and known a priori, and given a fixed d_i , then (2) is a linear program with a single budget constraint (i.e., a continuous linear knapsack problem [9]) and thus the optimal electric energy allocation can be easily computed by sorting the trips by

decreasing order of the values of p_{ij} and then activating the electric energy according to the sorted order. The utility function of each PHEB can thus be computed off-line. Particularly, for each bus i , we vary d_i between 1 and 100 in steps of 1 and compute the optimal $f_i(d_i)$. We note that (2) is a parametric linear program with parameter d_i , and thus $f_i(d_i)$ for all $i \in \mathcal{N}$ is a piece-wise concave function [10], and as such possibly non-differentiable. However, since the derivative $f'_i(d_i)$ is required by the optimisation algorithm that is proposed next, $f_i(d_i)$ for each bus i is approximated using cubic spline functions. The resulting (normalised) utility functions for the 16 PHEBs that are used in the illustrative example of this paper are shown in Fig. 3. Note that some utility functions, corresponding to the specific buses passing by the city centre, are not defined for small values of d_i , as some minimum budget was required anyway to travel in EV mode in the green zones.

Remark 3 - Feasibility: Problem (2) may not be feasible if the overall available energy is smaller than the electrical energy required to travel in the green zones. Accordingly, in the following we shall assume that $E_{av} \geq \sum_{i \in \mathcal{N}} \sum_{j \in \mathcal{T}_i} \underline{e}_{ij}$. Alternatively, one could relax the green zones hard constraints, and compute a best-effort solution (i.e., the buses try to travel in EV mode in the green zones as much as possible, given the scarce level of their batteries).

III. ALGORITHMS AND OPTIMAL SOLUTION

The optimisation problems (1) and (2) can in principle be easily solved in a centralised way adopting simple Linear Programming (LP) techniques. In order to do so, it is required that all the utility functions are known to the central agent. In this work however, we are interested in solving equations (1) and (2) in a distributed manner, to avoid having to reveal the utility functions to the central agent. Such a possibility has a number of advantages over the centralised one. In particular, this allows us to handle the privacy preservation and agent actuation aspects, as discussed in the introduction. More specifically, the utility functions (i.e., average energy consumption and pollution along a trip) depend on some publicly known information (e.g., road characteristics and traffic) and on other more private information (e.g., number of passengers on board, driving style of the driver) that may not be wanted to be revealed. Also, it could happen that a single energy provider serves different bus companies, which obviously may not be interested in sharing such data. Accordingly, in this paper we are interested in a distributed solution that is more flexible in handling a larger number of possible scenarios.

In principle, many different methods may be used to solve the optimisation problem (1) that arises in our work (for instance ADMM-like algorithms [11]). ADMM (Alternating Direction Method of Multipliers) is a popular optimisation algorithm, that has been recently proposed as an evolution of other well-known optimisation algorithms, like the dual ascent and the method of multipliers. As an alternative

²<http://gis.uba.de/website/umweltzonen/umweltzonen.php>

to ADMM-like algorithms, our choice here is to adopt an AIMD-like algorithm [12] to solve the problem in a distributed fashion. Such a choice is motivated by many reasons:

- **Low-communication requirements:** Although we have presented here a simple case study with a small number of buses, the same programme can be easily generalised to include hundreds of buses. Also, the batch optimisation formulation might be solved in real-time to account for non fully-predictable aspects (for example to respond to traffic peaks or weather forecast updates). In this context, it is convenient to consider the communication cost of solving the optimisation algorithm. *AIMD based optimisation can be solved using only intermittent binary feedback and can thus, unlike many other distributed optimisation techniques, be solved without the need to broadcast of the Lagrange multipliers in a pseudo-continuous manner.*
- **Privacy-preservation requirements:** In our application, the utility functions $f_i(d_i)$ for all $i \in \mathcal{N}$ potentially reveal sensitive private information. For example, these functions contain historical information of how good a particular driver is on a given route. This information is potentially very useful for an employer and could potentially be used in a nefarious manner. In addition, in unionised environments, revealing these functions to an employer could also be of concern and consequently impede the adaptation of ideas like SPONGE. Given this context, a natural question is whether the distributed optimisation can be solved without revealing private information. *As we shall see, AIMD has some very nice privacy properties.*
- **Agent actuation:** *AIMD requires very little actuation ability on the agent-side.* This is in contrast to ADMM where at each time step, agents must solve a local optimisation problem.
- **Algorithm parameterisation:** *In AIMD the gain parameters of the network are independent of network dimension; rather, they only depend on the largest derivative over all utility functions.* Thus, selecting a gain for the algorithm is extremely simple in the case of AIMD.

As we shall further discuss in the following section, AIMD is thus a convenient alternative to ADMM, when the previous aspects are relevant.

A. AIMD Algorithm

Additive Increase Multiplicative Decrease (AIMD) algorithms were originally applied for solving issues arising in network congestion in the Internet [13]. To date, this idea has been widely explored for the design of practical algorithms for other applications as well, as for instance, network applications see [14]–[16], and smart grid applications see [17]–[20]. More recently, an unsynchronised AIMD algorithm based on the nonhomogeneous place-dependent Markov chains model was

proposed in [12] to solve utility optimisation problems. The pseudo-code of the proposed algorithm is given in Algorithm 1. Note that the algorithm does not compute the optimal

Algorithm 1 Unsynchronised AIMD Algorithm

```

1: Initialisation:  $k = 1, d_i(k) = 0;$ 
2: Broadcast the parameter  $\Gamma$  to the entire networks;
3: while  $k < k_{\max}$  do
4:   if  $\sum_{i=1}^N d_i(k) < E_{av}$  then
5:      $d_i(k+1) = d_i(k) + \alpha$ 
6:   else
7:     generate uniform random number,  $0 < r_i < 1$ , and
8:     calculate  $p_i(k) = \Gamma \frac{1}{\bar{d}_i(k) f'_i(\bar{d}_i(k))}$ 
9:     if  $r_i < p_i(k)$  then
10:       $d_i(k+1) = \beta d_i(k)$ 
11:    else
12:       $d_i(k+1) = d_i(k) + \alpha$ 
13:    end if
14:  end if
15:   $k = k + 1$ 
16: end while

```

budgets d_i in a single step, but in an iterative fashion, as $d_i(k)$ represents the value of the unknown energy to be allocated to the i 'th PHEB, computed at time step k . For large values of k , $d_i(k)$ will eventually converge to the optimal solution that maximises the environmental benefits (while still satisfying the energy constraint). In Algorithm 1, k_{\max} represents the maximum number of iterations before the algorithm stops (e.g., after five minutes of iterations).

The basic idea of Algorithm 1 is that if the sum of the $d_i(k)$ of all PHEBs is smaller than E_{av} , then each PHEV increases its target energy consumption $d_i(k)$ at the next iteration $k+1$ by a quantity α . However, if the sum of the energy budgets of all PHEVs exceeds E_{av} (this situation is usually called as a congestion event), then each PHEB decreases its energy consumption by a multiplicative factor $0 < \beta < 1$ with probability $p_i(k) = \Gamma \frac{1}{\bar{d}_i(k) f'_i(\bar{d}_i(k))}$, where Γ is a constant common broadcast parameter, and $\bar{d}_i(k)$ is the time average of the sequence of $d_i(k)$ at congestion events, up to the last iteration. It is proved in [12] that $\bar{d}_i(k)$ approaches to the optimal solution of the problem when Algorithm 1 converges and where the optimisation is carried out over the $f_i(\bar{d}_i)$ for all $i \in \mathcal{N}$.

The philosophy underlying the AIMD algorithm is to adjust $p_i(k)$ and $d_i(k)$ at every time step k such that for large values of k , $f'_i(\bar{d}_i(k)) = f'_j(\bar{d}_j(k))$, $\forall i \neq j \in \mathcal{N}$, or in other words the PHEBs achieve consensus on the derivatives of their utility functions. This, with strict convexity of the utility functions, is both necessary and sufficient for optimality when feasibility is guaranteed. This property is known from elementary optimisation theory. Algorithm 1 was originally designed in [12] to *minimise* a cost function of interest, here we slightly adapt it to *maximise* CO_2 savings. Accordingly, given that each approximated utility function f_i in our case is strictly concave, and that the p_i are strictly non-increasing

in our problem, then we can adapt the algorithm in [12] so that consensus is achieved on $1/f'_i$, and the convergence and optimality properties of the algorithm are preserved.

The AIMD algorithm with finite window is an example of an *iterated function system* [21]. Such systems have been widely studied in the literature; see [22] and the references therein. Strict concavity of the utility functions is not necessary for convergence of the system to a unique stationary invariant measure (ergodicity). Strict concavity is however required in our context for convergence to the optimal point of the associated optimisation problem.

The AIMD algorithms as described are stochastic algorithms. Almost sure convergence of the long term average to the optimal point is proved in [12]. Consequently, for every convergent trajectory, the variance about the optimal point converges to zero asymptotically. Convergence to the optimum follows convergence of the long term average and can be slow when defined in terms of congestion epochs. However, as only low bit communication is required to ensure convergence, convergence measured in terms of communication effort may not be so bad. In fact, and as we shall see, when this is taken into account, simple experiments suggest that its convergence properties are comparable with other better known schemes.

B. Privacy aspects

We now make some brief comments concerning the privacy properties of AIMD based optimisation. Recall that we assume that the central agent may receive the value d_i from agent i , and performs the aggregation $A = \sum_{i=1}^N d_i$. We also assume that there are no incentives for an agent to cooperate with the central agent to help deduce the f'_i ; that is, all agents, other than the central agent, are honest. Given this basic setting, one may discern the following four basic levels of privacy.

- (i) **Absolute utility privacy (AUP):** Here, the central agent cannot deduce $f'_i(d)$ based on knowledge available to it. This is a basic level of privacy.
- (ii) **Relative utility privacy (RUP):** Here the central agent cannot deduce whether $f'_i(d) > f'_j(d)$. This again, is a basic level of privacy.
- (iii) **Absolute derivative privacy (ADP):** Here, the central agent cannot deduce $f'_i(d)$ based on knowledge available to it. This information is important since it allows the central agents to estimate the *price* elasticity of individual agents.
- (iv) **Relative derivative privacy (RDP):** Here the central agent cannot deduce whether $f'_i(d) > f'_j(d)$.

A more rigorous discussion on privacy preservation is clearly beyond the scope of this paper. However, we note briefly that the stochastic AIMD algorithm allows us to give guarantees regarding some of these privacy categories. First, since the optimisation is based on $f'_i(d_i)$, the AIMD algorithm can be considered *AUP*- and *ADP*-private. Deducing any $f'_i(d_i)$ would require estimation of the $p_i(k)$ in Algorithm 1. Clearly, this is

difficult (but not impossible) except at optimal points. However, since our algorithm only requires an implicit consensus among all derivatives, one may replace in the formula for $p_i(k)$ (i.e., line 8 of Algorithm 1), $f'_i(d_i)$ with $g(f'_i(d_i))$, where g is chosen so that the convergence conditions in [12] are satisfied. Clearly, without knowledge of the function g , the central agent cannot deduce $f'_i(d_i)$ even if the probabilities p_i 's are correctly estimated when the algorithm converges.

IV. SUMO SIMULATIONS

A. Simulation Set-up

In this section, we evaluate the performance of the proposed AIMD algorithm in a realistic traffic scenario³, where vehicular flows are simulated using the popular mobility simulator SUMO [23]. In doing so, we shall also compare the results obtained using AIMD with those obtained with the ADMM algorithm. All the simulations are performed over the road network of Dublin, Ireland, depicted in Fig. 2, imported from OpenStreetMap [24].

B. Simulation Results

We assume that 16 PHEBs participate to the SPONGE programme in Dublin city, Ireland. We further assume that weather forecasting tools predict an availability of 10MWh (about 55% of the energy required by the buses to travel in EV mode for the whole time) in the next charging period. Before starting their routes, the optimisation problem is solved using the described AIMD algorithm with parameters $\alpha = 1$ and $\beta = 0.5$, and the available energy is optimally allocated to the 16 different buses. Fig. 4 compares the overall energy that would be required for each of the 16 trips when travelling all the time in EV mode (blue bars) with the optimal allocated budgets (red bars). Note that the first 6 buses also need some minimum energy to travel in the green zones, which is reported with the yellow bars. Fig. 5 shows how much energy is expected to be required to travel in EV mode for each bus route, and the expected CO_2 emissions as well. Note that the quantity is not constant, as it depends at what time of the day a single trip will take place (i.e., with different expected traffic conditions). Finally, Fig. 6 shows the details of the final solution (i.e., how much energy is allocated per route per bus).

Fig. 7 and Fig. 8 show that the AIMD algorithm indeed converges to the optimal solution and that the necessary condition for optimality (KKT) is when AIMD converges (i.e., the derivatives of the utility functions converge to the same value), respectively. Comparatively, Fig. 9 demonstrates that ADMM converges to the optimal solution as well. We note that although ADMM requires less iterations to converge (15 iterations to the first time reaching 5% error of the optimal solutions) compared to AIMD (5720 iterations to the first time 90% of samples are within 5% error of the optimal solutions, by using a window of 1000 congestion events),

³This paper has supplementary downloadable material provided by the authors. This includes the trip of one PHEB simulated in SUMO and a readme file. This material is 77.7MB in size.

ADMM requires more data to be transmitted to the agents. For instance, if we consider that at each iteration ADMM needs to broadcast a packet (with multiplier) to each bus in 64 bits, then the total data required for algorithm convergence is $15 \cdot 16 \cdot 64 = 1.92\text{kB}$. On the other hand, AIMD needs to transmit one bit for all buses only on congestion events so the maximum data that is transmitted is 0.715kB . This shows that AIMD is competitive from the perspective of communication overhead when compared to the ADMM algorithm. Finally, Fig. 10 shows that the distributed solution obtains the same results of a centralised LP solution, even if the utility functions have been slightly changed (i.e., to make them strictly concave), for values of the available energy E_{av} ranging from 10% to 100% of the all energy required to travel in EV mode all the time.

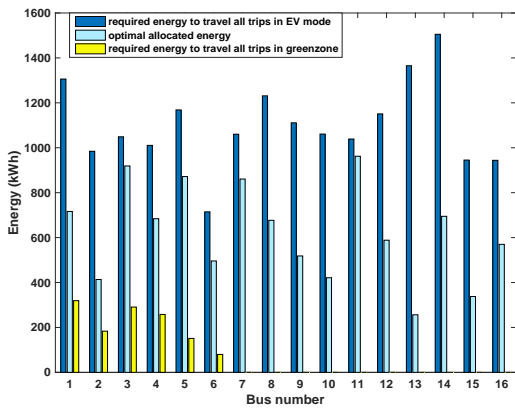


Fig. 4: Comparison of different energy consumption patterns for 16 PHEBs.

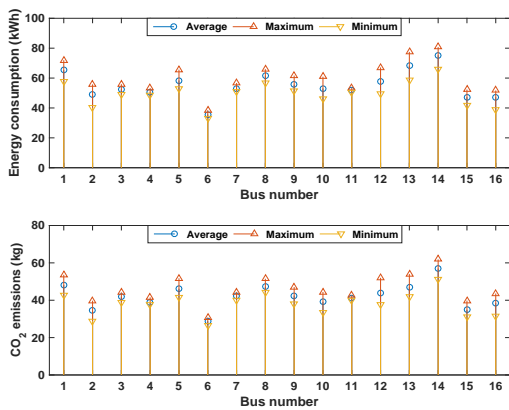


Fig. 5: Energy consumption and emission generation patterns for all bus routes.

V. COMMENTS ON THE PRACTICAL IMPLEMENTATION OF SPONGE

To conclude the paper we now briefly comment on the feasibility of the testing and implementation of a SPONGE

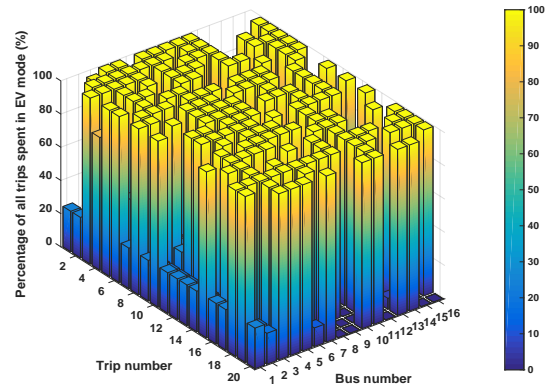


Fig. 6: Percentage of electric energy used for all PHEBs of all their trips.




Fig. 7: AIMD converges to the correct solution.




Fig. 8: The derivatives of the utility functions converge in a consensus fashion.




Fig. 9: Evolution of the distributed ADMM.

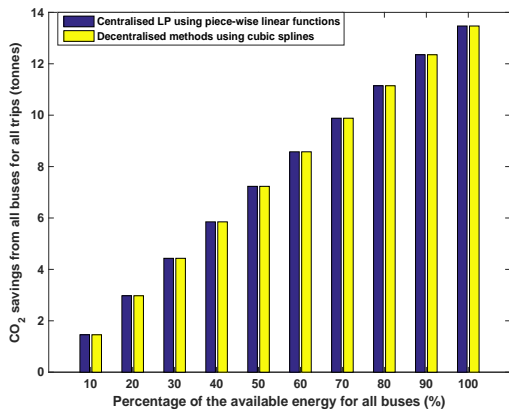


Fig. 10: Overall CO_2 savings of all PHEBs with respect to different percentage of energy available.

programme.

A. Large-Scale Traffic Simulator: As we have mentioned, all simulations are based on the SUMO simulation environment. SUMO [23] is an open source, microscopic road traffic simulation package primarily being developed at the Institute of Transportation Systems at the German Aerospace Centre (DLR). SUMO is designed to handle large road networks, and comes with a “remote control” interface, TraCI (short for Traffic Control Interface) [25], that allows one to adapt the simulation and to control singular vehicles on the fly. Based on this, large-scale hardware-in-the-loop emulations with both actual vehicles and (possibly) thousands of simulated vehicles can be easily performed as described in [26].

B. Test Vehicle: While we have not yet implemented SPONGE in a real bus, the algorithm has been implemented in a real test vehicle. Our test vehicle is a 2015 Toyota Prius VVTi 1.8 5DR CVT Plugin Hybrid vehicle and is pictured in Fig. 11. The engine management system of the Prius allows the vehicle to be powered by the ICE alone, the battery, or using a combination of both, and it is this degree of freedom that we exploit to implement SPONGE. For the purpose of



Fig. 11: Field-test vehicle: 2015 Toyota Prius.

this programme, we have made some important modifications to the basic vehicle to make it behave as a context-aware vehicle. First, we automate the switching of the vehicle from ICE to EV mode by adapting the EVmode button hardware in the vehicle. For this purpose, a dedicated Bluetooth-controlled mechanical interface was constructed to override the manual EV button based on signals from a smartphone. The switching is based on GPS location, external context information, and onboard signals such as speed and battery level. Second, special-purpose hardware was constructed to permit communication between a smartphone and the controller area network (CAN) bus. The Prius provides a CAN access on the vehicle diagnosis *On Board Diagnosis II (OBDII) interface*. Our hardware module acts as a gateway between this CAN interface and the smartphone. The module is directly connected to CAN and to the smartphone via Bluetooth. Communication to other vehicles, to GPS, and to a cloud server is also realised using a smartphone device. To control the driving mode, the software connects via Bluetooth to a mechanical switch to toggle driving mode between the EV mode and non-EV driving modes. In our application we use a Samsung Galaxy S III mini (model no. GT-I8190N) running the Android Jelly Bean operating system (version 4.1.2) and the OBD2 interface device that we used was the Kiwi Bluetooth OBD-II Adaptor by PLX Devices.⁴

C. Weather forecasting: An important component in any real practical implementation of the SPONGE programme is the ability to have a reasonably accurate, and cheap, prediction of the expected energy that will be available for charging E_{av} . To obtain a feeling for fidelity of such tools, we evaluated the accuracy of a *free* online forecasting tool over a 3 month period. The tool that we evaluated is provided by the Technical University of Crete and is described in [27], where the energy

⁴PLX Devices Inc., 440 Oakmead Parkway, Sunnyvale, CA 94085, USA. Phone: +1 (408) 7457591. Website: <http://www.plxdevices.com>

generated by a solar plant can be predicted (anywhere in the world) by simply providing the technical parameters of the plant. We collected real data on-site from PV panels mounted on the flat roof of the building in University College Dublin, Ireland. We recorded a total of 100 days and the predicted and the actual recorded energy are shown in Fig. 12. As also shown in Fig. 13 the predictions are relatively accurate with 80% of the predictions within 3% of Normalised Mean Absolute Error (NMAE) and the maximum NMAE is 7%. Thus, our data suggests that accurate predictions can be performed even for small powers, and even when a free online tool is employed. As for wind power forecasts, we note that a recent study in Germany reported that “typical wind-forecast errors for representative wind power forecasts for a single wind project are 10% – 15% root mean square error of installed wind capacity but can drop down to 6% – 8% for day-ahead wind forecasts for a single control area and to 5% – 7% for day-ahead wind forecasts for all of Germany”⁵. The accuracy may further be increased if other (commercial) tools are employed. From the previous discussion it appears reasonable to claim that on average the prediction error is below 10%, and this is consistent with other recent studies as well, see for instance [28] and [29].

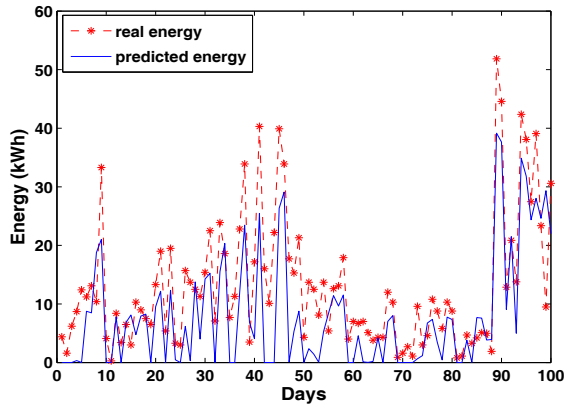


Fig. 12: Comparison between the real and the predicted energy generated from PV panels in UCD.

Comment: While the effect of uncertainty is beyond the scope of this paper, we note briefly that, it is simple to accommodate for forecasting errors by buying extra energy, if required, from the outer grid, or by appropriately using other storage devices, if available. However, interactions with the grid are not always convenient, either in terms of price, or in terms of environmental friendliness of the average power mix from the grid (see [30]). An alternative to this is to formulate an uncertainty description as part of the optimisation, and this will be part of future work.

VI. CONCLUSION

In this paper, we introduce an optimal energy allocation scheme for the SPONGE system in the context of PHEBs. We describe a distributed AIMD algorithm for solving the

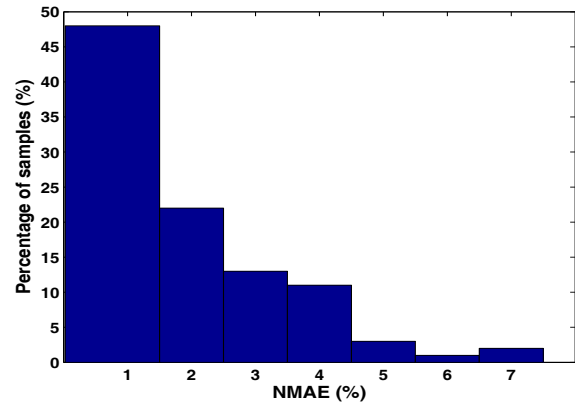


Fig. 13: Histogram of the percentage of NMAE.

optimisation problem. The main features of the proposed AIMD approach are the low-communication requirements and the privacy-preserving properties. The proposed approach is demonstrated on a case study with 16 buses with different energy requirements. The results demonstrate significant environmental benefits in terms of CO_2 emissions that can be achieved with optimal use of free renewable energies.

ACKNOWLEDGMENT

The authors gratefully acknowledge funding for this research provided by Science Foundation Ireland under grant 11/PI/1177. The work of E. Crisostomi was supported in part by the University of Pisa under Project PRA 2016 “Sviluppo di strumenti di analisi e gestione per l’approvvigionamento e la distribuzione dell’energia”.

REFERENCES

- [1] D. Lowell, “Comparison of modern CNG, Diesel and Diesel hybrid-electric transit buses: Efficiency & environmental performance,” *Report - M.J. Bradley & Associates, LLC*, 2013.
- [2] R. D. Jong, M. Ahman, R. Jacobs, and E. Dumitrescu, “Hybrid electric vehicle an overview of current technology and its application in developing and transitional countries,” *UNEP. Kenya*, 2009.
- [3] F. Tianheng, Y. Lin, G. Qing, H. Yanqing, Y. Ting, and Y. Bin, “A supervisory control strategy for plug-in hybrid electric vehicles based on energy demand prediction and route preview,” *IEEE Transactions on Vehicular Technology*, vol. 64, no. 5, pp. 1691–1700, 2015.
- [4] L. Li, C. Yang, Y. Zhang, L. Zhang, and J. Song, “Correctional DP-based energy management strategy of plug-in hybrid electric bus for city-bus route,” *IEEE Transactions on Vehicular Technology*, vol. 64, no. 7, pp. 2792–2803, 2015.
- [5] L. Li, B. Yan, C. Yang, Y. Zhang, Z. Chen, and G. Jiang, “Application oriented stochastic energy management for plug-in hybrid electric bus with AMT,” *IEEE Transactions on Vehicular Technology*, vol. accepted for publication, 2016.
- [6] C. Tessum, J. Hill, and J. Marshall, “Life cycle air quality impacts of conventional and alternative light-duty transportation in the United States,” *PNAS*, vol. 111, no. 52, pp. 18490–18495, 2014.
- [7] Y. Gu, F. Häusler, W. Griggs, E. Crisostomi, and R. Shorten, “Smart procurement of naturally generated energy (SPONGE) for PHEVs,” *International Journal of Control*, vol. 89, no. 7, pp. 1467–1480, 2016.
- [8] A. Schlote, F. Häusler, T. Hecker, A. Bergmann, E. Crisostomi, I. Radusch, and R. Shorten, “Cooperative regulation and trading of emissions using plug-in hybrid vehicles,” *IEEE Transactions on Intelligent Transportation Systems*, vol. 14, no. 3, pp. 1572–1585, 2013.
- [9] G. B. Dantzig, “Discrete-variable extremum problems,” *Operations research*, vol. 5, no. 2, pp. 266–288, 1957.
- [10] I. Adler and R. D. Monteiro, “A geometric view of parametric linear programming,” *Algorithmica*, vol. 8, no. 1-6, pp. 161–176, 1992.

⁵http://www.nrel.gov/electricity/transmission/resource_forecasting.html


- [11] 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.
- [12] F. Wirth, S. Stuedli, J. Y. Yu, M. Corless, and R. Shorten, "Nonhomogeneous place-dependent markov chains, unsynchronised AIMD, and network utility maximization," *arXiv preprint arXiv:1404.5064*, 2014.
- [13] R. Rejaie, M. Handley, and D. Estrin, "Rap: An end-to-end rate-based congestion control mechanism for realtime streams in the internet," in *INFOCOM'99. Eighteenth Annual Joint Conference of the IEEE Computer and Communications Societies. Proceedings. IEEE*, vol. 3. IEEE, 1999, pp. 1337–1345.
- [14] M. Corless and R. Shorten, "An ergodic AIMD algorithm with application to high-speed networks," *International Journal of Control*, vol. 85, no. 6, pp. 746–764, 2012.
- [15] Ł. Budzisz, R. Stanojević, R. Shorten, and F. Baker, "A strategy for fair coexistence of loss and delay-based congestion control algorithms," *Communications Letters, IEEE*, vol. 13, no. 7, pp. 555–557, 2009.
- [16] R. Shorten, D. J. Leith, J. Foy, and R. Kilduff, "Analysis and design of AIMD congestion control algorithms in communication networks," *Automatica*, vol. 41, no. 4, pp. 725–730, 2005.
- [17] M. Liu and S. McLoone, "Investigation of AIMD based charging strategies for EVs connected to a low-voltage distribution network," in *Intelligent Computing for Sustainable Energy and Environment*. Springer, 2013, pp. 433–441.
- [18] M. Liu, E. Crisostomi, Y. Gu, and R. Shorten, "Optimal distributed consensus algorithm for fair V2G power dispatch in a microgrid," in *Electric Vehicle Conference (IEVC), 2014 IEEE International*. IEEE, 2014, pp. 1–7.
- [19] E. Crisostomi, M. Liu, M. Raugi, and R. Shorten, "Plug-and-play distributed algorithms for optimized power generation in a microgrid," *Smart Grid, IEEE Transactions on*, vol. 5, no. 4, pp. 2145–2154, 2014.
- [20] S. Stüdl, E. Crisostomi, R. Middleton, and R. Shorten, "Optimal real-time distributed V2G and G2V management of electric vehicles," *International Journal of Control*, vol. 87, no. 6, pp. 1153–1162, 2014.
- [21] M. F. Barnsley, S. G. Demko, J. H. Elton, and J. S. Geronimo, "Invariant measures for markov processes arising from iterated function systems with place-dependent probabilities," in *Annales de l'IHP Probabilités et statistiques*, vol. 24, no. 3, 1988, pp. 367–394.
- [22] M. Corless, C. King, R. Shorten, and F. Wirth, *AIMD Dynamics and Distributed Resource Allocation*. SIAM, 2016.
- [23] D. Krajzewicz, M. Bonert, and P. Wagner, "The open source traffic simulation package sumo," *RoboCup 2006 Infrastructure Simulation Competition*, vol. 1, pp. 1–5, 2006.
- [24] M. Haklay and P. Weber, "Openstreetmap: User-generated street maps," *Pervasive Computing, IEEE*, vol. 7, no. 4, pp. 12–18, 2008.
- [25] A. Wegener, M. Piórkowski, M. Raya, H. Hellbrück, S. Fischer, and J.-P. Hubaux, "TraCI: an interface for coupling road traffic and network simulators," in *11th Communications and Networking Simulation Symposium, Ottawa, Canada, 2008*, pp. 155–163.
- [26] W. M. Griggs, R. H. Ordóñez-Hurtado, E. Crisostomi, F. Hausler, K. Massow, and R. N. Shorten, "A large-scale sumo-based emulation platform," *Intelligent Transportation Systems, IEEE Transactions on*, vol. 16, no. 6, pp. 3050–3059, 2015.
- [27] A. A. Panagopoulos, G. Chalkiadakis, and E. Koutroulis, "Predicting the power output of distributed renewable energy resources within a broad geographical region," *ECAI-2012/PAIS-2012: 20th European Conference on Artificial Intelligence, Prestigious Applications of Intelligent Systems Track*, 2012. [Online]. Available: <http://eprints.soton.ac.uk/341638/>
- [28] B.-M. Hodge, D. Lew, M. Milligan, H. Holtinen, S. Sillanpää, E. Gómez Lázaro, R. Scharff, L. Söder, X. Larsén, G. Giebel, D. Flynn, and J. Dobschinski, "Wind Power Forecasting Error Distributions: An International Comparison," in *11th Annual International Workshop on Large-Scale Integration of Wind Power into Power Systems as well as on Transmission Networks for Offshore Wind Power Plants, Lisbon, Portugal, 2012*.
- [29] J. Zhang, B.-M. Hodge, A. Florita, S. Lu, H. Hamannk, and V. Banunarayanan, "Metrics for Evaluating the Accuracy of Solar Power Forecasting," in *3rd International Workshop on Integration of Solar Power into Power Systems, London, UK, 2013*.
- [30] S. Stüdl, W. Griggs, E. Crisostomi, and R. Shorten, "On optimality criteria for reverse charging of electric vehicles," *IEEE Transactions on Intelligent Transportation Systems*, vol. 15, no. 1, pp. 451–456, 2014.



Joe Naoum-Sawaya is currently an Assistant Professor in Management Science at Ivey Business School. He received a B.E. in Computer Engineering from the American University of Beirut and a M.A.Sc. and Ph.D. in Operations Research from the University of Waterloo. His research interests include large scale optimisation methods for practical problems arising in the industry.



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



Mingming Liu received his B.E. degrees in Electronic Engineering with first class honours from National University of Ireland Maynooth in 2011. He obtained his Ph.D. degree from the Hamilton Institute, Maynooth University in 2015. He is currently a post-doctoral research fellow in University College Dublin, working with Prof. Robert Shorten. His current research interests are nonlinear system dynamics, distributed control techniques, modelling and optimisation in the context of smart grid and smart transportation systems.



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 Professor Shorten is currently Professor of Control Engineering and Decision Science at University College Dublin, and holds a position at IBM Research. 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.