Managing Electric Vehicles in the Smart Grid Using Artificial Intelligence: A Survey

Emmanouil S. Rigas, *Student Member, IEEE,* Sarvapali D. Ramchurn, and Nick Bassiliades, *Member, IEEE*

***Abstract*—Along with the development of Smart Grids, the wide adoption of Electric Vehicles (EVs) is seen as a catalyst to the reduction of CO2 emissions and more intelligent transportation systems. In particular, EVs augment the grid with the ability to store energy at some points in the network and give it back at others and therefore help optimise the use of energy from intermittent renewable energy sources and let users refill their cars in a variety of locations. However, a number of challenges need to be addressed if such benefits are to be achieved. On the one hand, given their limited range and costs involved in charging EV batteries, it is important to design algorithms that will minimise costs while avoid users being stranded. On the other hand, collectives of EVs need to be organized in such a way as to avoid peaks on the grid that may result in high electricity prices and overload local distribution grids. In order to meet such challenges, a number of technological solutions have been proposed. In this paper, we focus on those that utilise artificial intelligence techniques to render EVs and the systems that manage collectives of EVs smarter. In particular, we provide a survey of the literature and identify the commonalities and key differences in the approaches. This allows us to develop a classification of key techniques and benchmarks that can be used to advance the state-of-the art in this space.**

***Index Terms*—AI, electric vehicles, smart grid**

I. INTRODUCTION

FACED with dwindling fossil fuels, and the increasingly negative impact of climate change on society, several countries have instigated national plans to reduce carbon emissions [1]. In particular, the electrification of transport is seen as one of the main pathways to achieve significant reductions in CO_2 emissions. In the last few years EVs have gained ground, and, to date, more than 180 thousand of them have been deployed worldwide. Despite this number corresponding to only 0.02% of all vehicles on the roads, an ambitious target of having over 20 million EVs on the roads by 2020 has been set by the International Energy Agency [2].[1]

In order to ensure that the large-scale deployment of EVs results in a significant reduction of CO_2 emissions, it is important that they are charged using energy from renewable sources (e.g., wind, solar). Crucially, given the intermittency of these sources, mechanisms (e.g., [3], [4]), as part of a Smart Grid [5], need to be developed to ensure the smooth integration of such sources in our energy systems. EVs could potentially help by storing energy when there is a surplus, and feed this energy back to the grid when there is demand for it. [6], [7].

Indeed, the ability of EVs to store energy while being used for transportation [8] represents an enormous potential to make energy systems more efficient. On the one hand, given that vehicles drive only for a small percentage of the day ($4-5\%$ in the US), and a large percentage of the vehicles stay unused in parking lots (90% in the US) [9], and considering the fact that EVs are equipped with large batteries, they could be used as storage devices when parked (i.e., as part of Vehicle-to-Grid (V2G) schemes [6], [10]), and thus dramatically increase the storage capacity of the network. Indeed, studies [10] have shown that if one fourth of vehicles in the US were electric, this would double the current storage capacity of the network. On the other hand, given that large numbers of EVs need to charge on a daily basis, (40% of EV owners in California travel daily further than the range of their fully charged battery [11]) if EVs charge as and when needed, they may overload the network. For this reason, new mechanisms are required to be able to manage the charging of EVs –Grid-to-Vehicle (G2V)– in real time while considering the constraints of the distribution networks within which EVs need to charge. Moreover, EV routing systems should consider the ability of EVs to recuperate energy while braking and/or when driving downhill, and choose routes that fully utilise this ability. By so doing, it may be possible for EVs to charge less often, thus maximising their range, reducing the costs for their owners, and minimising the peaks they cause on energy grids.

Against this background, a number of techniques and mechanisms to manage EVs, either individually or collectively, have been developed [12], [13], [14]. For example, a number of web and mobile-based applications have been developed to provide information to EV drivers about the locations of charging points[2] where available charging slots exist. Moreover, prototype systems for energy efficient routing have been developed,[3,4] while new types of chargers that can fully charge an EV battery in less than an hour are becoming commonplace. Thus, while a number of advances have been made in terms of the physical infrastructure and technologies for EVs, these may not be sufficient to manage the dynamism and uncertainty underlying the behaviour of individual and collectives of EVs. Controlling the activities of EVs will demand algorithms that can solve problems that involve a large number of heterogeneous entities (e.g., EV

E. S. Rigas and N. Bassiliades are with the Department of Informatics, Aristotle University of Thessaloniki, 54124, Thessaloniki, Greece (e-mail: erigas, nbassili@csd.auth.gr).

S. D. Ramchurn is with the AIC Group, School of Electronics and Computer Science, University of Southampton, Southampton, SO17 1BJ, UK (email: sdr1@soton.ac.uk).

[1]https://www.iea.org/

[2]http://ev-charging.com.

[3]http://www.greenav.org.

[4]http://evtripplanner.com.

owners, charging point owners, grid operators), each one having its own goals, needs and incentives (e.g., amount of energy to charge, profit maximisation), while they will operate in highly dynamic environments (e.g., variable number of EVs, variable intentions of the drivers) and having to deal with a number of uncertainties (e.g., future arrival of EVs, future energy demand, energy production from renewable sources). Some of these challenges have recently been tackled by the Artificial Intelligence community, and in this paper we survey the state of the art of such AI approaches in the following EV application issues:

Energy efficient EV routing and range maximisation: algorithms and mechanisms have been developed to route EVs in order to minimise energy loss and maximise energy harvested during a trip. In particular, building upon existing search algorithms, solutions have been developed to adapt to the needs and the characteristics of EVs, so as to take advantage of their energy recuperation ability and maximise the driving range. For example, [15] and [16] propose algorithms for energy efficient EV routing with or without recharging, while [17] provides an algorithm for calculating reachable locations from a certain starting point given an initial battery level. Moreover, [18] enhances the use of supercapacitors with machine learning and data mining techniques to maximise the range of EVs.

Congestion management: algorithms have been designed to manage and control the charging of the EVs, so as to minimise queues at charging points, and the discomfort to the drivers. For example [19] and [20] propose algorithms for routing EVs to charging points where the least congestion exists, considering the preferences and the constraints of the drivers (e.g., final destination, amount of electricity to charge), while [21] presents a heuristic algorithm to place charging points given a certain topology so that an EV is able to travel between any two locations without running out of energy.

Integrating EVs into the Smart Grid: a number of mechanisms have been developed to schedule and control the charging of the EVs (G2V) so that peaks and possible overloads of the electricity network may be avoided, while minimizing electricity cost. Moreover, we also survey approaches that utilise the storage capacity of the EVs (V2G) in order to balance the electricity demand of various locations in the network, or to ease the integration of intermittent renewable energy sources to the grid. For example [22] and [23] propose algorithms that schedule the charging of collectives of EVs considering the needs of the drivers and the limits of the distribution network, while [24] and [25] use price signals in order to incentivize EVs not to charge at locations, or during periods of high demand. Moreover, mechanisms such as [26] and [27], allow aggregations of EVs to bid for electricity in markets in order to minimise cost, while [4] and [28] present mechanisms to manage the integration of renewables into the grid.

In order to clarify the intersections and differences between the above challenges at a conceptual level, we provide an abstract description of the research landscape in Figure I. While we use a tree representation (signifying a delineation between the concepts), it is clear that there are overlaps (e.g., in terms of congestion management) between the different nodes of the tree (which we consider later in the paper – see Section III).

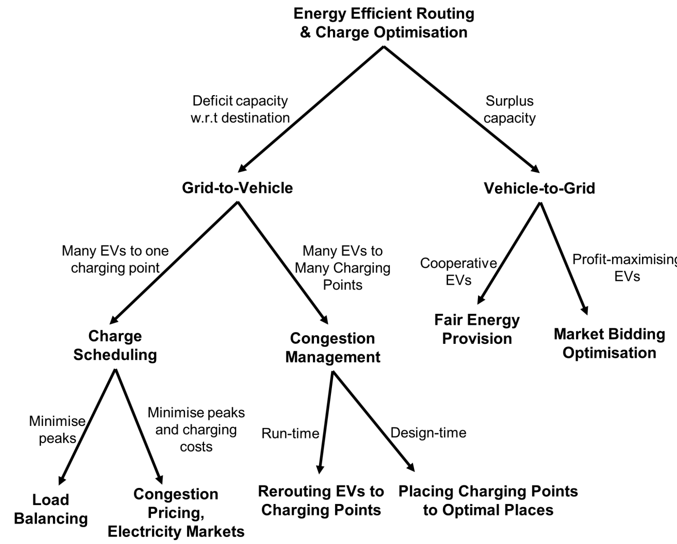

Fig. 1. The electric vehicles research landscape

Thus, from this representation of the research landscape it can be seen that there are different considerations depending on whether the EVs can travel or not based on their battery level (i.e., they need to route to their destination or charge), which in turn, gives rise to challenges for Grid-to-Vehicle and Vehicle-to-Grid systems in terms of load balancing or congestion management among others. Coupled with such issues, is the problem of incentivising EV owners to take certain routes, charge at certain times (e.g., to avoid peaks) or to form part of EV collectives to trade on the energy markets. Finally, the infrastructure also needs to be designed in order to handle large numbers of EVs (e.g., by placing charging points in appropriate places), whichever mechanism is used to charge EVs or sell their spare capacity to the grid.

In what follows, we elaborate on the above challenges. By comparing and contrasting, and by critically evaluating these techniques, we identify areas that need further research, and thus develop a classification of key techniques and benchmarks that can be used to advance the state-of-the art in this space.

The rest of the paper is structured as follows: Section II presents work on energy efficient EV routing and range maximisation, and Section III work on congestion management. Moreover, Section IV presents work on methods and techniques for the efficient integration of EVs to the Smart Grid, both G2V and V2G, while Section V summarizes and discusses a classification scheme of the reviewed papers, identifying areas that need further research.

II. Energy Efficient EV Routing and Range maximisation

Due to the limited range and the long charging times, a number of techniques to optimise the battery usage and to maximise the range of an EV have been developed. Two key research challenges are considered:

1) Energy efficient EV routing (considering or not recharging), where established search algorithms are adapted to the characteristics of EVs so as to calculate routes that utilise the EVs' energy recuperation ability in order to maximise driving range.
2) Battery efficiency maximisation where techniques to maximise the utilisation of the energy stored by an EV are considered.

We elaborate on these challenges in the following subsections (Table I summarises the key papers of this section).

A. Energy Efficient EV Routing

In contrast to conventional vehicle routing that is concerned with minimising travel time and distance travelled, EV routing is concerned with finding 'energy optimal' routes: routes that maximise energy recuperation (through regenerative braking[5]), or routes that pass through charging points that minimise the cost of charging.

Now, approaches to EV Routing typically represent the road network as a weighted directed or undirected graph. In such a graph, the edge weights represent the amount of energy that is needed, or the amount of energy that will be recuperated while an EV is driving over an edge. Whereas in non-EV routing the weights are positive values (e.g., distance or time), in EV routing, energy recuperation induce negative edge costs. This makes it harder to apply standard routing algorithms (e.g., Dijkstra's algorithm) and hence, recent work has looked at algorithms that can take into account such graphs. We elaborate on them below:

1) Energy Efficient EV Routing Without Considering Recharging Events: Using the predefined graph representation and considering energy recuperation, Artmeier et al. [15] and Eisner et al. [29] recently proposed initial solutions for EV routing. In particular, [15] extend the shortest path problem with a set of hard (the battery cannot be discharged below zero) and soft (points where energy could be recuperated but the battery's capacity will be exceeded should be avoided as the extra energy will be lost) constraints making it a special case of a Constrained Shortest Path Problem (CSPP). They proposed a general algorithmic framework for computing trees of shortest paths and present four variations of this framework. These variations differ in the strategy they use to choose the next node to expand in the tree, and they prove their algorithm to run in polynomial time ($O(n^3)$).

The authors in [29] on the other hand, have managed to reduce the time complexity to $O(nlogn+m)$ after an $O(nm)$ preprocessing phase (n is the number of nodes and m the number of edges). In more detail, they model the overcharging (charging beyond the maximum capacity of the battery is not possible) and the battery usage constraints as cost functions on the edges which obey the FIFO [30] property ($O(nm)$). Then, by applying a generalization of the Johnson Potential Shifting Technique [31] to the (partly) negative cost functions they render Dijkstra's algorithm applicable to the shortest path finding problem with negative edge weights ($O(nlogn+m)$). For graphs $G(V,E)$ with constant (negative or positive) edge costs, Johnson's shifting technique tries to determine a potential function $\phi : V \rightarrow \mathbb{R}$ in order to replace the edge costs $const_e$ of an edge $e = (v,w)$ by $con\acute{s}t_e = const_e - \phi(w) + \phi(v)$. If no negative cycles exist, there is a ϕ such as $con\acute{s}t_e \geq 0$ $\forall e \in E$. Note that, in this EV routing scenario no negative cycles exist, as it is not possible for an EV to take a round trip and end up with more energy than the initial one. Moreover, this technique does not affect the structure of the shortest paths, as the potential cost of a certain path, does not depend on the path itself. Johnson's technique also lets the authors use a speed-up strategy for shortest path queries. This strategy is based on the *construction hierarchies* technique [32], which removes nodes in an iterative manner while it perceives the shortest path distances between the remaining nodes.

In contrast, Sachenbacher et al. [33] use the A* search algorithm and they achieve an $O(n^2)$ runtime. This solution uses a detailed vehicle model where the authors consider parameters such as weight and aerodynamic efficiency amongst others, making the results even more applicable to a real world deployment. However, using this representation of the problem, the computation of edge costs is complex and changes dynamically with parameters such as vehicle payload, power demand of auxiliary consumers (e.g., A/C), and battery constraints (treated by dynamically adapting edge costs), therefore making it harder to use preprocessing techniques such as the Johnson algorithm [31]. For this reason, the A* algorithm is chosen as the best solution as it expands the least number of verticies compared to all other search algorithms using the same heuristics.

In terms of evaluating these algorithms, Artmeier et al. [15] show that the Dijkstra and the Bellman-Ford-based variants have reasonable execution times and therefore, practical usability, whereas Eisner et al. [29] compare their algorithm against [15] and prove that it has a better performance in terms of complexity and execution time and can handle bigger graphs (see last column of Table I). Sachenbacher et al. [33] also test their A* algorithm against the two best variations of [15] and prove it to be faster especially when the distance between the source and the destination vertices was short. Note that all of [15], [29] and [33] use real data from the OpenStreetMap[6] and the Altitude Map NASA SRTM[7] projects. Furthermore, [15] have developed a prototype system[8] for energy efficient routing based on this data.

2) Energy Efficient EV Routing with Recharging Events: the works discussed in the previous section do not consider the fact that EVs can recharge en-route. However, recharging en-route is sometimes necessary in order for the EV to be able to reach its final destination, particularly when it has to travel beyond its maximum range.

Sweda and Klabjan [34], considered a setting where no recuperation of energy is performed (edge costs represent

[5]Regenerative braking is a braking technology that can recapture much of the vehicle's kinetic energy and convert it into electricity, so that it can be used to recharge the vehicle's batteries.

[6]http://www.openstreetmap.org.

[7]http://www2.jpl.nasa.gov/srtm/.

[8]http://www.greenav.org.

energy loss), but recharging can take place en-route at some nodes. They model the problem of finding a minimum-cost path for an EV when the vehicle needs to recharge along the way as a dynamic program and they prove that the optimal (EV charging) control and state space (set of nodes the EV can visit while the battery capacity remains within a certain threshold) are discrete under some assumptions. By so doing, standard recursive techniques can be applied to solve the program. The authors prove that in a directed acyclic graph there exists an optimal path, in terms of cost, between any two nodes such that charging (which is modeled to be instantaneous) takes place at every node. Then, by applying a backward recursion[9] algorithm, they decide on the amount of energy that will be charged at each node.

In some cases, it can turn out that the most energy efficient route may be considerably longer than the shortest and/or fastest one. This is because EVs may be able to recuperate energy over longer routes that involve downward slopes. In contrast to [15], [29], [33] and [34] that only focus on calculating the most energy efficient routes, Storand [16] considers additional criteria in defining the value of chosen routes. In more detail, apart from the energy cost of a route, it takes into account time constraints of the driver by trying to balance the travel time against energy consumption and the number of required recharging events. More specifically, they consider two variants:

1) Limiting the number of recharging events, and
2) Minimizing the number of recharging events under a distance constraint.

These optimisation problems are instances of a constrained shortest path problem which they show to be NP-hard. However, they provide preproccesing techniques for fast query answering. Indeed, the authors test their algorithm on graphs based on road networks in Germany (using the OpenStreetMap[6] and the Altitude Map NASA SRTM[7]) and it is shown to compute solutions for networks with 5M nodes in less than 20 msecs.

Finally, Storand and Funke [17] address the problem of EV routing with the goal of finding which destinations are reachable from a certain location based on the current battery level of the EV and the availability of charging stations or battery swap stations.[10] This information is very important for EV drivers when it comes to planning their journey and therefore, reduces their likelihood of running out of energy. The authors introduce the notion of EV-reachable (going from point A to B) and strongly EV-connected (going from point A to B and back to A) paths and they prove that their algorithm for calculating these paths has an $O(nlogn + m)$ time complexity. Moreover, they model battery swap stations as nodes that instantaneously give a certain amount of energy to an EV when it passes through them. Despite this being a simple model of battery swap stations, as it does not consider the delays incurred in queues at charging stations, this is the only model that considers battery swapping and not only recharging. The authors evaluate their algorithm in a similar setting to [16] and it is shown to compute solutions for networks with 5M nodes and 200 battery swap stations in under 0.2 seconds.

[9]In order to solve a problem of size N, you assume a solution of size N - 1 and then you use this solution to solve the problem of size N.

[10]In a battery swap station the battery is not recharged but instead it is replaced by an already charged one. Such stations can reduce battery reloading time significantly, but they come with a high cost [35].

Now, the techniques discussed above typically ignore the physics of electric batteries that dictate how much energy can be stored or extracted from a battery and how these affect its lifetime. Hence, in the next section, we provide a short discussion of existing techniques that specifically focus on this aspect.

B. *Battery Efficiency Maximisation*

The trend in energy storage technology for EVs (to maximise lifetime and allow for fast charging) is to use a chemical battery in conjunction with supercapacitors [18]. In a supercapacitor, energy is stored electrostatically on the surface of the material, and does not involve chemical reactions. Supercapacitors can be charged quickly, and they can last for millions of charge-discharge cycles, but they have a relatively low energy density [36]. Supercapacitors can discharge a large current at short notice (e.g., when accelerating), thus reducing the stress on the chemical battery. When no current is drawn from the supercapacitor, it may then recharge, at a slower rate, from the attached battery. By so doing, the supercapacitor acts as a buffer for sudden energy demands on the battery. In such systems, the management of the charging and discharging of the capacitors and the energy flow from the capacitors to the battery needs to be optimised in order to maximise battery lifetime. To this end, [18] develop a stochastic planning algorithm using dynamic programming. Their algorithm has quadratic complexity in the discredited capacity levels of the supercapacitor, but requires an accurate prediction of future energy requirements. To this end, they apply machine learning techniques to predict future energy consumption (using data about commuter trips collected across the United States) and use such predictions within a Markov Decision Process to determine a charging/discharging policy. The authors evaluated their policy against the policies taking part at the Chargecar[11] algorithmic challenge and show that it marginally outperforms the previous best algorithm designed for this problem. (see details in [37]).

The techniques presented so far focus on individual EV routing, ignoring the effects of the collective behaviour of EVs might have on the charging network. We elaborate on this in the next section.

III. Congestion Management

Existing work addresses congestion in EV systems in two main ways. First, congestion can be managed by individually guiding EVs to charging points in order to minimise queues. Second, charging points (and the associated charging slots) may be placed at specific locations to distribute the load evenly across the routes usually taken by EVs. In both cases, most existing work represent the road network as a (weighted)

[11]http://www.chargecar.org.

TABLE I
CLASSIFICATION OF PAPERS - EV ROUTING AND RANGE MAXIMISATION

cit.	Specific Goal	Problem Solving Technique	Control Scheme	Complexity	Evaluation Method
[15]	Energy efficient routing with energy recuperation	Graph-based Search	Decentralised	$O(n^3)$	Theoretical evaluation & Simulation dataset road network OpenStreetMap and Altitude Map NASA SRTM (0.78M nodes and 1.7M edges)
[29]	Energy efficient routing with energy recuperation	Graph-based Search	Decentralised	$O(nlogn + m)$ with $O(nm)$ preprocessing	As above (5.5M nodes and 11.7M edges)
[33]	Energy efficient routing with energy recuperation	Graph-based Search	Decentralised	$O(n^2)$	As above (2M nodes and 4.9M edges)
[34]	Energy efficient routing without recuperation, with en-route charging	Mathematical Programming	Decentralised	Polynomial	Theoretical evaluation
[16]	Energy efficient routing considering comfort factors	Graph-based Search	Decentralised	NP-hard	Theoretical evaluation & Simulation dataset road network OpenStreetMap and Altitude Map NASA SRTM (5.5M nodes and 11.7M edges)
[17]	Finding of reachable locations from a certain location	Graph-based Search	Decentralised	$O(nlogn + m)$	Same as the previous
[18]	Battery efficiency maximisation	Mathematical Programming, Machine Learning	Decentralised	$O(n^2)$	Theoretical evaluation & Participation in ChargeCar competition based on a dataset of 1984 EV trips

directed or undirected graph. Moreover, while in the first area AI techniques such as stochastic optimisation, utility-based agent coordination, or mathematical programming are utilised, in the second area, graph-based search is proven to be NP-hard and heuristic optimisation algorithms are used instead. (Table II summarises the key papers of this section).

A. Routing EVs to Minimise Congestion

Initial work by de Weerdt et al. [19] proposed a navigation system that can predict congestion at charging stations and suggests the most efficient route, in terms of travel time, but not energy efficiency, to its user. In order to achieve this, they proposed an Intention Aware Routing System (IARS) which is implemented as a software agent. The agent exchanges intentions with other agents, where the intentions are probabilistic information about which stations the EVs will go to and when, thus making it possible for each agent to predict congestion levels. Note that their system can route EVs using only historical data and can update the routes online as more accurate information about EVs' intentions become available. The authors tested their algorithm (assuming all cars can fully charge in 30 mins) against other similar approaches which do not use intentions, and empirically proved that it outperforms them in terms of waiting time by up to 80%.

Now, a key assumption in [19] is that the communication between EVs and charging points is reliable, if not continuous. Instead, Qin and Zhang [20], propose a distributed charging scheduling algorithm where EVs communicate only with charging points, but are not able to update their decision en-route. In more detail, the authors consider a setting of a highway network with charging points at the exits, modelled as a graph. For every EV that needs charging, the set of charging points that exist between its current position and its final destination is calculated. Based on the preferences of the owner of the EV, every charging point from this set reports the minimum waiting time (queuing and charging) that can be achieved, and the EV selects the one with the minimum waiting time. The waiting time for the selected charging point is then compared to the waiting times for the rest of the charging points, and based on past data, a probability of an EV driver deviating from the plan and going to another charging point is calculated. These probabilities are then used for more accurate predictions on future waiting times. The authors evaluated their algorithm (assuming EVs minimise distance travelled) in a simple setting mostly using synthetic data, and show that it is able to achieve solutions (waiting times) that are up to less than 10% of the optimal.

While [19] and [20] consider only time as a cost to the system, Rigas et al. [38] instead introduce pricing mechanisms as a method to reduce congestion at charging points. Under their pricing scheme, EVs (modelled as agents with utility functions capturing time and monetary costs) are incentivised to avoid charging at congested charging points. Thus, using prices reported by charging points over time, EVs book charging slots at the charging point that minimises their delays (e.g., walking from a charging point to their final destination) and provides enough charge to route to its final destination. Bessler and Gronbaek [39] also work on a model similar to [38], but they consider charging points that are not, necessarily, close to the drivers' final destination and therefore require drivers to use other means of transport (including walking as in [38]). This approach has the advantage that the set of feasible charging points can be larger, compared to one where no multi-modal transportation is taken into consideration, and therefore, congestion at charging points can be more efficiently handled. Indeed, the authors test their algorithm on a road network in Wien, Austria, and prove that they can achieve up to 75% more charging options compared to a setting where no multi-modal routes are taken into account.

We next discuss the placement of charging points as an alternative mechanism to reduce congestion.

B. Charging Point Placement

Initial work by Storand and Funke [21] address the problem of charging point placement on a road network under the constraint that the energy spent for return trips between any pair of nodes is never larger than an EV's battery capacity. The problem is shown to be NP-Hard and heuristic solutions are developed and tested on road networks from Germany (using data from OpenStreetMap and SRTM). Similarly, Lam et al. [40] propose a greedy algorithm which, compared to an optimal solution that uses Mixed-Integer programming techniques (using synthetic data), is faster while producing solutions up to 5% from the optimal, but in considerably lower computation time. Unfortunately, both of these approaches do not guarantee that detours will not be imposed on the EV drivers. However, recent work by Funke et al. [41] investigate methods for placing charging points, where, given any shortest path between any two nodes, there are enough stations for an EV to recover enough energy to continue its journey (assuming it starts with a fully charged battery). In more detail, this problem is defined as the EV Shortest Path Cover problem (SPC) and is modelled as an instance of the Hitting Set Problem [42].[12] Moreover, they adapt existing (for the Hitting Set Problem) heuristic algorithms to solve the SPC problem and prove that near optimal results within a factor of $O(logn)$ of the optimum (n being the number of nodes in the network) can be achieved.

In general, the efficient placement of charging points is a necessary but not sufficient condition for the mainstream adoption of EVs. Along with the placement of such charging points, it is important to consider the peaks in demand they can individually handle (by installing enough charging slots) due to EVs that arrive in different numbers at different times of the day. Initial work by Bayram et al. [43], introduces the concept of *effective power* which is a deterministic quantity related to the aggregated stochastic demand for electricity at an EV charging station. The aim of this work is to minimise the electric power delivered to the station, as well as the number of charging slots that must be installed in the station, while the EVs that remain uncharged are kept to a minimum. The authors use predictions of the actual demand for electricity, as a percentage of the maximum demand, given a fixed number of charging slots. The authors evaluate their methodology using numerical examples and mathematically prove that it can lead to up to 40% of savings in the total required power, while the infrastructure cost can be reduced by up to around 30%, while 10% of EVs are not able to charge.

The solutions discussed in this section point to the fact that the load induced by EVs at different charging points will stress not only the transportation network but also the electricity network that delivers energy to each of the charging points. Alternatively, however, EVs could be used to power local grids to satisfy demand (from any consumer, including EVs) as part of a smart grid that permits such serendipitous charging and discharging events. Hence, in the next section, we elaborate on the integration of EVs into the Smart Grid.

[12]Given a set system (U, S) with U being a universe of elements, and S a collection of subsets of U, the goal is to find a minimum cardinality subset $L \subseteq S$ such that each set S' is hit by at least one element in L, i.e., $\forall S' \in S : L \cap S' \neq 0$.

IV. Integrating EVs into the Smart grid

The IEEE Intelligent System Applications (ISA) subcommittee[13] has recently recognised the usefulness of AI approaches in solving key power systems challenges involved in balancing loads on the electricity grid. Hence, in this section we discuss a number of AI-based solutions that have been developed to address both Grid-to-Vehicle (G2V) and Vehicle-to-Grid (V2G) problems. We discuss these solutions in turn. (Tables III- VI summarise the key papers of this section).

A. Grid to Vehicle (G2V)

Here we focus on solutions that address the scheduling of charging cycles to minimise the load on transformers and distribution lines. We identified three main categories of solutions: (i) load-balancing: techniques to predict future loads and schedule charging cycles to minimise possible peaks, (ii) congestion pricing: financial incentives used to manage demand dynamically, and (iii) electricity markets: allow competing energy providers and consumers to converge on efficient allocations of energy that minimise peaks in the network. In all of these solutions we find commonalities in the AI techniques used, ranging from agent-based solutions to electronic auctions. In particular, in the first category works typically aim to optimise (minimise) either cost (for the electricity network, and/or for the EVs), or load on the network, or both using mathematical programming. In the second category, individuals or collectives of EVs (formulated as agents) minimise charging cost using agent-based coordination techniques that also consider load on the grid, and in a few cases apply game theoretic concepts. Finally, in the third category individuals or collectives of EVs optimise their participation in electronic auctions and try to minimise charging cost. Here, works typically use either mathematical programming, or utility-based agent coordination combined with concepts from auctions theory, and in some cases they also use mechanism design. We elaborate on each of these categories in what follows.

1) Load Balancing: In [44], the authors present a simple analysis of the impact the uncontrolled charging of Plug-in Hybrid Electric Vehicles (PHEVs) can have on the distribution network and develop a dynamic programming solution that computes the charging schedule for individual EVs across a network in order to minimise peaks and carbon taxes. They do so using predictions of EV consumption in future time slots where such predictions are liable to uncertainty. Their algorithm is shown (when applied to an IEEE 34-node test grid using load profiles from a Belgian distribution network) to reduce losses by up to 2.2% and power deviations by up to 3%, in spite of errors in predicting future consumption from EVs. In a similar vein, Anh et al. [45] address the same problem with a decentralised algorithm where each EV computes its own schedule (but assuming no prediction error) that is shown to achieve near-optimal performance (using data from the Detroit area). Similar techniques have also been proposed in [12], and evaluated in a Portuguese electricity network. As

[13]http://sites.ieee.org/pes-iss/.

TABLE II
CLASSIFICATION OF PAPERS - CONGESTION MANAGEMENT

cit.	Specific Goal	Problem Solving Technique	Control Scheme	Evaluation Method
[19]	Navigation system predicting congestion at charging points targeting minimisation of journey time	Stochastic optimisation	Decentralised	Evaluation through simulation against similar algorithms - Results up to 50 EVs are presented (the authors claim it can work up to 1000)
[20]	Managing congestion at charging points to minimise charging and waiting time	Graph-based search, and Mathematical programming, and Probabilities theory	Decentralised	Evaluation through simulation - $300km \times 300km$ road network, with 100 charging stations and 5 types of of EVs
[38]	Managing congestion at Charging Points to minimise load	Graph-based search, and Utility-based agent coordination	Decentralised	Simulation against an optimal algorithm (Locations of car parks in a UK city taken from Google maps and up to 2000 EVs)
[39]	Selecting charging points under a multi-modal transportation setting to minimise load	Graph-based search and mathematical programming	Decentralised	Simulation with dataset from a road network in Austria taken from OpenStreetMap
[21], [41]	Optimal Placement of Charging Points	Graph-based search, and heuristic optimisation	Centralised	Theoretical evaluation & Simulation dataset road network OpenStreetMap and Altitude Map NASA SRTM (15.015.877 nodes and 30.771.648 edges)
[40]	Optimal Placement of Charging Points	Graph-based search, and mathematical programming, and heuristic optimisation	Centralised	Theoretical evaluation & Simulation (up to 200 nodes)
[43]	Optimisation of the number of charging slots at a charging point, and of the amount of energy provided to it	Probabilities theory	Centralised	Simulation through numerical examples (50K EVs)

opposed to the previous works where large numbers of EVs are managed, Halvgaard et al. [46] develop an economic Model Predictive Control (MPC) method to minimise the cost of electricity for a single EV. They propose a dynamic programming algorithm to calculate an optimal charging plan which achieves up to 60% cost savings as opposed to uncontrolled charging when evaluated in a setting using real data taken from the Danish distribution network.

Vandael et al. [22] also propose a decentralised algorithm but specifically consider transformer limits and imbalance costs which are caused by unpredictable changes in production and consumption. By modelling EVs, transformers and BRPs (Balancing Responsible Party)[14] as agents that express their individual requirements (charging needs and departure time for EVs, power limits for transformers, and predicted loads for BRPs), they can coordinate the schedule of charging EVs. In particular, they develop an approach that distributes imbalances across the network and this is shown to reduce imbalances by 44% (on a dataset from the Belgian distribution network). In the same vein, Li et al. [23] propose an online decentralised algorithm that myopically (i.e., with no predictions of future system states) schedules charging cycles using only the present power system state. Hence, it is more robust than solutions that rely on, possibly erroneous, predictions of future system states (e.g., [45], [22]). They achieve coordinated charging cycles using a charging reference signal that is computed by an aggregator (i.e., the utility company) that aims to maximise the SoC (State of Charge) of the vehicles while it is penalized based on the load at each time point. The authors prove, both theoretically and empirically, using data from a Californian distribution network and simulating EV charging over a long time period, that this algorithm asymptotically matches a static optimal one, and also show that it is robust to forecasting errors. However, they assume that each EV is available to charge for more than the minimum needed time.

In contrast to the papers presented so far, the work proposed by Bayram et al. [47] assumes a large number of charging points, each of them having pre-ordered a certain amount of energy. In this setting, a centralised mechanism utilises mathematical programming techniques to optimally allocate the energy to EVs (based on individual preferences on charging rate, and amount of energy needed), so as to maximise the social welfare by serving the maximum number of EVs. The authors evaluate the mechanism in a setting where both selfish (want to charge at the nearest charging point), and cooperative EVs exist using data regarding traffic traces from the Seattle area, and prove that, up to 10% of energy savings can be achieved, while only 5% of EVs remain unserviced.

Now, the above solutions typically ignore the fact that ultimately, EVs may be powered using uncontrollable renewable energy sources (e.g., wind or solar). In turn, [28] propose dynamic programming algorithms that schedule the charging of EVs according to the availability of energy while guaranteeing the intended journeys can be completed (assuming knowledge of future traffic conditions). They also show that their solutions can adapt to fluctuations in energy generation from renewable sources and that this allows up to 61% penetration of EVs (using network and energy generation data from Portugal).

Note that the algorithms by [28] are purely reactive and do not try to model the uncertainty in energy production. In contrast, [3] develop a probabilistic model for wind forecasting (based on [49]) and additionally consider network constraints. Thus, they solve an Optimal Power Flow problem (to minimise system generation costs) that guarantees that demand is met by supply while respecting thermal limits on distribution lines. By modelling collectives of EVs at individual nodes as one large battery, their charging algorithm is shown to be robust

[14]The electricity grid consists of the transmission grid and the distribution grid. The transmission grid carries electricity from the producers to the distribution grid which then transfers electricity to the individual customers. The transmission system operator (TSO) keeps a balance between supply and demand. In order to achieve this, predictions of the energy that will be injected to, or withdrawn from each access point of the transmission network must be made. The predicted load schedule of the consumers and/or producers behind its access point is provided by the BRP that exists at each access point.

TABLE III
CLASSIFICATION OF PAPERS - INTEGRATING EVS INTO THE SMART GRID - *Load Balancing*

cit.	Specific Goal	Problem Solving Technique	Control Scheme	Evaluation Method
[44]	Load balancing to minimise peaks and carbon taxes using predictions (with possible errors) on future EV consumption	Mathematical programming	Centralised	Simulation with data from household load profiles and predicted EV penetration from Belgium and the IEEE 34-node test grid
[45]	Coordination of power generation and charging to reduce costs and carbon emissions using predictions (with no errors) on future power reserves	Mathematical programming	Decentralised	Simulation with data from DTE Energy in Michigan, USA (2M EVs)
[12]	Load balancing to minimise peaks using predictions (with no errors) on future loads	Mathematical Programming	Centralised	Simulation with data from a medium voltage distribution grid in Portugal (6.500 EVs)
[46]	Charging scheduling of a single EV to minimise costs using predictions (with no errors) on future loads	Mathematical programming	Decentralised	Simulation with data from a distribution grid from Denmark, where a 24h prediction on electricity consumption is assumed
[22]	Charging scheduling of a set of PHEVs to minimise peaks assuming no predictions on future loads are available	Utility-based agent coordination	Decentralised	Simulation with consumption profiles data from 200 households obtained from the Belgian distribution grid provider Infrax
[23]	Charging scheduling of a set of PHEVs to minimise peaks assuming no predictions on future loads are available	Mathematical programming	Decentralised	Simulation on IEEE 37-bus and IEEE 123-bus system and using data from residential distribution networks in southern California
[47]	Optimal allocation of a fixed amount of energy to a set of EVs to minimise peaks and maximise social welfare using stochastic predictions on future loads	Mathematical programming	Centralised	Simulation using data on traffic traces from the Seattle area (1M EVs)
[28]	Scheduling of EVs according to the availability of energy for load balancing considering Renewables, but no prediction of their future energy production	Mathematical programming	Centralised	Simulation using data derived from the low voltage distribution grid, and data on power generation from renewables from Portugal for up to 230 EVs
[3]	Charging scheduling of EVs for load balancing and system generation costs reduction using a probabilistic model to forecast production from renewables	Mathematical programming	Centralised	Simulation based on data from the Swiss transmition operator Swissgrid, driving patterns from an agent-based transport simulator (MATSim [48]) and wind output data forecast from the National Renewable Energy Laboratory (1M EVs)

to errors in wind prediction, but a trade-off between flexibility and cost minimisation is identified.

We next discuss congestion pricing approaches to managing EV charging that also consider constraints imposed by the distribution network.

2) Congestion Pricing: Sundstrom and Binding [24], propose algorithms that price energy consumption according to the time of day (i.e., time of use tariffs) under the assumption that demand will be time dependent. Thus, they develop an EV charging scheduling algorithm, using MIP (Mixed Integer Programming), that uses these prices and power constraints and thermal limits of the network. Taking real data from distribution grids (in Denmark and Germany) and assuming that a single wind-powered electricity generator exists, they show that with their solution only 0.04% of the grid is overloaded by more than 10%, compared to purely myopic charging (i.e., as and when needed) where up to 4% of the grid is overloaded by more than 10%.

While [24] assume energy demands are centrally known and can be used for scheduling (and hence less robust to failures), [50] develop a decentralised solution where EVs react to a price signal broadcast by the utility a day-ahead. In more detail, two alternative tariffs are explored, one where the same price profile applies system-wide, and another where different prices can be defined at different nodes. By shifting their charging cycles to minimise cost (solving a constrained Optimal Power Flow problem), the EVs also reduce congestion on the distribution network. Crucially, they show that their decentralised algorithm produces solutions that are up to 97% of a centralised algorithm (with known EV profiles and schedules). Echoing results in another study [51], they show that their solution mainly balances schedules at individual nodes rather than across the network. Rigas et al. [38] and Karfopoulos and Hatziargyriou [52] present solutions to this problem. In particular, [38] applies congestion pricing across nodes in the network using pricing functions that are demand-dependent (at each node rather than across the network). By minimising charging costs (and time the drivers spent waiting and/or walking to their actual destination), the EVs (acting as self-interested agents), automatically schedule themselves to minimise congestion across the network but also at individual charging points. Thus they are able to show (using data of car park locations in Southampton, UK), that their agent-based congestion management algorithm is able to scale to thousands of agents, producing good-enough solutions, compared to a centralised scheme that assumes complete information about the future arrivals of EVs. Moreover, [52] formulate the problem as a single-objective, non-cooperative, dynamic game and apply a number of price signals across a set of regions of a distribution network. The authors prove that a Nash-equilibrium can be achieved under the assumption that the EV agents are (weakly) coupled (they take into consideration the strategies of others when deciding on their charging). Moreover, by simulating their mechanism in a setting using data from a distribution network in Greece, they show that as opposed to uncoupled agents, weakly coupled ones can achieve up to 13% reduction on the maximum line load. Note however, that real time pricing comes with a higher infrastructure cost compared to time of use pricing [53].

In contrast to [50], [38] and [52], Bayram et al. [54] propose the use of fixed prices up to a certain number of EVs that charge at one charging point, and once this threshold is exceeded, congestion pricing is used in order to incentivize EVs to charge at other points. By so doing, they are able to reduce the need to continuously communicate prices to EVs (as in [38] for example). In particular, their solution focuses on maximising revenue for the operator while minimising the number of EVs priced out of the market. However, as their mechanism is only tested on synthetic data, it is unclear whether such results would port to situations where EV arrival rates are unpredictable.

In contrast to the above, a number of studies [25], [55] use Game theoretic analysis to study the performance of the system when EVs and charging points adopt simple strategies to minimise their individual cost. In particular, they cast the problem as a game and attempt to predict the Nash equilibrium of the game. Specifically, while [25] shows that EVs competing for charging slots across a network would end up minimising congestion costs across the network, [55] instead shows that when charging points belong to different stakeholders, despite the competition between them, EVs can easily be exploited if they simply go to the nearest charging point (rather than choosing the cheapest one).

Apart from the above approaches that only price charging slots, a number of approaches have recently studied how charging rates can be throttled using congestion pricing. In particular, we note the work of [56] that applies Internet congestion control techniques to throttle charging rates at different points in the network. They further decentralise their solution using Lagrangian decomposition techniques. While they make some significant assumptions (e.g., residential load is constant and a fixed number of EVs are connected to chargers), it is interesting to see how such congestion management techniques that are popular in communication networks can be transferred to electricity networks.

Using more traditional agent-based negotiation techniques, Gan et al. [13], implement an iterative procedure to allow EVs to negotiate the charging rate (at different time points) with a utility company (that broadcasts a price signal to control charging). Crucially, they show that, should the charging characteristics of all EVs are known, an optimal solution is reached in a decentralised fashion. They further validate their approach empirically and show (using data from a Californian distribution network) that it impressively outperforms a standard benchmark for this domain [25].

In the settings we have discussed so far, EVs do not have the option to negotiate on the congestion price (as this is set by the utility company or charging point owners). Instead, in the next subsection, we discuss market-based price setting techniques.

3) Electricity Markets: Initial work by Caramanis and Foster [26], investigate market-based control techniques for load balancing and to provide regulation services that allow renewable energy sources to be integrated.[15] Specifically, they assume that EVs join an aggregator that directly participates in day-ahead[16] electricity markets where different generators (including renewable) participate. Crucially, they develop a bidding strategy, using stochastic dynamic programming techniques, for the aggregator to account for uncertain demand from the EVs while maximising regulation service revenues (by efficiently absorbing unpredictable surges of wind energy into the EV batteries). In [57] they further develop a new bidding strategy (using mathematical programming) for the EV aggregator to operate in hour-ahead (real time) markets[17] and show (using data from US power exchange) that it outperforms typical benchmarks by up to 15% (in cost reduction for the EVs).

In the same vein, González Vayá and Andersson [58] propose a bidding strategy, using MIP techniques, for a day-ahead market having as an objective to minimise charging costs, while satisfying the EVs' demand for electricity. The setting is studied over a period of time, thus making it an intertemporal problem, and therefore a multi-period optimisation is used. In addition to this, a single-period optimisation is carried out in order to allocate energy to individual vehicles. In [27], the same authors go a step further, as the bidding strategy is modelled as a two-level problem (implemented as a mixed-integer linear program), where the upper-level is in charge of minimising the aggregator's charging cost (a set of EVs is represented by an aggregator), while the lower-level represents the market clearing (the price on which electricity is sold), where the bids of other participants are not known in advance. The bidding strategy is evaluated based on historic data on electricity pricing from Germany and Austria, and driving patterns taken from the MATSim [48] simulator based on Swiss transport data, and the results show an up to 37% reduction costs. Additionally, Yang et al. [59] propose a centralised charging scheduling framework which also considers the load mismatch risk between the day-ahead and the real-time market.[18] The framework is based on the day-ahead prices and on statistical information of the EVs' driving patterns and the risk-aware day-ahead scheduling is modelled as a two stage stochastic linear problem which is solved using the L-shaped method [60]. Using day-ahead electricity prices from a distribution network provider in New England and random vehicle travel activities to simulate a realistic scenario, they evaluate their risk-aware algorithm and show that it reduces the total cost by up to 20%, while it also reduces peaks.

[15]Regulation service corrects for short-term changes in electricity use that might affect the stability of the power system. It helps match generation and load and adjusts generation output to maintain the desired frequency. Energy from renewable sources come with a certain amount of intermittency and, therefor, regulation service might need to be increased by up to 20%.

[16]Day-Ahead Market is a forward market in which prices are calculated for the next operating day based on generation offers, demand bids and scheduled bilateral transactions.

[17]Real-Time Market is a spot market in which current prices are calculated at five-minute intervals based on actual grid operating conditions.

[18]An entity (e.g., an EV aggregator) buys electricity in the day ahead market based on predictions on the next day's consumption. Then in the real time market it can buy (or sell) electricity to cover the actual demand. However, real-time markets are more expensive compared to day-ahead ones, and therefore, the amount of energy bought in the real-time market must be minimised.

TABLE IV
CLASSIFICATION OF PAPERS - INTEGRATING EVS INTO THE SMART GRID - *Congestion Pricing*

cit.	Specific Goal	Problem Solving Technique	Control Scheme	Evaluation Method
[24]	Charging scheduling through time of use pricing to minimise peaks and cost	Mathematical programming	Centralised	Simulation using data from a distribution grid in Denmark correlated with consumer profiles from Germany (3500 EVs)
[50]	Charging scheduling through price tariffs broadcasted a day ahead to minimise peaks and cost	Utility based agent coordination, and mathematical programming	Decentralised	Simulation based on a model of the Swiss transmition network with 191 nodes, 246 lines and 21 transformers based on data from Swissgrid (1M EVs)
[38]	Charging point selection and charging scheduling to minimise congestion, peaks and costs using dynamic pricing	Utility-based agent coordination	Decentralised	Simulation using real car parks locations from a UK city (2000 EVs)
[52]	Charging scheduling across a set of regions using price tariffs and game theoretic analysis to minimise peaks and cost	Utility-based agent coordination, and game theoretic analysis.	Decentralised	Theoretical evaluation and simulation using data from a distribution network in Greece (1200 EVs)
[54]	Charging scheduling using fixed prices up to a number of EVs and congestion pricing afterwords to minimise peaks and cost	Utility based agent coordination.	Decentralised	Simulation using synthetic data
[25]	Minimisation of congestion costs through EV competition for charging slots across a network	Utility based agent coordination, and game theoretic analysis	Decentralised	Theoretical evaluation and simulation using real data on a distribution network from the Midwest Independent System Operator (USA) (10M EVs)
[55]	Charging scheduling over a set of charging points to minimise peaks and cost using game theoretic analysis	Utility based agent coordination, and game theoretic analysis	Decentralised	Theoretical evaluation and simulation using synthetic data (1350 EVs)
[56]	Scheduling time and rate of charge to minimise peaks and costs using congestion pricing	Utility-based agent coordination, and mathematical programming	Decentralised	Theoretical evaluation and simulation based on an IEEE 13-bus test feeder
[13]	Scheduling time and rate of charge to minimise peaks and costs using price tariffs	Utility-based agent coordination, and mathematical programming	Decentralised	Simulation with real data from a distribution network in USA

Aside from mechanisms that allow collectives of EVs to participate in electricity markets, new mechanisms have recently been developed to manage congestion at a local level, while in all these mechanisms the incentives and allocations are set to ensure the agents have, as their best strategy, to reveal their preferences for charging times and reserve prices. In particular, we note the work of [61], [62] and [63] that use mechanism design techniques to incentivise self-interested EV agents (that hold their owners utility function) to book charging slots in order to achieve system-wide objectives (e.g., cost reduction, network stability). Specifically, [61] propose a mechanism for allocating electric power units to self-interested agents, aiming to maximise the social welfare of the agents. In order to generate efficient electricity unit allocation decisions, the authors use a modified version of the Consensus algorithm [64]. Moreover, they use the concept of pre-commitment (the mechanism pledges that it will charge the EV by its departure time, but has the flexibility to choose when and at what rate the charging will take place), they prove that their mechanism incentivizes truthful reporting of the preferences of the agents. Assuming that one charging point exists, and that probabilistic knowledge on future EV arrivals exists, they evaluate their mechanism within a scenario involving 100 EVs and show that their mechanism achieves an 93% or more of that of an offline optimal mechanism. Instead, in [62], agents state time windows within which they will be available to charge, and bid for units of electricity in a periodic multi-unit auction (one auction per time step). In order to ensure truthfulness, the authors developed a mechanism that occasionally leaves units of electricity unallocated (burned), even if there is demand for them. These units are burned either at the time of allocation or at the time of departure of the agent. Moreover, in [63], a two-sided market (between charging points and EVs) is proposed. In particular, the agents report their preferences and their value for the electricity and the charging points report their availability and costs, and then they are allocated the slot that maximises the difference between their value and the sellers' cost. Both [62] and [63] show that their mechanisms can achieve performance up to 95% of the optimal.

Finally, Tushar et al. [65] cast the problem of the provision of electricity to collectives of EVs from a smart electricity grid in a distributed manner as a Stackelberg game[19] [66]. In particular, the electricity grid acts as a leader and aims to maximise its revenues by setting prices for a certain amount of electricity available for EV charging. In turn, the collectives of the EVs act as followers and need to decide on their charging strategies so as to optimise a tradeoff between the benefit from battery charging and the associated cost. The authors prove that with the use of variational inequalities,[20] the proposed game reaches a socially optimal Stackelberg equilibrium. In this state, the grid optimises its price, while the EVs choose their equilibrium strategies. They show that the equilibrium reached in the game results a higher average utility than typical

[19]The Stackelberg leadership model is a strategic game where two (or more) players (firms) offer an undifferentiated product with known demand. Players have to compete by choosing the amount of output to produce, but one of them goes first (leader). The other player(s) (follower) observe what amount player 1 has chosen, and choose their amount accordingly to maximise profits. In this setting, player 1 knows that player 2 will follow this strategy since it can rely on the other player's economic rationality. In a Stackenberg model, equilibrium is reached when player 1 pre-emptively expands output and secures larger profits.

[20]Given $X \subseteq \Re^n$ and $F : \Re^n \to \Re^n$, the $VI(X, F)$ consists of finding a vector $z^* \in X$ such that $\langle F(z^*), z - z^* \rangle \geq 0$, for all $z \in X$.

benchmarks in this domain [67], [68].

We next turn to mechanisms that permit EVs to sell their stored energy back to the grid as part of V2G programmes.

B. Vehicle-to-Grid (V2G)

EVs' large batteries can, if well managed, become a valuable asset to a smart electricity grid. As discussed earlier (see Section I), the stored energy can be used to smooth out the fluctuating production of electricity from renewable sources. Moreover, the provision of V2G services can potentially be very profitable for EV owners. V2G services can be provided either in an one-to-one basis (each EV will sell its own spare energy to the grid), or in the form of collectives of EVs which act as one entity and trade electricity. Indeed, unless operating through an aggregator, it is impossible for individual EVs to sell V2G services in electricity markets where buyers typically buy energy in Megawatt-hours rather than kilowatt-hours [69].

Now, a number of different AI-based approaches have been developed to manage V2G programmes. For example, some seek to optimise, using mathematical programming, the use of stored energy to cater for low energy production periods from renewables [70], [14]. Others, instead have applied coalition formation techniques,[21] to coalesce EVs into efficient groupings that can make profitable V2G trades [72], [73] and [74]. We elaborate on all these approaches in the following paragraphs.

In terms of using V2G to balance supply against demand, Chatzivasileiadis et al. [70] developed a solution which mimics inertia[22] techniques using battery storage from a fleet of PHEVs, and applied Q-learning in order to learn the optimal controller placement strategy. They provide theoretical results that show that the system reaches a stable state where demand is always balanced by supply, and that costs are reduced as fewer controllers need to be installed. [4] instead, use particle swarm optimisation (PSO)[23] [75] to optimise energy trades (charging and discharging of EVs). In comparison to [70], [4] also consider CO2 emission reduction as a key aspect. V2G service can reduce CO2 emissions by giving energy to the grid when demand is higher than supply, and in this way highly energy consuming reserve power plants (or peaking plants) are not activated. Thus, PSO is used to schedule the charging and/or discharging activities of the vehicles so as cost and emissions to be minimised. The authors also show (using data from [76]) that their mechanism can trade-off emission reduction for cost reduction. The intuition is that when charging cost is low, EVs prefer to charge, and therefore emissions are high, whereas, when charging cost is high, emissions are low as EVs prefer to discharge.

[21]Coalition formation allows groups of autonomous rational agents to form stable (i.e., in a state of equilibrium) teams [71].

[22]In case the power fed to an electricity network is suddenly reduced, the generators deliver to the network an amount of stored energy called inertia. In this way the frequency of the network remains stable for some time, letting the controllers handle the change in the level of available energy. Electricity networks based on renewable energy sources lack such kinds of techniques.

[23]PSO is a bio-inspired algorithm for the optimisation of non-linear functions. It is based on the behavior of flocks of birds, or schools of fish, and it has similarities with other population-based evolutionary algorithms.

Galus and Andersson [14] propose algorithms for an aggregator to trade energy in the energy market by both managing the charging and discharging of the EVs. Based on the current SOC of the vehicles, the desired SOC and the time of departure, it is able to optimise the amount charged in the batteries in order to make a profit by reselling a quantity that leaves the EVs with enough to go onto their onward journeys. Moreover, the authors claim that by using a Model Predictive Control (MPC) approach, their solution has the ability to cope with the error in the forecast of energy output from renewable energy sources, and in particular wind energy. Crucially, they show (using driving patterns derived from MATSim [48], and prediction on wind speeds from the Cosmo-2 model [77]) that their algorithm can handle the infeed error ($<$ 300 MW) from a 500 MW wind park.

Having as a target to maximise the profit of EVs by providing V2G services, Wehinger et al. [72] modeled the German wholesale electricity market and studied the effect of storage devices and a PHEV cluster on the spot prices (prices at particular time points at the day-ahead market). The agents (PHEVs) participate in a day-ahead electricity market, where they submit bidding curves that represent the agent's power output for a specific spot price range. The authors propose a Q-learning based method called "model predictive bidding" (a variation of this algorithm where reinforcement learning in combination with a genetic algorithm is used is presented in [78]) in order to predict future spot prices and maximise their profit. This method initially predicts future spot prices that are later adjusted to incorporate market-power. Finally, the bidding curves are optimised using dynamic programming. The model is evaluated on data taken from a German distribution network and show that an aggregation of PHEVs can lead up to 116% increase in profit compared to a reference scenario with no PHEVs.

Couillet et al. [79] tackle the same problem from a different perspective: they investigate the competitive interaction between EVs or PHEVs in a Cournot market, which consists of electricity transactions to or from a distribution network. In more detail, they formulate the problem as a mean field game and prove that a Nash, or mean field in this case, equilibrium can be reached. In this formulation, EVs and PHEVs trade electricity at prices that change dynamically with the time of the day. In this way, they are incentivized to buy or sell electricity at specific time periods so as the energy demand for other appliances attached to the grid to be met, and overloads to be avoided. Indeed, through some simple simulations, the authors proved that peak time electricity demand is reduced and it is shifted to periods of the day with lower demand.

Kamboj et al. [73] consider a special case of V2G programmes where EVs can be grouped into coalitions to participate in electricity markets and make profit by selling electricity. The formation of coalitions is needed in order the minimum amount of energy Transmission System Operators (TSOs) are willing to buy to be reached. In so doing, four types of agents are used:

1) Vehicle agent: captures the preferences of the EV owners (minimum SOC, profit from V2G services).
2) Aggregator agent: forms the best coalitions, trades on

TABLE V
CLASSIFICATION OF PAPERS - INTEGRATING EVS INTO THE SMART GRID - *Electricity Markets*

cit.	Specific Goal	Problem Solving Technique	Control Scheme	Evaluation Method
[26]	Market-based EV charging control techniques for load balancing and regulation services provision for the integration of renewables	Mathematical programming, and stochastic optimisation, and auctions	Centralised	Simulation using real data from a low voltage residential feeder in Texas, USA
[57]	Bidding strategy for participation in hour-ahead electricity markets to minimise cost	Mathematical programming, and stochastic optimisation, and auctions	Centralised	Theoretical evaluation and simulation using data from various distribution networks in the USA (100 EVs)
[58]	Bidding strategy for participation in a day-ahead market to minimise cost	Mathematical programming, and auctions	Centralised	Simulation based on data from German and Austrian distribution operators prices (1M EVs)
[27]	Bidding strategy for participation in a day-ahead electricity market to minimise cost	Mathematical programming and auctions	Centralised	Simulation using data from the Swiss transmition operator Swissgrid
[59]	Day-ahead charging scheduling mechanism considering the load mismatch risk between the day-ahead and the real-time market	Mathematical programming, and auctions	Centralised	Simulation using data on day-ahead prices from an ISO in New England, USA
[61]	Allocation of electric power units to self-interested agents (EV owners) with private preferences	Utility-based agent coordination, and auctions, and mechanism design	Decentralised	Simulation using real data from a EVs-trial in the UK (CABLED project) (100 EVs)
[62]	Online auction protocol for allocating resources (electric power) to agents (EVs)	Utility-based agent coordination, and auctions, and mechanism design	Decentralised	Same as above (200 EVs)
[63]	Two-sided online market for EV charging where multiple charging points are considered	Utility-based agent coordination, and auctions, and mechanism design	Decentralised	Theoretical evaluation and numerical examples (100 EVs)
[65]	Modelling of the problem of the provision of electricity to collectives of EVs from a smart electricity grid as a Stackelberg game	utility-based agent coordination, and game theoretic analysis, and auctions.	Decentralised	Theoretical evaluation and simulation using synthetic data (1000 EVs)

the wholesale energy market, and calculates fair payoffs to the EVs based on their Shapley [80] values.[24]

3) TSO agent: mediates between power systems and aggregator agents to ensure limits on the lines are respected.
4) Battery charger agent: charges EVs based on their individual characteristics.

Vehicle agents send *join* requests to the aggregator agents that are in close proximity (different coalitions are formed in different geographical areas), and choose to participate in the coalition that offers them the largest profit. A key contribution of this work is the actual deployment of the system in Delaware in collaboration with PJM.[25] The system has been shown to result in \$2400 annual profit for each EV owner that provides V2G services.

Taking inspiration from [73], Ramos et al. [74] present a mechanism to form EV coalitions under distribution network constraints. Thus they propose an algorithm to coalesce EVs (based on *join* requests sent) within the same region as well as an incentive scheme that rewards larger coalitions. Individual agents are rewarded with a price that is commensurate with their contribution (i.e., energy contributed) to the coalition. Thus they are able to show (using data from the Brazilian electricity network) that they are able to provide high quality solutions (95% of the optimal coalition structure (one that maximises the sum of the values of all coalitions)) and that their algorithm can scale to large numbers of agents. It is important to note here that the incentives in this mechanism ignore the fact that different coalitions of EVs may need to find buyers for their energy in an energy market (possibly other coalitions of G2V EVs). In light of this, Saad et al. [81] use non-cooperative game theoretic techniques, specifically double auctions, to ensure that an efficient allocation is reached (i.e., one where buyers and sellers pay their reserve prices). In particular, they prove this outcome is a Nash Equilibrium of the system.

In this section, we have presented a number of V2G management approaches that ultimately aim to smooth the integration of renewables into the smart grid. An open challenge remains, however, in incentivising EV owners to participate in such schemes where their needs for energy may be unpredictable. In the next section a classification scheme of the reviewed papers, as well as a discussion on a number of open research issues are presented.

V. DISCUSSION AND OPEN RESEARCH ISSUES

In this paper, we have analyzed the application of Artificial Intelligence techniques to address the major challenges that arise in the deployment and management of Electric Vehicles. In particular, we have studied AI techniques for energy-efficient EV routing and charging point selection, as well as for the integration of EVs into the smart grid.

In order to summarize our study and to provide a concise yet comprehensive framework for characterizing the reviewed papers, we classified key goals, techniques, control schemes and evaluation benchmarks according to the specific research lines (i.e., EVs routing, Congestion management, Integration of EVs into the Smart Grid) (See Tables I-VI).

[24]Shapley value assigns an expected marginal contribution to each player in a coalitional form game with respect to a uniform distribution over the set of all permutations on the set of players.

[25]http://www.pjm.com/.

TABLE VI

CLASSIFICATION OF PAPERS - INTEGRATING EVS INTO THE SMART GRID - *V2G*

cit.	Specific Goal	Problem Solving Technique	Control Scheme	Evaluation Method
[70]	Balance supply against demand and minimise costs	Machine learning	Centralised	Theoretical evaluation and simulation based on synthetic data
[4]	Charge and discharge of EVs to minimise costs and $CO2$ emissions	PSO optimisation, and mathematical programming	Centralised	Simulation using real data on loads and estimations on emission coefficients and generators production (50K EVs)
[14]	Charge and discharge of EVs managing the forecast on energy generation from renewables	Mathematical programming	Centralised	Simulation using driving patterns from MATSim [48] and prediction on wind speeds from the Cosmo-2 model [77] (9000 EVs)
[72]	Q-learning based method to predict future spot prices in a day ahead market and provide financially efficient V2G services	Mathematical programming	Centralised	Simulation based on data from a German distribution grid (8M EVs)
[79]	Game theoretic analysis of V2G services	Game theory and auctions	Decentralised	Theoretical evaluation and simulation using synthetic data
[73]	EV coalitions to participate in electricity markets and maximise revenues	Game theory and auctions	Centralised	Real world deployment in Delaware in collaboration with PJM (5EVs)
[74]	EV coalitions to participate in electricity markets and maximise revenues	Game theory and heuristic optimisation	Decentralised	Simulation using data from the Brazilian electricity network (40 EVs)
[81]	Non-cooperative game to solve the problem of collectives of PHEVs selling energy to the grid	Game theory and auctions	Decentralised	Theoretical evaluation and simulation using synthetic data (20K EVs)

The main areas that were reviewed in this paper were selected based on the main activities of an EV: Driving (Section II, Table I), selecting a charging point (Section III, Table II), and charging and/or discharging (Section IV, Tables III-VI). All of the papers included in this work can be classified under one of these three categories. Apart from these categories, though, we further classify the papers along a number of dimensions:

- **Specific Goal:** This dimension extracts specific challenges addressed within individual papers. For example most of the work on EV routing is related to the energy-efficient routing problem, while the majority of the work regarding the integration of EVs into the smart grid study G2V management.
- **Problem-solving Technique:** For example most work that focus on energy efficient EV routing use graph-based search algorithms, while papers that manage large collectives of EVs either in a G2V or in V2G mode mostly use mathematical programming.
- **Control Scheme:** This is determined by the needs of the application. For example, as we can see from Table IV, utility-based agent coordination is achieved mostly under a decentralised control scheme, while the management of aggregations of EVs for smooth renewable energy integration takes place under a centralised scheme.
- **Evaluation Method:** Most of the works presented in this paper use either theoretical evaluation (i.e., in terms of time complexity, or proof of Nash equilibrium), or simulations using real and rarely synthetic data or, in some rare cases (as in [73]), real-world deployments.

Based on our analysis, we identify a number of key considerations for individual aspects of EV management:

EVs Routing: work here has focused on energy efficient EV routing, and typically represents the road network as directed or undirected graphs. Most of the solutions proposed, consider the energy recuperation ability of the EVs, and the negative edge costs that are derived from it, and adapt existing search algorithms such as the Dijkstra's or the A* in order to find the path that minimises the energy consumption of the EV. Moreover, some of these approaches also consider recharging, but simply as passing through some specific nodes. The algorithm presented in [29] is proven to have the best performance in terms of time and computational complexity. (Note that in this section in addition to the aforementioned dimensions, the complexity of the algorithms is also taken into consideration as it is an aspect thoroughly analyzed in the papers –Table I).

Charging Point Selection: here a number of problem formulations and solution techniques are proposed. They typically focus on the minimisation of the traffic congestion and the delays incurred by the EVs. The most promising solutions are able to select the best charging point and update this decision en-route. Moreover, incentives should be given to EVs in order to avoid charging at few central and congested charging points. Therefore, a combination of the solutions presented in [19] and [38] could prove to handle the problem of the efficient selection of charging points in the most efficient manner. Finally, the work presented in [41] achieves the optimal placement of charging points, while long detours of EVs are avoided.

Integrating EVs to the Smart Grid: here work is typically divided in terms of G2V and V2G algorithms (to the exception of [14] which considers both in one system).

We have identified important benchmarks in the context of G2V, namely (i) [3] which uses mathematical programming techniques to solve the problem of charging EVs from intermittent renewable sources (ii) [38] and [54] which use congestion pricing techniques to balance the load across the network and (iii) [63] and [57] that use market based techniques to ensure EVs satisfy grid constraints and participate in energy markets.

In terms of V2G solutions, most of them focus on profit maximisation as a key goal. Solution techniques range from mathematical programming techniques to optimise trading decisions on the energy market [14] and coalition formation to devise discharging policies [73].

Against this background we can identify some key scientific dimensions of the problems that need to be tackled:

1) *Uncertainty*: while several algorithms have been proposed to account for uncertainty in renewable energy production (e.g., [82], and [49]), very few tackle the uncertainty in arrival and departure times of EVs and the load they will impose on the distribution network (some initial work can be found in [83] and in [84]), as well as uncertainties in the reliability of communication systems used to coordinate collectives of EVs. The challenge here is to produce predictions at short notice as late decisions could potentially result in major disruption to the transportation network. Hence, efficient machine learning algorithms need to be developed to predict behaviours in the system. In particular, we believe predictions could be improved by constructing better models of human mobility [85], [86] as well as by fusing data from across the transportation network [87], [88]. To this end, large scale deployments of EVs are essential, should future machine learning algorithms be trained and evaluated in big enough datasets so as their efficiency to be maximised.
2) *Dynamism*: The state of the electricity grid, the production of renewable sources, the charging point availability, the congestion at communication and transportation networks and the number of EVs available to provide V2G services, change quickly while a large number of EVs are either driving or charging. Under such a dynamic setting, fail-safe mechanisms and approximation algorithms will be required to solve optimisation problems at short notice, while minimising communication bandwidth. While we have noted a number of solutions that use stochastic dynamic programming, it will be interesting to see if such solutions can be decentralised to ensure the system is more robust (e.g., to [89]).

Furthermore, we can identify some key engineering dimensions of the problems that need to be tackled:

1) *Interoperability*: There is a need EV technologies to be able to work seamlessly and efficiently together. Different types of chargers should be able to work with all EV models, and data exchanged between entities (EVs, charging points, network operators) should have an understandable by all format and meaning. For example, Semantic Web technologies, such as XML, RDF and ontologies, can provide a structured and consistent way to represent the data being exchanged, and therefore, make the collaboration of various technologies more efficient.
2) *Privacy*: In order for EVs to be efficiently managed in terms of driving, charging and/or discharging, data on the location and the preferences of them must, in many cases, be obtained by a central mechanism. This creates issues of privacy and data protection, as drivers might not be willing to disclose such information.
3) *Real world validation*: Currently, most of the mechanisms and the technologies related to the management of the EVs remain at a theoretical or at a pilot deployment level, and thus, their effectiveness in a large scale deployment has not been validated. The design of effective interfaces for human-EV (agent) interaction to be smooth and efficient, but also the research on ways to motivate and incentivize people to follow the instructions given to them by systems (e.g., a routing system giving instruction on an energy efficient route to take, or a charging point to charge at) is crucial. Moreover, the complexity of the coordination of a large number of entities (e.g., EVs, charging points, electricity network managers), and the ability of the systems to react to unexpected situations (e.g., a large number of EVs wanting to charge within a short period of time) and prevent negative events (e.g., overloading of the electricity network) must be carefully studied, analyzed, and verified.

Our study has shown that several AI-based approaches are emerging in all areas of EV management: from battery charging algorithms to network congestion management algorithms. We believe in a concerted effort, involving transportation engineers, power systems experts, and AI researchers, in order to bring the benefits of such solutions to the real world. Hence, we advocate more joint deployments of novel AI solutions in field trials, with users of different types, in order to unpack more specific challenges that remain to be addressed before EVs can be deployed at scale. Moreover, AI techniques being exploratory in nature can help EV researchers to quickly explore optimisation search spaces using heuristics. Since modern AI is being based on solid scientific approaches all its experiments and results are verifiable and reproducible. Thus, engineers can base the development of standards, which are crucial should a systematic management of EVs activities be achieved, on the results of AI research on EVs [90]. Currently, a number of EV related standards already exist,[26] and others are under way.

Acknowledgment

The authors would like to thank Dr. George Koutitas (International Hellenic University) for the fruitful comments and discussions. We would like to also acknowledge support from the EPSRC-Funded International Centre for Infrastructure Futures (EP/K012347/1).

References

[1] U. E. I. Administration, "Annual energy review 2011," Tech. Rep. DOE/EIA-0384(2011), 2012.
[2] IEA, "Global ev outlook," Tech. Rep., 2013.
[3] M. Gonzalez Vaya and G. Andersson, "Integrating renewable energy forecast uncertainty in smart-charging approaches for plug-in electric vehicles," in *PowerTech (POWERTECH), 2013 IEEE Grenoble*, 2013, pp. 1–6.
[4] A. Saber and G. Venayagamoorthy, "Plug-in vehicles and renewable energy sources for cost and emission reductions," *Industrial Electronics, IEEE Transactions on*, vol. 58, no. 4, pp. 1229–1238, April 2011.
[5] H. Farhangi, "The path of the smart grid," *Power and Energy Magazine, IEEE*, vol. 8, no. 1, pp. 18–28, January 2010.

[26] http://electricvehicle.ieee.org/standards.

[6] W. Kempton and J. Tomic, “Vehicle-to-grid power fundamentals: Calculating capacity and net revenue,” *Journal of Power Sources*, vol. 144, no. 1, pp. 268 – 279, 2005.

[7] M. D. Galus, M. G. Vay, T. Krause, and G. Andersson, “The role of electric vehicles in smart grids,” *Wiley Interdisciplinary Reviews: Energy and Environment*, vol. 2, no. 4, pp. 384–400, 2013. [Online]. Available: http://dx.doi.org/10.1002/wene.56

[8] W. J. Mitchel, C. E. Borroni-Bird, and L. D. Burns, *Reinventing the automobile: Personal urban mobility for the 21st century*. MIT Press, 2010.

[9] J. Tomic and W. Kempton, “Using fleets of electric-drive vehicles for grid support,” *Journal of Power Sources*, vol. 168, no. 2, pp. 459 – 468, 2007.

[10] W. Kempton and J. Tomic, “Vehicle-to-grid power implementation: From stabilizing the grid to supporting large-scale renewable energy,” *Journal of Power Sources*, vol. 144, no. 1, pp. 280 – 294, 2005.

[11] M. A. Nicholas, G. Tal, and J. Woodjack, “California statewide charging survey: What do drivers want?” *92nd Annual Meeting of the Transportation Research Board*, 2013.

[12] J. Peas Lopes, F. Soares, and P. Almeida, “Identifying management procedures to deal with connection of electric vehicles in the grid,” in *PowerTech, 2009 IEEE Bucharest*, 2009, pp. 1–8.

[13] L. Gan, U. Topcu, and S. Low, “Optimal decentralized protocol for electric vehicle charging,” *Power Systems, IEEE Transactions on*, vol. 28, no. 2, pp. 940–951, 2013.

[14] M. Galus and G. Andersson, “Balancing renewable energy source with vehicle to grid services from a large fleet of plug-in hybrid electric vehicles controlled in a metropolitan area distribution network,” *Cigrè 2011, Bologna, 13-15 Sep*, 2011.

[15] A. Artmeier, J. Haselmayr, M. Leucker, and M. Sachenbacher, “The shortest path problem revisited: Optimal routing for electric vehicles,” in *KI 2010: Advances in Artificial Intelligence*. Springer Berlin Heidelberg, 2010, vol. 6359, pp. 309–316.

[16] S. Storandt, “Quick and energy-efficient routes: computing constrained shortest paths for electric vehicles,” in *Proceedings of the 5th ACM SIGSPATIAL International Workshop on Computational Transportation Science*, ser. IWCTS ’12, 2012, pp. 20–25.

[17] S. Storandt and S. Funke, “Cruising with a battery-powered vehicle and not getting stranded,” in *26th Conf. on Artificial Intelligence (AAAI)*, 2012.

[18] S. Ermon, Y. Xue, C. Gomes, and B. Selman, “Learning policies for battery usage optimization in electric vehicles,” in *Machine Learning and Knowledge Discovery in Databases*, ser. Lecture Notes in Computer Science, P. Flach, T. Bie, and N. Cristianini, Eds. Springer Berlin Heidelberg, 2012, vol. 7524, pp. 195–210.

[19] M. M. De Weerdt, E. H. Gerding, S. Stein, V. Robu, and N. R. Jennings, “Intention-aware routing to minimise delays at electric vehicle charging stations,” in *Proceedings of the Twenty-Third International Joint Conference on Artificial Intelligence*, ser. IJCAI’13. AAAI Press, 2013, pp. 83–89.

[20] H. Qin and W. Zhang, “Charging scheduling with minimal waiting in a network of electric vehicles and charging stations,” in *Proceedings of the Eighth ACM International Workshop on Vehicular Inter-networking*, ser. VANET ’11. New York, NY, USA: ACM, 2011, pp. 51–60.

[21] S. Storandt and S. Funke, “Enabling e-mobility: Facility location for battery loading stations,” in *27th Conf. on Artificial Intelligence (AAAI)*, 2013.

[22] S. Vandael, N. Boucké, T. Holvoet, K. De Craemer, and G. Deconinck, “Decentralized coordination of plug-in hybrid vehicles for imbalance reduction in a smart grid,” in *The 10th International Conference on Autonomous Agents and Multiagent Systems - Volume 2*, 2011, pp. 803–810.

[23] Q. Li, T. Cui, R. Negi, F. Franchetti, and M. D. Ilic, “On-line decentralized charging of plug-in electric vehicles in power systems,” *arXiv preprint arXiv:1106.5063*, 2011.

[24] O. Sundstrom and C. Binding, “Planning electric-drive vehicle charging under constrained grid conditions,” in *Power System Technology (POWERCON), 2010 International Conference on*, 2010, pp. 1–6.

[25] Z. Ma, D. Callaway, and I. Hiskens, “Decentralized charging control for large populations of plug-in electric vehicles,” in *Decision and Control (CDC), 2010 49th IEEE Conference on*, 2010, pp. 206–212.

[26] M. Caramanis and J. Foster, “Management of electric vehicle charging to mitigate renewable generation intermittency and distribution network congestion,” in *Decision and Control, 2009 held jointly with the 2009 28th Chinese Control Conference. CDC/CCC 2009. Proceedings of the 48th IEEE Conference on*, 2009, pp. 4717–4722.

[27] M. Gonzalez Vaya and G. Andersson, “Optimal bidding strategy of a plug-in electric vehicle aggregator in day-ahead electricity markets,” in *European Energy Market (EEM), 2013 10th International Conference on the*, 2013, pp. 1–6.

[28] J. P. Lopes, F. J. Soares, P. Almeida, and M. M. da Silva, “Smart charging strategies for electric vehicles: Enhancing grid performance and maximizing the use of variable renewable energy resources,” in *EVS24 Intenational Battery, Hybrid and Fuell Cell Electric Vehicle Symposium, Stavanger, Norveška*, 2009.

[29] J. Eisner, S. Funke, and S. Storandt, “Optimal route planning for electric vehicles in large networks,” in *25th Conf. on Artificial Intelligence (AAAI)*, 2011.

[30] B. C. Dean, “Shortest paths in fifo time-dependent networks: Theory and algorithms,” Tech. Rep., 2004.

[31] D. B. Johnson, “Efficient algorithms for shortest paths in sparse networks,” *J. ACM*, vol. 24, no. 1, pp. 1–13, Jan. 1977.

[32] R. Geisberger, P. Sanders, D. Schultes, and D. Delling, “Contraction hierarchies: Faster and simpler hierarchical routing in road networks,” in *Experimental Algorithms*, ser. Lecture Notes in Computer Science, C. McGeoch, Ed. Springer Berlin Heidelberg, 2008, vol. 5038, pp. 319–333.

[33] M. Sachenbacher, M. Leucker, A. Artmeier, and J. Haselmayr, “Efficient energy-optimal routing for electric vehicles,” in *25th Conf. on Artificial Intelligence (AAAI)*, 2011.

[34] T. Sweda and D. Klabjan, “Finding minimum-cost paths for electric vehicles,” in *Electric Vehicle Conference (IEVC), 2012 IEEE International*, 2012, pp. 1–4.

[35] I. Bayram, G. Michailidis, M. Devetsikiotis, F. Granelli, and S. Bhattacharya, “Smart vehicles in the smart grid: Challenges, trends, and application to the design of charging stations,” in *Control and Optimization Methods for Electric Smart Grids*, ser. Power Electronics and Power Systems, A. Chakrabortty and M. D. Ili, Eds. Springer New York, 2012, vol. 3, pp. 133–145.

[36] M. Winter and R. J. Brodd, “What are batteries, fuel cells, and supercapacitors?” *Chemical reviews*, vol. 104, no. 10, pp. 4245–4270, 2004.

[37] P. Dille, M. Duescher, I. Nourbakhsh, G. Podnar, and J. Schapiro, “Evaluating the urban electric vehicle,” Tech. Rep., 2010.

[38] E. S. Rigas, S. D. Ramchurn, N. Bassiliades, and G. Koutitas, “Congestion management for urban ev charging systems,” in *Smart Grid Communications (SmartGridComm), 2013 IEEE International Conference on*, 2013, pp. 121–126.

[39] S. Bessler and J. Grønbæk, “Routing ev users towards an optimal charging plan,” in *International Battery, Hybrid and Fuel Cell Electric Vehicle Symposium*, 2012.

[40] A. Lam, Y.-W. Leung, and X. Chu, “Electric vehicle charging station placement,” in *Smart Grid Communications (SmartGridComm), 2013 IEEE International Conference on*, Oct 2013, pp. 510–515.

[41] S. Funke, A. Nusser, and S. Storandt, “Placement of loading stations for electric vehicles: No detours necessary!” in *Twenty-Eighth AAAI Conference on Artificial Intelligence*, 2014.

[42] R. Greiner, B. A. Smith, and R. W. Wilkerson, “A correction to the algorithm in reiter’s theory of diagnosis,” *Artificial Intelligence*, vol. 41, no. 1, pp. 79 – 88, 1989.

[43] I. Bayram, G. Michailidis, and M. Devetsikiotis, “Electric power resource provisioning for large scale public ev charging facilities,” in *Smart Grid Communications (SmartGridComm), 2013 IEEE International Conference on*, Oct 2013, pp. 133–138.

[44] K. Clement-Nyns, E. Haesen, and J. Driesen, “The impact of charging plug-in hybrid electric vehicles on a residential distribution grid,” *Power Systems, IEEE Transactions on*, vol. 25, no. 1, pp. 371–380, 2010.

[45] C. Ahn, C.-T. Li, and H. Peng, “Decentralized charging algorithm for electrified vehicles connected to smart grid,” in *American Control Conference (ACC), 2011*, 2011, pp. 3924–3929.

[46] R. Halvgaard, N. Poulsen, H. Madsen, J. Jorgensen, F. Marra, and D. Bondy, “Electric vehicle charge planning using economic model predictive control,” in *Electric Vehicle Conference (IEVC), 2012 IEEE International*, 2012, pp. 1–6.

[47] I. Bayram, G. Michailidis, M. Devetsikiotis, and F. Granelli, “Electric power allocation in a network of fast charging stations,” *Selected Areas in Communications, IEEE Journal on*, vol. 31, no. 7, pp. 1235–1246, July 2013.

[48] M. Balmer, K. W. Axhausen, and K. Nagel, “Agent-based demand-modeling framework for large-scale microsimulations,” *Transportation Research Record: Journal of the Transportation Research Board*, vol. 1985, no. 1, pp. 125–134, 2006.

[49] P. Pinson, H. Madsen, H. A. Nielsen, G. Papaefthymiou, and B. Klckl, "From probabilistic forecasts to statistical scenarios of short-term wind power production," *Wind Energy*, vol. 12, no. 1, pp. 51–62, 2009.

[50] M. Gonzalez Vaya and G. Andersson, "Centralized and decentralized approaches to smart charging of plug-in vehicles," in *Power and Energy Society General Meeting, 2012 IEEE*, 2012, pp. 1–8.

[51] H.-J. Kim, J. Lee, and G.-L. Park, "Constraint-based charging scheduler design for electric vehicles," in *Intelligent Information and Database Systems*, ser. Lecture Notes in Computer Science, J.-S. Pan, S.-M. Chen, and N. Nguyen, Eds. Springer Berlin Heidelberg, 2012, vol. 7198, pp. 266–275.

[52] E. Karfopoulos and N. Hatziargyriou, "A multi-agent system for controlled charging of a large population of electric vehicles," *Power Systems, IEEE Transactions on*, vol. 28, no. 2, pp. 1196–1204, May 2013.

[53] T. P. Lyon, M. Michelin, A. Jongejan, and T. Leahy, "Is smart charging policy for electric vehicles worthwhile?" *Energy Policy*, vol. 41, no. 0, pp. 259 – 268, 2012.

[54] I. Bayram, G. Michailidis, I. Papapanagiotou, and M. Devetsikiotis, "Decentralized control of electric vehicles in a network of fast charging stations," in *Global Communications Conference (GLOBECOM), 2013 IEEE*, Dec 2013, pp. 2785–2790.

[55] J. Escudero-Garzas and G. Seco-Granados, "Charging station selection optimization for plug-in electric vehicles: An oligopolistic game-theoretic framework," in *Innovative Smart Grid Technologies (ISGT), 2012 IEEE PES*, Jan 2012, pp. 1–8.

[56] O. Ardakanian, C. Rosenberg, and S. Keshav, "Distributed control of electric vehicle charging," in *Proceedings of the Fourth International Conference on Future Energy Systems*, ser. e-Energy '13. New York, NY, USA: ACM, 2013, pp. 101–112.

[57] J. Foster and M. Caramanis, "Optimal power market participation of plug-in electric vehicles pooled by distribution feeder," *Power Systems, IEEE Transactions on*, vol. 28, no. 3, pp. 2065–2076, Aug 2013.

[58] M. G. Vayá and G. Andersson, "Locational marginal pricing based smart charging of plug-in hybrid vehicle fleets," in *Smart Energy Strategies Conference*, 2011, pp. 21–23.

[59] L. Yang, J. Zhang, and D. Qian, "Risk-aware day-ahead scheduling and real-time dispatch for plug-in electric vehicles," in *Global Communications Conference (GLOBECOM), 2012 IEEE*, 2012, pp. 3026–3031.

[60] R. Van Slyke and R. Wets, "L-shaped linear programs with applications to optimal control and stochastic programming," *SIAM Journal on Applied Mathematics*, vol. 17, no. 4, pp. 638–663, 1969. [Online]. Available: http://epubs.siam.org/doi/abs/10.1137/0117061

[61] S. Stein, E. Gerding, V. Robu, and N. R. Jennings, "A model-based online mechanism with pre-commitment and its application to electric vehicle charging," in *Proceedings of the 11th International Conference on Autonomous Agents and Multiagent Systems - Volume 2*, ser. AAMAS '12. Richland, SC: International Foundation for Autonomous Agents and Multiagent Systems, 2012, pp. 669–676.

[62] E. H. Gerding, V. Robu, S. Stein, D. C. Parkes, A. Rogers, and N. R. Jennings, "Online mechanism design for electric vehicle charging," in *The 10th International Conference on Autonomous Agents and Multiagent Systems - Volume 2*. Richland, SC: International Foundation for Autonomous Agents and Multiagent Systems, 2011, pp. 811–818.

[63] E. H. Gerding, S. Stein, V. Robu, D. Zhao, and N. R. Jennings, "Two-sided online markets for electric vehicle charging," in *Proceedings of the 2013 International Conference on Autonomous Agents and Multi-agent Systems*. Richland, SC: International Foundation for Autonomous Agents and Multiagent Systems, 2013, pp. 989–996.

[64] R. Bent and P. Van Hentenryck, "The value of consensus in online stochastic scheduling." in *ICAPS*, vol. 4, 2004, pp. 219–226.

[65] W. Tushar, W. Saad, H. Poor, and D. Smith, "Economics of electric vehicle charging: A game theoretic approach," *Smart Grid, IEEE Transactions on*, vol. 3, no. 4, pp. 1767–1778, Dec 2012.

[66] H. Von StackelberG, D. Bazin, R. Hill, and L. Urch, *Market structure and equilibrium*. Springer, 2010.

[67] W. Su and M.-Y. Chow, "Performance evaluation of a phev parking station using particle swarm optimization," in *Power and Energy Society General Meeting, 2011 IEEE*, July 2011, pp. 1–6.

[68] C. Chan and Y. S. Wong, "Electric vehicles charge forward," *Power and Energy Magazine, IEEE*, vol. 2, no. 6, pp. 24–33, Nov 2004.

[69] W. Kempton, V. Udo, K. Huber, K. Komara, S. Letendre, S. Baker, D. Brunner, and N. Pearre, "A test of vehicle-to-grid (v2g) for energy storage and frequency regulation in the pjm system," *Results from an Industry-University Research Partnership*, p. 32, 2008.

[70] S. Chatzivasileiadis, M. Galus, Y. Reckinger, and G. Andersson, "Q-learning for optimal deployment strategies of frequency controllers using the aggregated storage of phev fleets," in *PowerTech, 2011 IEEE Trondheim*, 2011, pp. 1–8.

[71] G. Chalkiadakis, E. Markakis, and C. Boutilier, "Coalition formation under uncertainty: Bargaining equilibria and the bayesian core stability concept," in *Proceedings of the 6th International Joint Conference on Autonomous Agents and Multiagent Systems*, ser. AAMAS '07. New York, NY, USA: ACM, 2007, pp. 64:1–64:8. [Online]. Available: http://doi.acm.org/10.1145/1329125.1329203

[72] L. A. Wehinger, G. Hug, M. D. Galus, and G. Andersson, "Assessing the effect of storage devices and a phev cluster on german spot prices by using model predictive and profit maximizing agents," in *Proceedings of the Power Systems Computation Conference (PSCC)*, 2011.

[73] S. Kamboj, W. Kempton, and K. S. Decker, "Deploying power grid-integrated electric vehicles as a multi-agent system," in *The 10th International Conference on Autonomous Agents and Multiagent Systems - Volume 1*, ser. AAMAS '11. Richland, SC: International Foundation for Autonomous Agents and Multiagent Systems, 2011, pp. 13–20.

[74] G. de O Ramos, J. Rial, and A. Bazzan, "Self-adapting coalition formation among electric vehicles in smart grids," in *Self-Adaptive and Self-Organizing Systems (SASO), 2013 IEEE 7th International Conference on*, Sept 2013, pp. 11–20.

[75] J. Kennedy, R. Eberhart *et al.*, "Particle swarm optimization," in *Proceedings of IEEE international conference on neural networks*, vol. 4, no. 2. Perth, Australia, 1995, pp. 1942–1948.

[76] T. Ting, M. V. C. Rao, and C. Loo, "A novel approach for unit commitment problem via an effective hybrid particle swarm optimization," *Power Systems, IEEE Transactions on*, vol. 21, no. 1, pp. 411–418, Feb 2006.

[77] J. Steppeler, G. Doms, U. Schttler, H. W. Bitzer, A. Gassmann, U. Damrath, and G. Gregoric, "Meso-gamma scale forecasts using the nonhydrostatic model lm," *Meteorology and Atmospheric Physics*, vol. 82, no. 1-4, pp. 75–96, 2003.

[78] L. Wehinger, M. D. Galus, and G. Andersson, "Agent-based simulator for the german electricity wholesale market including wind power generation and widescale phev adoption," in *Energy Market (EEM), 2010 7th International Conference on the European*, 2010, pp. 1–6.

[79] R. Couillet, S. Perlaza, H. Tembine, and M. Debbah, "Electrical vehicles in the smart grid: A mean field game analysis," *Selected Areas in Communications, IEEE Journal on*, vol. 30, no. 6, pp. 1086–1096, July 2012.

[80] L. S. Shapley, "A value for n-person games," *Contributions to the Theory of Games*, vol. 2, 1953.

[81] W. Saad, Z. Han, H. Poor, and T. Basar, "A noncooperative game for double auction-based energy trading between phevs and distribution grids," in *Smart Grid Communications (SmartGridComm), 2011 IEEE International Conference on*, Oct 2011, pp. 267–272.

[82] A. Panagopoulos, G. Chalkiadakis, and E. Koutroulis, "Predicting the power output of distributed renewable energy resources within a broad geographical region," in *ECAI*, 2012, pp. 981–986.

[83] N. Ghiasnezhad Omran and S. Filizadeh, "Location-based forecasting of vehicular charging load on the distribution system," *Smart Grid, IEEE Transactions on*, vol. 5, no. 2, pp. 632–641, March 2014.

[84] S. Shahidinejad, S. Filizadeh, and E. Bibeau, "Profile of charging load on the grid due to plug-in vehicles," *Smart Grid, IEEE Transactions on*, vol. 3, no. 1, pp. 135–141, March 2012.

[85] R. Becker, R. Cáceres, K. Hanson, S. Isaacman, J. M. Loh, M. Martonosi, J. Rowland, S. Urbanek, A. Varshavsky, and C. Volinsky, "Human mobility characterization from cellular network data," *Commun. ACM*, vol. 56, no. 1, pp. 74–82, Jan. 2013. [Online]. Available: http://doi.acm.org/10.1145/2398356.2398375

[86] K. Lee, S. Hong, S. J. Kim, I. Rhee, and S. Chong, "Slaw: A new mobility model for human walks," in *INFOCOM 2009, IEEE*, April 2009, pp. 855–863.

[87] G.-Z. Yang, J. Andreu-Perez, X. Hu, and S. Thiemjarus, "Multi-sensor fusion," in *Body Sensor Networks*, G.-Z. Yang, Ed. Springer London, 2014, pp. 301–354.

[88] D. Hall and J. Llinas, "An introduction to multisensor data fusion," *Proceedings of the IEEE*, vol. 85, no. 1, pp. 6–23, Jan 1997.

[89] R. Stranders, A. Farinelli, A. Rogers, and N. R. Jennings, "Decentralised coordination of continuously valued control parameters using the max-sum algorithm," in *Proceedings of The 8th International Conference on Autonomous Agents and Multiagent Systems - Volume 1*, ser. AAMAS '09, 2009, pp. 601–608.

[90] S. Russell, "Norvig (2003)," *Artificial intelligence: a modern approach*, pp. 25–26, 2003.

Emmanouil S. Rigas (S' 12) received the BSc degree in Computer Science from the Department of Informatics of the Aristotle University of Thessaloniki, Greece in 2011, and the MSc degree in Web Technology from the University of Southampton, UK in 2012. Currently he is working towards his PhD degree at the Department of Informatics of the Aristotle University. His research interests include the use of Artificial Intelligence and in particular multi-agent systems, for the management of Electric Vehicles in the Smart Grid.

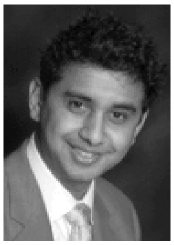

Sarvapali D. Ramchurn received the BEng (Hons.)in Information Systems Engineering from Imperial College London in 2001 and the PhD in Multiagent Systems from the University of Southampton in 2004. Currently, he is a Roberts Fellow in the Agents, Interaction, and Complexity research group with the University of Southampton, Southampton, U.K., where he carries out research on the design of autonomous agents and multi-agent systems for realworld sociotechnical applications, including energy systems, disaster management, and crowdsourcing. His research touches on a number of fields including machine learning, data science, and game theory. Specifically, he has pioneered the development of agent-based coordination algorithms for distributed task allocation that have been deployed on real-world unmanned aerial vehicles and in the Premier Leagues Fantasy Football game, where his approach has been shown to outperform more than two million human players. His papers have been cited more than 3000 times (according to Google scholar), and his work has been featured in various media including BBC News, New Scientist, EPSRC Pioneer, and Wired.

Nick Bassiliades (http://tinyurl.com/nbassili) (M' 93) received the M.Sc. degree in applied artificial intelligence from the Department of Computing Science, University of Aberdeen, Aberdeen, U.K., in 1992 and the Ph.D. degree in parallel knowledge base systems from the Department of Informatics, Aristotle University of Thessaloniki, Thessaloniki, Greece, in 1998.

He is currently an Associate Professor with the Aristotle University of Thessaloniki, teaching logic programming, knowledge systems, knowledge management, and semantic Web. He is also an Academic Coordinator with the School of Science and Technology, International Hellenic University, Thessaloniki. He has published more than 160 papers in journals, conferences, and books and has coauthored two books. His published research has received over 1000 citations (h-index 18). His research interests include knowledge-based and rule systems, multiagent systems, ontologies, linked open data, and semantic Web.

Dr. Bassiliades was on the Program Committee of more than 70 conferences/ workshops and on the Organizational Committee of six conferences/workshops. He has been the Program Cochair of RuleML-2008, RuleML-2011@IJCAI, and WIMS-2014. He has been involved in 30 RD projects leading nine of them. He has been the General Secretary of the Board of the Greek Artificial Intelligence Society, the Director of RuleML, Inc., and a member of the Greek Computer Society and the Association for Computing Machinery