

Multi-layer event-based Vehicle-to-Grid (V2G) scheduling with short term predictive capability within a modular aggregator control structure

Hannes Krueger, and Andrew Cruden

Abstract—In this work a novel method of event-based V2G scheduling is devised that is suitable for dynamic real time aggregator control in large scale V2G applications within centrally controlled EV car parks. The method is applicable in deterministic systems where a V2G network provides or receives electricity in reoccurring and predictable patterns (events). The scheduling strategy shown is based on a robust modular high-level aggregator control structure and a proposed communications and data management system. The scheduling consists of three algorithm layers, differentiating between predictive scheduling for in-event periods, smart charging for out-of-event periods and reactive scheduling for ongoing adjustments in real-time to account for uncertainty. The scheduling process is described in terms of its underlying rules for prioritising EVs to be either charged or discharged. It's behaviour is then analysed using a simulated car park of up to one thousand connected EVs for an example application in which a V2G network is used to support nearby electrified rail infrastructure, providing power for train acceleration and accepting power from regenerative braking. The departure or arrival of a train of known type and speed pattern can be regarded as a reoccurring event and its effect on the V2G network is therefore predictable due to train schedules and tracking.

Index Terms—Vehicle-to-Grid, aggregator control, predictive scheduling

I. INTRODUCTION

The desire to reduce our dependency on fossil fuels and resulting environmental impacts in individual transportation has led to significant increases in the number of electric vehicles (EVs) in recent years. According to the International Energy Agency, the global EV fleet grew by 2 million, exceeding 5.1 million in 2018 alone [1]. The global automotive industry is investing heavily in EV technology to increase market penetration. This includes industry heavy-weights such as Ford [2], Toyota [2], Volkswagen [2][3], Nissan [2] and Porsche [4].

While generally regarded as a positive development, this electrification creates major challenges for power grids world wide. Shifting individual transportation away from fossil fuel combustion increases the electricity demand for EV battery charging and creates power demand peaks if charging is not

managed effectively. In order to mitigate these challenges much work has been done in the field of EV smart charging [5][6][7].

The concept of V2G goes further in an attempt to turn EVs into a valuable resource for power grids rather than just a burdensome load. It describes the use of EV battery packs as aggregated distributed grid-based energy storage. This requires bi-directional power flow between parked EVs and the power grid [8][9]. The power flow is generally managed by a so-called aggregator, which arbitrates between the needs of EVs and the Grid (ensuring a net flow of energy to the EV battery pack to recharge it for driving purposes). A V2G network needs to serve both the grid and the connected EVs sufficiently to a) fulfil the designated grid service provided, and b) to battle potential range anxiety from EV owners and encourage participation in V2G schemes [10].

To manage the power flow between an EV and the power grid, the aggregator requires information about the current state of the EV's battery pack. At a minimum, the information required is the State-of-Charge (SOC) as well as the overall capacity and possible rate of charge/discharge. For this data to be transmitted, the EV must be equipped with suitable hardware and software that allows for access from the aggregator. As outlined in previous work [11], the author's assume that each EV and EV charger on the network is equipped with communication hardware able to transmit and receive information via Ethernet or Wi-Fi connection using the REST API [12][13]. However, no generally accepted standard for such hardware and software is currently adopted in the electric vehicle industry.

V2G has been proposed for various applications; each different in scale and requirements towards overall system responsiveness (i.e. how quickly a V2G network adapts to changes in power demand or EV population). In [14] V2G is being proposed for load shifting with scheduling taking place for a whole day ahead using hourly time intervals. Similarly, [15] proposes a decentralised V2G dispatch strategy, where schedules are determined at the beginning of every 30 minute time interval. [16] proposes the usage of V2G for aggregator profit maximisation during peak load shaving while lowering the cost of EV charging to the customer. The model used includes factors such as battery degradation and battery replacement costs and uses 30 minute time intervals for scheduling. The maximisation of aggregator profits is also the aim of [17], which is using a 5 minute resolution. [18] aims at minimising charging costs of EVs using time

The authors are with the EPSRC Centre for Doctoral Training in Energy Storage and its Applications, University of Southampton, United Kingdom (email: hk5g11@soton.ac.uk). This work was supported by the Engineering and Physical Sciences Research Council (EPSRC) Centre for Doctoral Training in "Energy Storage and Its Applications" (EP/L016818/1) and the "Transenergy – Road to Rail Energy Exchange (R2REE)" (EP/N022289/1) projects.

Manuscript received XXX, XX, 2015; revised XXX, XX, 2015.

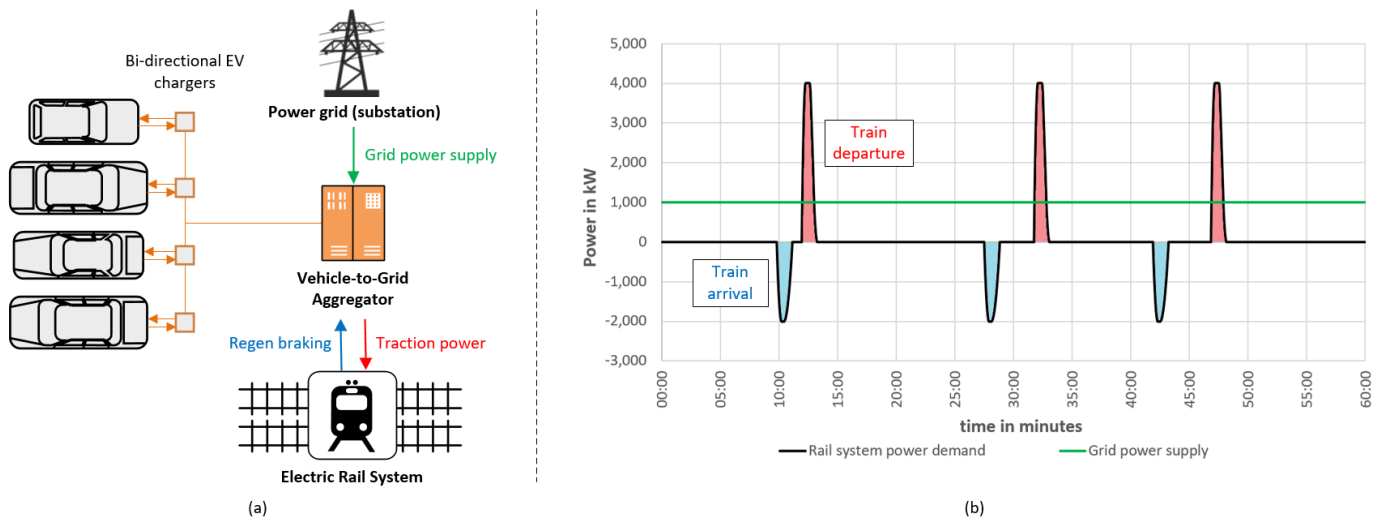


Fig. 1. (a) V2G for support of local electric rail systems, system overview and power flows: EV population acts as buffer between the grid connection and the fluctuating rail system power demands to ensure steady power flow from the grid; (b) Rail system power demand over one hour (3 trains arriving/departing sequentially, varying dwell time): traction power drawn for train acceleration - positive/red, power supplied from regenerative braking - negative/blue.

intervals of 1 hour for a large scale V2G network of up to 400 distributed EVs. Also aimed at minimising the charging costs of individual EVs is the work presented in [19] where a resolution of 10 minutes was identified as the best compromise between precision and computational cost. In [20] EV users are offered a variety of contracts when connecting to a V2G network stating how much energy will be charged, how much will be made available for V2G as well as the charging costs and compensation for V2G participation. Such an approach gives a lot of control to the EV user but adds significant complexity and constraints to the scheduling process.

In the examples above system responsiveness is not a major consideration as changes in power demand are gradual and a delay of a few seconds in making scheduling decisions would not significantly effect the system performance. In [21] a V2G network is used to provide frequency regulation in a power grid with integrated renewable electricity generation while minimising costs. An EV population of 1,000 EVs was simulated, but execution times for decision making processes have not been reported. Similarly in [22] a V2G network also provides frequency regulation but the focus lies on fairness criteria in the treatment of different EVs (fair distribution of power within the EV population). Again, execution times have not been reported. The response times required for frequency regulation can vary widely between "milliseconds" up to 20 minutes [23], whereas primary frequency response requires energy storage systems to deliver rated power within 10 seconds [24].

The V2G concept has the potential to provide system response times within a few seconds or even on a sub-second time-scale, but certain challenges of vehicle management, communication and decision making have to be addressed. These challenges include:

- the scale of EV population on the network
- the changing energy storage potential (connected capacity, power available, SOC of EV battery packs)
- predictability of EV availability and power grid demands

- the necessity to charge EVs over time (constraining scheduling process)
- aggregator-to-EV communication delays (exacerbated by the need for encryption to address security and privacy concerns [25])
- complexity of underlying scheduling rules

The authors believe that the challenges of efficient scheduling in large scale, centrally controlled V2G networks (whole network working towards a common goal) with regards to system responsiveness are not yet sufficiently addressed in literature. Therefore a novel approach to V2G scheduling is being proposed that is aimed at increasing system responsiveness by utilising both, predictive scheduling (charge and discharge decisions taken ahead of predictable changes in power demand) and reactive scheduling (decisions taken in reaction to changes in power demand).

Uncertainty is a major challenge of a purely predictive scheduling approach whereas a purely reactive scheduling approach leads to a lag in system response. When combined, rather than using complex stochastic programming (modelling/optimisation involving uncertainties and probabilities), the predictive scheduling element can ignore uncertainties and assume perfect knowledge on the connected EV population and the changes in power demand over the period being scheduled. The reactive scheduling element can then account for uncertainties and refine the schedule over time (the better the initial predictive scheduling, the less interference is needed from the reactive scheduling element).

The author's novel approach to combine predictive and reactive scheduling requires careful consideration within the high level aggregator control strategy to ensure system integrity and inter-compatibility. As multiple algorithms would share responsibility over the scheduling process sequencing and communication between the algorithms need to be carefully designed to avoid issues linked to race errors (where the behaviour of an algorithm may change depending on the order

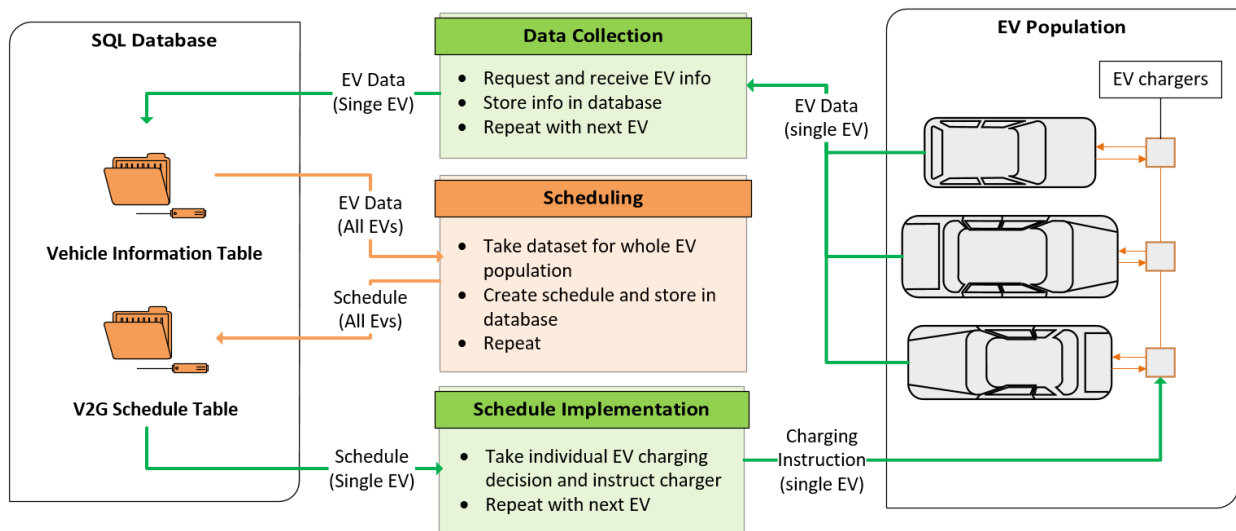


Fig. 2. Modular aggregator control structure with tasks of data collection, V2G scheduling and schedule implementation split into separate modules accessing a mutual database [11].

in which sub-routines are executed [26]).

For example, an EV charger might receive two instructions in the wrong order because a communication delay elsewhere on the network delayed the execution of one of the scheduling algorithms. Hence, the outlined V2G scheduling approach is embedded within a modular aggregator control structure which allows for real-time aggregator control with dynamic response to changes in power demand or connected EV population.

As an example on how multi-layer event-based V2G scheduling could be used, a scenario is presented in which a V2G network supports nearby electric rail traffic. Traction power demands of accelerating electric trains or excess power from decelerating trains utilising brake energy recovery are causing rapidly changing load patterns on the local power grid. A V2G network could support rail infrastructure locally by acting as a power buffer between rail system and power grid (directly represented by the local substation) as shown in Figure 1. This V2G application is being explored as part of the authors' involvement in the "TransEnergy - Road to Rail Energy Exchange (R2REE)" research project [27]. Various types of energy storage either on-board or along track lines have been proposed to accept power from regenerative braking for later use during acceleration [28]. Another approach is dwell time optimisation in which the departures and arrivals of trains on a network are synchronised so that accelerating trains accept power from decelerating ones [29] (significantly constraining train schedules).

In this V2G application, as electric trains accelerate (causing a spike in power demand for traction) the connected EV population (or parts thereof) would be discharged, feeding into the rail system and reducing the load on the local substation. As arriving trains decelerate using regenerative braking, the resulting spike in power from the rail system would be fed into the V2G network. Both these operations reduce fluctuations in the power demand experienced by the substation (similar in effect to the concept of dwell time optimisation [29] but without the constraints to train schedules). In periods without

rail traffic, the EV population can draw power from the shared grid connection for battery charging, thereby maintaining a steady power flow from the grid. Using a V2G network for electric rail support could offer a number of advantages:

- lowering grid connection upgrade requirements for new rail electrification projects (support from the V2G network could lower peaks in power demand from train acceleration experienced by the grid)
- lowering grid connection upgrade requirements for new EV charging infrastructure (assuming an EV car park can share the connection with the electrified rail infrastructure)
- enabling regenerative braking for electric trains where it has not been available before (leading to energy savings and potentially cost savings due to reduced wear on the mechanical brake systems[28])

II. MODULAR AGGREGATOR CONTROL STRUCTURE

The nature of V2G networks differs from most other energy storage technologies as the storage capacity as well as the power the system can provide vary significantly with the number and state of the connected EVs. Depending on the application, these parameters can be very difficult to predict as individual EVs may connect or disconnect from the network at any time. This is particularly true if V2G is implemented in a public setting rather than in a limited, strongly controlled environment (i.e. managing a commercial vehicle fleet with known driving schedules). An ideal aggregator control strategy is expected to be:

- **Dynamic:** The system can quickly adapt to changes in both the EV population and the power grid in real time.
- **Scalable:** The system can manage a wide range in numbers of connected EVs.
- **Robust:** A faulty communication route (i.e. to a single EV with malfunctioning hardware) does not "break" the whole system.

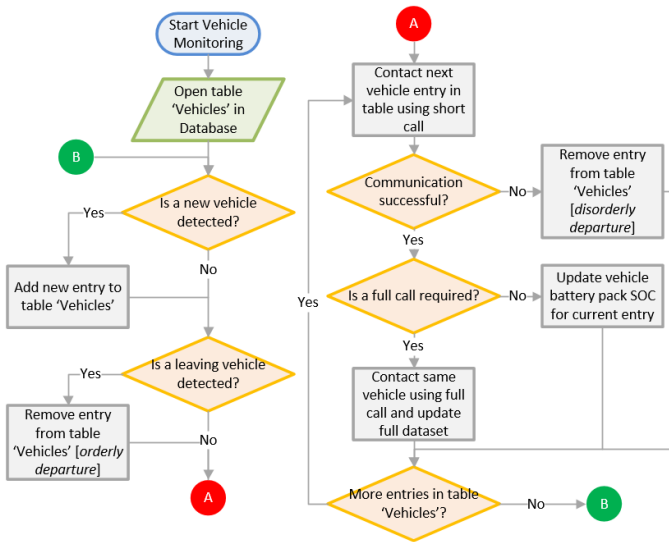


Fig. 3. Flow diagram of the data collection module responsible for regularly updating EV data in SQL database.

- **Computationally efficient:** The execution and sequencing of tasks is optimised to ensure low computational costs.
- **Compatible:** All parts of the V2G network comply with common hardware and software standards to ensure inter-compatibility.

In order to achieve the above, the authors are using the modular aggregator control structure described in [11] to simulate aggregator-to-EV communications and control algorithms overseeing a large-scale car parked based V2G network. The approach separates V2G aggregator control processes into data collection, scheduling and schedule implementation modules (see Figure 2) where each task can be shared between multiple modules (significant for the multi-layer scheduling approach outlined later in Section III). Any communication between modules is indirect through the use of a mutually accessible Structured Query Language (SQL) database.

The system is designed to enable dynamic real-time aggregator control and as such individual modules operate continuously without ever reaching a set end-state. Further the high level control structure does not force any fixed resolution or time steps on individual modules. This aids system responsiveness as modules are not forced to wait for each other, but could lead to computational inefficiencies and redundancy (for example, EV data might be updated more often than actually needed).

A. Data collection; input for scheduling modules

The scheduling routines make use of EV data from the mutual database to determine how each individual EV should be charged or discharged. This information is sourced and entered into the database beforehand through data collection modules which are executed independently following the routine in Figure 3. The necessary aggregator-to-EV communication uses either Ethernet or Wi-Fi connections along with the REST API [12][13] in a master-slave configuration. The

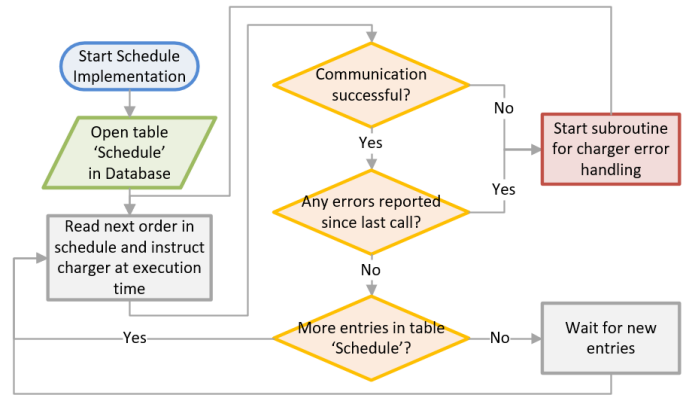


Fig. 4. Flow diagram of the schedule implementation module responsible for executing orders as specified in the schedule.

master (aggregator) requests EV data while the slave (EV) only responds to requests and never initiates communication. Here it is assumed that all connected EVs possess REST API compatible communication hardware and that EV battery packs are continuously monitored by the on-board battery management system.

Requests for EV data can take the form of a "full call" where the aggregator demands all relevant EV data (i.e. ID, battery pack capacity, maximum charging rate, etc.) or a "short call" where only the (continuously changing) battery pack state of charge is expected. While this data might be older for some EVs than for others (EVs are contacted sequentially), for scheduling it is assumed that any information in the database is accurate and up-to-date at the point of scheduling. The dataset for each EV on the V2G network is assumed to have the following format:

- **EV ID:** A unique integer value (using internal auto-increment database function) used to identify connected EVs.
- **IP address:** Unique internet protocol (IP) address of the charging hardware corresponding to each EV. It is assumed that each EV has its own network connected charger (either on-board or external).
- **Capacity:** EV battery pack capacity (in kWh) as an unsigned float value.
- **SOC:** EV battery pack State of Charge (in %) as an unsigned float value.
- **Charging rate:** Current charging rate of this EV (in kW) as a signed float value. A negative value represents discharging of the EV battery pack.
- **Maximum charging rate:** Assumed current maximum charging rate for this EV (in kW) as an unsigned float value. The value depends on the EV's battery pack (and may change with SOC) and the charging hardware. The lower limit shall apply.
- **Maximum discharging rate:** Assumed current maximum discharging rate for this EV (in kW) as an unsigned float value. Again, the value depends on the EV's battery pack and charging hardware. The lower limit applies.
- **Event status:** Binary value to signal if this EV is currently assigned to provide grid services for an event

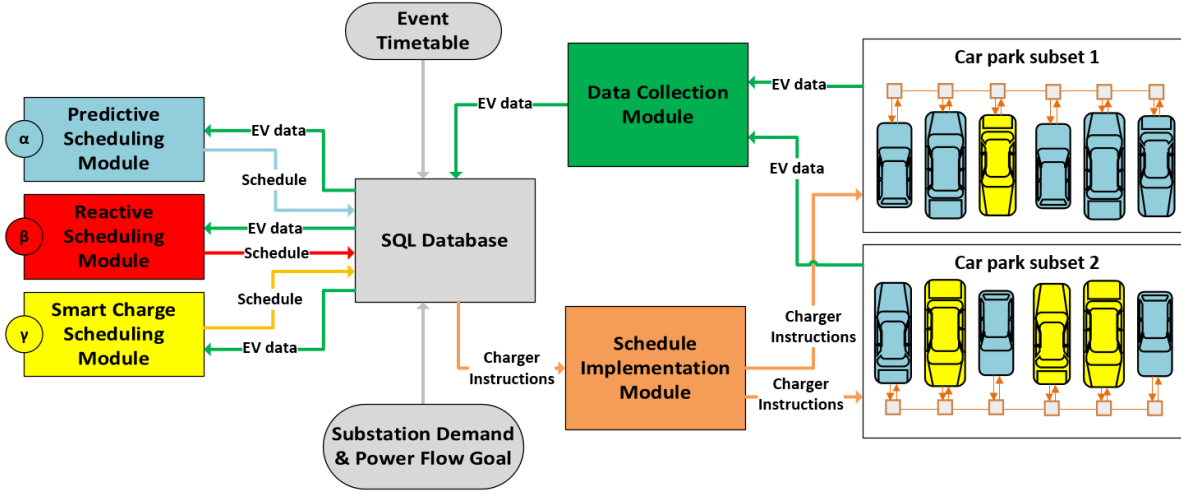


Fig. 5. Modular aggregator control structure using multiple independent scheduling modules.

(as described in section III). A value of 1 means the EV is assigned to an event (limiting its usage for other scheduling operations). A value of 0 means the EV is not currently used for grid services and can be assigned freely during scheduling.

- **Charge Weighting (CW)*:** An unsigned float value quantifying the EV's suitability to receive power (the higher the value, the higher the chance of this EV to be allocated to charging) - see equation 1. Value is zero when SOC is 100 % and thus an EV cannot be charged further. Derived from other parameters and calculated within the database.
- **Discharge Weighting (DCW)*:** An unsigned float value quantifying the EV's suitability to deliver power (the higher the value, the higher the chance of this EV to be allocated to discharging) - see equation 2. Value is zero when SOC is zero and thus an EV cannot be discharged further. Derived from other parameters and calculated within the database.

The underlying functions used to evaluate the dimensionless CW and DCW scores are highly dependent on the specific application. Generally, in situations when power has to be fed into the car park, EVs are particularly useful to the aggregator if their battery pack SOC is low, their battery pack capacity is high and their current maximum charging rate is high. Similarly, when power has to be supplied by the car park, EVs are particularly useful if their battery pack SOC is high, their battery pack capacity is high and their current maximum discharging rate is high. Hence, for the scope of this paper we can define:

$$CW = (1 - SOC) * \frac{Capacity}{Base\ Capacity} * \frac{Max\ Charging\ Rate}{Base\ Power\ Rating} \quad (1)$$

$$DCW = SOC * \frac{Capacity}{Base\ Capacity} * \frac{Max\ Discharging\ Rate}{Base\ Power\ Rating} \quad (2)$$

Base Capacity and Base Power Rating are chosen to be 1 kWh and 1 kW respectively. Depending on the application for V2G, these equations may be altered to include factors that are currently unaccounted for (i.e. information on future journeys, battery state of health, user preferences, etc.) or to increase the relative importance of specific parameters (i.e. increasing the weight of the maximum charging/discharging rates in high power applications). Using this approach of determining an EV's suitability to supply or receive power prior to the scheduling process aids overall system responsiveness as it reduces complexity during the scheduling process.

B. Schedule implementation; output from scheduling modules

Before discussing the scheduling process, the format of the anticipated output is to be determined. The schedule as defined in this paper is a list of charging instructions, or "orders", that are fed into the SQL database for subsequent handling by the schedule implementation module as shown in Figure 4. Again, aggregator-to-EV communication uses either Ethernet or Wi-Fi connections and the REST API in a master-slave configuration. Each order represents the instruction to change the charging rate of a specific EV to a given value at a given time. The SQL table "schedule" contains all orders in the format outlined below.

- **Order ID:** A unique integer value (using internal auto-increment database function) used to identify orders in cases where the schedule has to be revised.
- **Execution time:** Time at which this order is to be implemented in the datetime format. It is assumed that all devices on the network are synced following the universal coordinated time (UTC) time standard.
- **IP address:** Unique IP address of the charging hardware to receive the order in binary format. It is assumed that each EV has its own network connected charger.
- **Charging rate:** Anticipated charging rate (in kW) as a signed float value. A negative value represents an

instruction to discharge the EV battery pack at the given rate.

- **Order status:** Data string used for schedule implementation/revision. If empty, this order is not yet due for execution and could be revised/cancelled if necessary. If "queued", the (or one of multiple) implementation module(s) picked up and is about to execute this order (to avoid contradictory instructions and race errors [26] this order cannot be altered). If "cancelled", this order was revised by subsequent scheduling and will not be executed.

III. EVENT-BASED V2G SCHEDULING

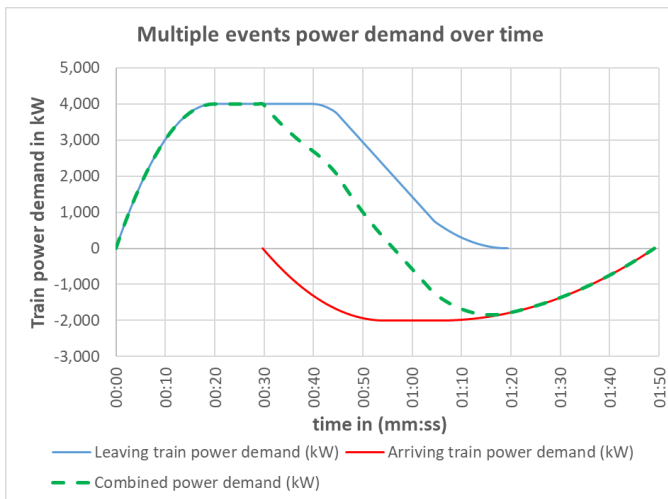


Fig. 6. Combined effects to two stacked "counter-events" on the rail station power demand

In general, V2G scheduling algorithms have to consider two sides: the power grid and the EV population being managed. The requirements of both sides need to be satisfied to perform a given grid service and simultaneously ensure EVs are sufficiently charged for mobility purposes. Further, both sides are generally subject to uncertainties, although some re-occurring patterns in power demand and supply might be distinguishable.

This work aims to exploit such patterns on the power demand side to enable dynamic fast-response V2G scheduling which is particularly important in applications where power demands change rapidly and a large number of connected EVs have to be managed. This is achieved by combining multiple scheduling algorithm layers within the modular aggregator control structure (shown in Figure 5) combining both predictive and reactive scheduling to mitigate the disadvantages of each approach (discussed later in this section).

The example application used in this work is the support of nearby electrified rail infrastructure where the V2G network provides power to electric trains leaving a train station and accepts power from arriving trains utilising regenerative braking. In this application, the timing and magnitude of the rail power demands is predictable due to timetabled train operation and ongoing train tracking. The departure or arrival of a train with

known type and speed pattern can be regarded as a reoccurring event.

Delays and changes to train schedules are still a source of uncertainty but the tracking of trains ensures predictive capability in the short term. Here, it is assumed that the position of trains on the rail network is continuously monitored. While this is not always the case yet, positional tracking is commonly used as part of Positive Train Control (PTC) systems [30] and might utilise GPS (Global Positioning System) [31][32] or alternative systems such as the Galileo satellite navigation system [31].

Within this work an event is defined as either the arrival or the departure of an electric train at a train station, where train type and speed pattern are known. It is assumed that two similar events, for example two identical trains accelerating at the same rate will result in similar power demands over time, even if the two events are hours apart. It is further assumed, that each event has a known beginning and end time (i.e. the time at which a train moves out of the V2G network's range and into the next track section supplied by another substation). Examples of assumed power demand curves for trains arriving at or leaving a train station are shown in Figure 6.

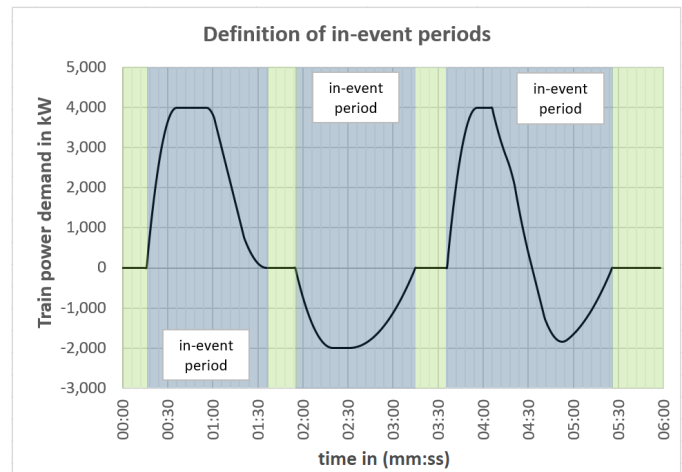


Fig. 7. Visualisation of in-event (blue) and out-of-event (green) periods for a sequence of events (left to right: single train accelerating, single train braking, train accelerating while another is braking).

As will be shown later, scheduling for events takes place sequentially so the power demands of two (or more) overlapping events (happening simultaneously) may stack or, partially, cancel each other out (as exploited where dwell-time optimisation is employed [29]). For further discussion, two events are defined as "co-active" when either, both require power flow from the EV car park (i.e. both departing trains) or both require power flow into the car park (i.e. both arriving trains). In contrast, two events are defined as "counter-active" when one event has a positive power demand and the other one a negative power demand (i.e. one train departing, one train arriving). It follows that scheduling differentiates between in-event periods (during which an event is ongoing) and out-of-event periods where no events are happening and different scheduling rules may apply, see Figure 7. Further, any predictive scheduling is based on current conditions of

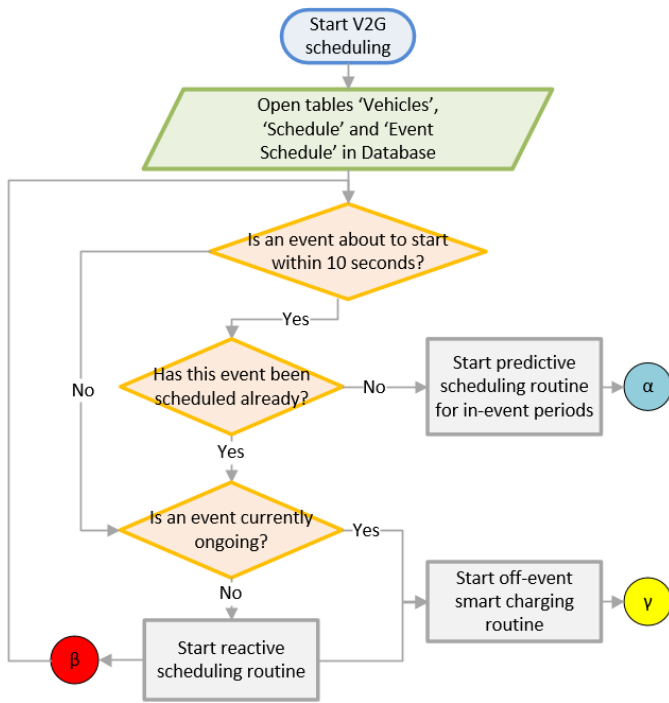


Fig. 8. Three layer scheduling, top-level control routine

the connected EV population. Hence, a degree of uncertainty remains and predictive scheduling will not yield perfect results (EVs may unexpectedly disconnect or faulty communication prevents timely schedule implementation).

This is mitigated by combining predictive scheduling with a reactive element used to adjust for inaccuracies and the eventual mismatch in power demand and supply in real-time. As a result, V2G scheduling is separated into three layers, operating at times as shown in Figure 7 and prioritised as such:

- **High priority:** In-event predictive scheduling (first layer, rough scheduling)
- **Medium priority:** Ongoing reactive scheduling (second layer, finer adjustments)
- **Low priority:** Off-event smart charging scheduling (third layer, rough scheduling for bulk of EVs)

The overarching control over how and when these layers are engaged follows the routine shown in Figure 8. For simplification all the following discussion assumes that all connected EVs are forced to participate in the V2G operation and that EVs are already assessed in terms of their suitability to receive or provide electricity (As quantified by CW and DCW). From the latter assumption it follows that no further analysis of or comparison between EVs is necessary during scheduling - the order in which EVs are assigned to charging/discharging is directly linked to the respective CW and DCW scores. As a further simplification, unlike in most other scheduling strategies, economic factors such as electricity prices are not considered in the scheduling process. However, charging costs and compensation for participation in this V2G application would be an interesting topic in itself and should be discussed in future work in this field.

A. Scheduling for in-event periods, first layer

The first scheduling layer is responsible for managing the V2G network's response to events occurring in the near future (i.e. in a few seconds). It creates the schedule for the whole duration of an anticipated event before its occurrence. This layer can be initialised once the starting time of the next event is known. Within this work it is assumed that due to railway signalling systems and GPS tracking train departure/arrival times are known at least 10 seconds beforehand, which is when this algorithm is initialised. In practice, this time may vary between different rail systems and does not have to be constant as long as enough time is available for the execution of this scheduling layer (see Section IV).

A train departure event begins as the train starts accelerating, causing a spike in power demand. The peak value can vary widely depending mainly on the train's weight and acceleration, but is usually on the scale of a few megawatts (peaks of about 3.6 MW for both traction power drawn and brake energy provided in [29], about 2 MW and 1.2 MW respectively in [33]). As the power demand increases quickly, system responsiveness is paramount. However, due to the scale of power required, a large number of EVs is needed to match the demand and each aggregator-to-EV communication attempt takes time (around 15 milliseconds [11]) and has to happen sequentially. Thus the number of communication attempts is to be kept low. In order to do this, the first layer of scheduling only assigns EVs to start charging/discharging at their respective maximum charging/discharging rates, rather than slowly and continuously ramping up charging/discharging rates for all suitable EVs. At this stage, it is assumed that any EV can switch from its previous charging/discharging rate to its current maximum instantaneously at the time determined in the schedule. Any delay is being mitigated later as the reactive scheduling adjusts power demands in real-time.

A flow diagram of the algorithm used in this layer is shown in Figure 9. As a first step the expected power demand profile for the next event is being loaded into memory. Next all relevant EV data is being obtained from the network-wide mutual SQL database (IP address, maximum charge/discharge rate, current charge/discharge rate, CW/DCW). For train departure events, only EVs suitable for discharge (DCW above a given threshold value) are considered and EV data is sorted by DCW from high to low so that the most suitable EVs are being assigned first. Similarly, for train arrival events EV data is sorted by CW from high to low - only EVs with a CW of zero (signalling full battery) are ignored.

Next the algorithm iterates through each time step in the power demand curve. The temporal resolution is a major factor determining computational complexity of this layer as it determines how often the main loop shown in Figure 9 is executed. In this work one second time step are used. Shorter time steps could be used for more refined scheduling at the expense of increased computational cost. The second major driver of computational complexity is the number of EVs loaded into memory. This impacts the execution time of each iteration of the main loop (as each step in the red block in Figure 9 requires the manipulation of a larger dataset).

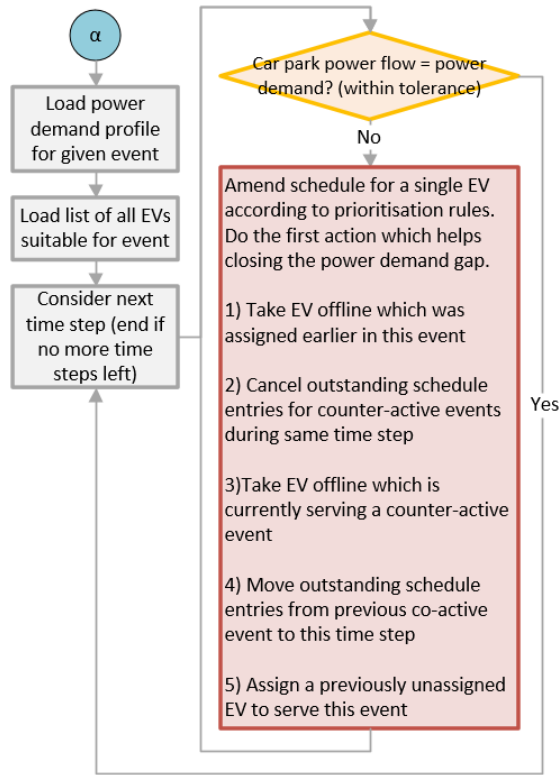


Fig. 9. Predictive scheduling layer flow diagram

Assuming only one event is happening at the time of execution, the algorithm will sequentially assign EVs to charge/discharge at maximum rate to match the change in power demand between each time step. Further, assigned EVs are marked as "busy" in the mutual data base (set "Event Status" to 1) to prevent conflicting instructions being given by another scheduling algorithm. As the event comes to an end and power demands decrease, these EVs are set to return to their previous charge/discharge rates. As a result, for each EV being used in an event two new charging instructions are added to the schedule table.

In more complicated cases where multiple events overlap, this scheduling layer also has to consider previous schedule entries. Each event is being scheduled without consideration of future events. Consequently, schedule entries from a previous event that is still ongoing may contradict and interfere with the event currently being scheduled. The predictive scheduling layer is designed to correct previous schedule entries if applicable before assigning new EVs. This avoids situations where a subset of the EV population is being discharged to support a departing train while another subset is simultaneously being charged to accept power from an arriving train.

Thus any available actions to close the gap between power demand and power flow from the EV car park are prioritised in the following order:

- 1) Take any EVs offline (return to pre-event charging rate) that were assigned in this event. This is the top priority to ensure that every EV assignment is reversed as the event comes to an end.

- 2) Cancel any outstanding charge/discharge orders scheduled for this time step to serve counter-event. This revision of a previous schedule prevents conflicting charging instructions from being executed in the future. Only possible for order not yet queued by the implementation module(s).
- 3) Take any EVs offline that are already charging/discharging at this time step to serve counter-event. This schedule revision moves forward the execution time of an already planned order to end an EV assignment.
- 4) Move orders from any previous co-active event forward/backward. This revision leads to an already planned EV assignment to be executed earlier or cancelled later. Only the extra time period of the EV's assignment is counted towards the current event.
- 5) Assign an EV that is not currently serving any event to charge/discharge at maximum rate. Only when no useful revisions of previous scheduling decisions are available are additional EVs being assigned.
- 6) If there is still a power gap, the current population of EVs cannot fully serve the event as not enough suitable EVs are available. Previous assignments are still valid and the scheduling algorithm will still try to meet demands as far as possible for subsequent time steps.

B. Reactive scheduling for schedule refinement, second layer

The second scheduling layer is a reactive one continuously implementing minor corrections in real-time to ensure the power flow into/out of the EV car park matches the application's power demand (within a given tolerance, here a value of 5 kW is being used as tolerance). The corresponding algorithm routine is shown in Figure 10. Corrections may be necessary due to uncertainty at the time of the predictive schedule creation (EV data may be outdated, an EV might have unexpectedly disconnected in the mean time, the power demand may differ from the initial predictions, etc.).

In contrast to the first layer, the second one does not follow fixed time steps, but instead loops continuously (the temporal resolution does not determine execution times but depends on them). Further, this layer is capable of assigning charging rates other than an EVs respective maximum charging or discharging rate at the time of scheduling (if a lower rate for a single EV is enough to close the gap between power demand and supply). This gives it the ability for finer adjustments to the power flow into and out of the EV population. Revisions to the existing charging schedule or new charging instructions are fed into the database to be executed by the implementation module.

As was the case in the first layer, the size of the EV population impacts computational complexity through the size of the dataset being manipulated. However, any schedule revision only takes place if the mismatch between car park power flow and application power demand exceeds the stated tolerance (see Figure 10). It follows that computational complexity of this second layer is dependent on the accuracy of the preceding predictive scheduling. If the predicted power demand in the previous layer matches the actual power demand (and no active

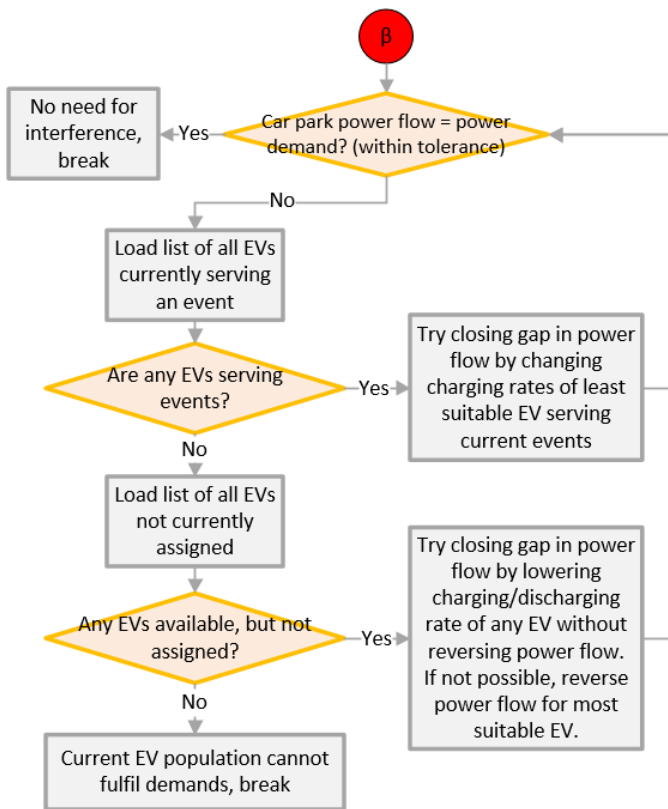


Fig. 10. Reactive scheduling layer flow diagram

EVs have suddenly disconnected since) the reactive layer will not interfere (hence, no need for further communication with database to load EV data or manipulate the existing schedule).

C. Scheduling for non-event periods (smart charging), third layer

The third scheduling layer applies a different set of rules outside of events. Depending on the V2G application these rules may differ. Considering the electric rail support, this layer is used as a smart charging scheduler. This is necessary as the rules of the other scheduling layers inevitably lead to a situation where, on average, more power is drawn from the EV population during events than is supplied to the EVs (arriving trains are expected to supply significantly less energy from regenerative braking than they require for acceleration - around a third of the traction energy according to [34]). However, the aggregator is responsible for ensuring EVs are sufficiently charged for mobility purposes and must (at least over time) receive a net charge.

A flow diagram describing the smart charging algorithm routine is shown in Figure 11. This layer requires two additional parameters (both of which may change throughout the day): a minimum charging rate per EV and an optimal total power flow from the grid (the power drawn by the EV car park plus train station through a shared grid connection). Both values are obtained from the database and have to be chosen based on car park size and V2G application. Ideally, the minimum charging rate on its own is sufficient to ensure EVs are being charged over time.

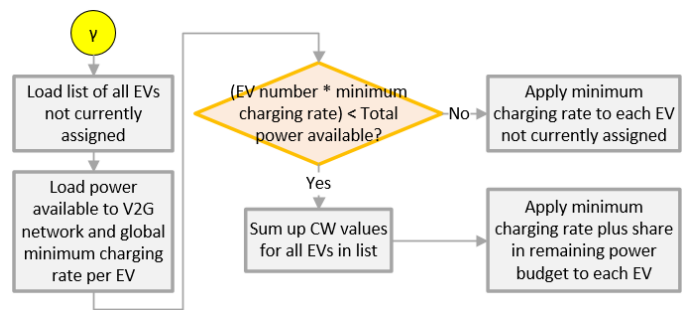


Fig. 11. Smart charging scheduling layer flow diagram

The algorithm is designed to make full use of the power available to the EV car park. Individual EVs are assigned the minimum charging rate plus a share of the remaining power available that is proportional to each EV's CW value (up to the maximum charging rate of each EV). The higher the CW, the more power is being allocated. The computational cost of this algorithm depends mainly on the size of the connected EV population. As decisions are based on the pre-assigned CW scores of EVs no in-depth analysis of EVs is required. As this layer operates in non-event periods where system wide power flows only change gradually and system responsiveness is less significant, computational cost are not a major concern. Therefore alternative, more complex smart charging algorithms may be employed instead (taking into account additional factors such as EV battery degradation, anticipated EV journeys or fairness criteria).

IV. INTERACTION BETWEEN PREDICTIVE AND REACTIVE SCHEDULING LAYERS

In order to analyse the behaviour of the scheduling algorithms presented, the same combination of events (a train departure followed by a train arrival 30 seconds later) are scheduled and simulated using A) only predictive scheduling and B) only reactive scheduling. Finally, the same event combination (altered with a braking manoeuvre unknown to the predictive layer) is being scheduled for using C) combined predictive and reactive scheduling.

The scheduling algorithms have been tested using two desktop PCs connected to the same computer network via Ethernet connection. The computer network is not a dedicated one so traffic outside of the authors' control can have a minor impact on the results reported. All aggregator control algorithms (data collection, scheduling, schedule implementation) were handled on a machine equipped with an Intel Core i5-2400 CPU (4x3.1 GHz) and 4 gigabyte DDR3 Random Access Memory (RAM). The memory usage was carefully monitored to ensure algorithm execution is not "bottlenecked" and slowed down by a lack of memory (typically around 70 % of memory was in use during operation).

The EV population was simulated on another machine equipped with an Intel Core i5-4590 CPU (4x3.3 GHz) and 16 gigabyte DDR3 RAM. Available memory is the major restriction for the number of EVs that can be simulated (typically around 90 to 95 % of the available RAM is reserved while simulating 1,000 EVs).

Each simulated EV has a unique IP address and network port combination so that aggregator-to-EV communication via the REST API realistically mirrors the communication delays that could be expected on a network with a real EV population. It was found that data collection (requesting information from an EV, receiving information and storing it in the database) as well as schedule implementation (retrieving charging instructions from the database, submitting instructions to EV charger and receiving confirmation) each take about 15 milliseconds per EV using an Ethernet connection.

A. Predictive scheduling only




Fig. 12. Simulated V2G network response to two counter-active events (one train departing at 0 seconds, one train arriving at 30 seconds) using predictive scheduling only

Initially, only the predictive scheduling layer is being tested (i.e. the reactive layer is not enabled). The predictive layer is set up to begin scheduling 10 seconds before an event begins and to create the schedule for the whole duration of an event in 1 second time steps. The execution times for this layer (time passed between the program initialisation and the schedule for a single event being fully passed onto the database) vary significantly with the number of EVs on the network (see Table I).

The relationship between execution time and number of EVs on the network is non-linear as for each time step the algorithm

Number of EVs	Min. time (s)	Average time (s)	Max. time (s)
50	0.007	0.034	0.496
100	0.011	0.353	0.765
200	0.147	0.527	0.844
500	0.328	0.649	1.141
1000	0.469	0.837	1.820

TABLE I
EXECUTION TIMES FOR PREDICTIVE SCHEDULING ALGORITHM FOR VARYING NUMBER OF CONNECTED EVs (OVER 1,000 SCHEDULING CYCLES)

stops when a solution is reached or no more EVs are available. This means that for a low number of EVs, the scheduling process may be fast, but the power demand may not have been met. Similarly, for a large number of EVs, solutions might be reached with less EVs than available and the execution times only increase as more EV data had been loaded into memory from the database.

It should be noted that even the longest execution time reported in Table I with 1.82 seconds is well below the 10 seconds made available to the algorithm before the beginning of the event. Thus the predictive scheduling could take place closer to the beginning of an event (which might lead to less uncertainty at the point of scheduling) or the time available could be used for more complex scheduling (smaller time steps, more complex scheduling rules, smaller increments in charging rates assigned per EV, etc.).

Figure 12 shows the V2G network response to the two counter-active events for a connected EV population of 500 simulated EVs. The rail traffic power demand over time represents an input to the scheduling process. The network response is represented by the sum of all individual EV power flows over time. This sum changes whenever a new charging/discharging instruction is being sent to an EV by the schedule implementation modules (i.e. when the V2G aggregator is taking action in response to a change in power demand).

For the purely predictive scheduling approach, the simulations show that the power supplied by the EV population matches the rail power demand very well and without significant delay. Delays are being avoided here as the schedule has been entered into the database and made available to the schedule implementation algorithm well in advance (10 seconds minus execution time). Minor mismatch in power supply and demand exist due to the predictive layer's limitation of assigning EVs at maximum charging or discharging rate.

However, it must be noted that the predictive scheduling only followed the predicted power demand curves defined for departing or arriving trains. Any noise or deviation from these predictions on the power system have not been part of the simulation, thus the initial prediction was in fact a perfect one within this simulation (deviation in power demand from the initial prediction will be addressed when combining predictive and reactive scheduling in subsection C).

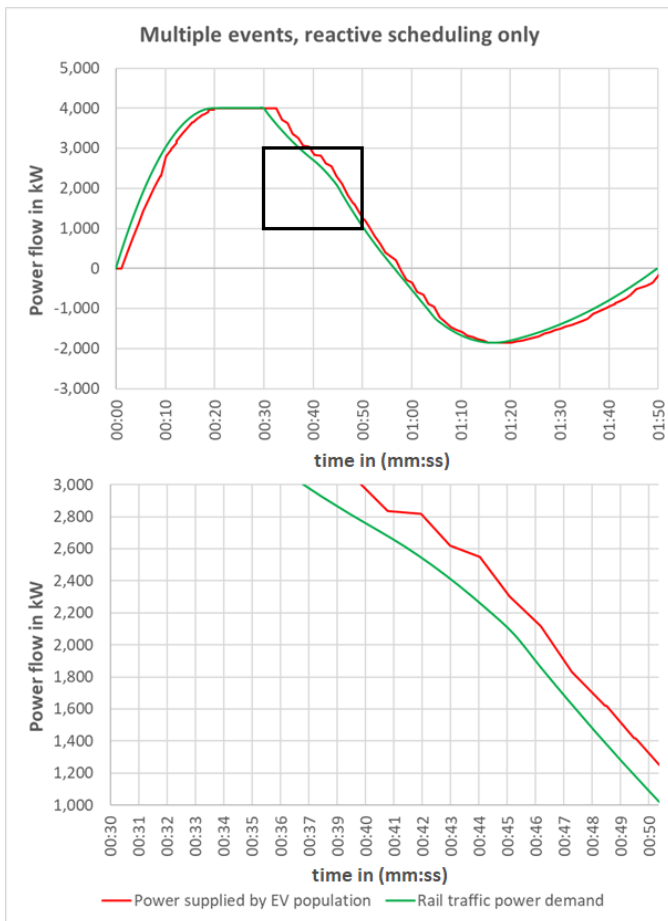


Fig. 13. Simulated V2G network response to two counter-active events (one train departing at 0 seconds, one train arriving at 30 seconds) using reactive scheduling only

B. Reactive scheduling only

Similar to the test in subsection A, the reactive scheduling layer has been tested by scheduling for the same event combination while the predictive layer has been disabled. The results are shown in Figure 13 (an identical simulated EV population of 500 EVs has been used). The figure shows that the V2G network response is following the power demand of the rail traffic, but that the response is continuously lagging behind by between 1 and 2.5 seconds. This lag is only in parts due to the actual schedule creation and primarily caused by communication delays within the network.

The execution times of the reactive scheduling layer were found to vary between about 47 and 187 milliseconds with an average execution time of 73 milliseconds (measured over 10,000 scheduling cycles, in-event periods only with 500 connected EVs). On average, 48 milliseconds of this time (between 31 and 140 milliseconds) was required just to determine the difference in power flow from the EV population and the power demand from the rail application (see Section III, Figure 10).

This information is sourced from the SQL database and delays are therefore due to database communication, database-internal computation and sequencing of SQL queries (as other

algorithms simultaneously access the same database - see Section II, Figure 5).

In each cycle the algorithm made between 1 and 21 changes to the schedule – on average 12.88 schedule changes per cycle. These new schedule entries in the database have to be processed by the implementation module and communicated to the EV chargers (see Section II B.). At around 15 milliseconds per EV, implementing 13 schedule entries is expected to take about 195 milliseconds.

The last source of delay in system response is the time taken to detect the effects of any changes in power flow within the EV population. The power flow data relied upon for scheduling originates from the data collection module (see Section II A.) which updates EV data sequentially. For a population of 500 EVs it was found that, on average, 3.04 seconds pass before all EVs have been checked. This delay is highly situational depending on when an EV is being contacted. Thus, any schedule change made may not be accounted for in subsequent scheduling cycles for a few seconds.

C. Combined scheduling with predictive and reactive scheduling layer

In the next test both the reactive and predictive scheduling layers are enabled. Again, the algorithms are determining the V2G network's response to a departing train followed by an arriving train 30 seconds later. To show the ability of this combined scheduling approach to adjust to uncertainty, the rail traffic power demand has been altered to differ from the predicted power demand curves.

As shown in Figure 14, the simulation now features a drop in the power provided from brake energy recovery of the arriving train from 1 minute and 17 seconds onwards. This could present a situation in which a train engages its mechanical brakes for an emergency stop, drastically reducing brake energy recovery. As this drop was unknown to the predictive scheduling layer, the initial schedule for charging/discharging the population of 500 EVs has been identical to that in subsection A (thus created at the same computational cost). This initial schedule has then been adjusted by the reactive layer once a mismatch above the 5 kW tolerance between power supply and demand has been detected.

Figure 14 shows that the V2G network's response closely matches the power demand without any significant delay until the unpredictable braking manoeuvre, followed by a relatively consistent delay in network response of about 1 second as the reactive scheduling layer makes ongoing adjustments. Thus, the system responsiveness has been improved as long as the initial prediction closely matches the actual power demand. Yet, the ability to account for uncertainty has been maintained by the reactive scheduling layer.

V. CONCLUSION

In this work a novel V2G scheduling strategy was presented that exploits repetitive power demand patterns and the resulting reduced uncertainty from reoccurring events to enable predictive scheduling capability. In this context, an event was




Fig. 14. Simulated V2G network response to two counter-active events (one train departing at 0 seconds, one train arriving at 30 seconds with sudden unpredicted braking) using both predictive and reactive scheduling

defined as a time limited, predictable pattern of power demand over time with known start and end time.

The scheduling strategy is proposed for large scale, centrally controlled V2G networks and consists of three distinct scheduling layers. The predictive scheduling takes place shortly before an event commences with each event being scheduled independently. As perfect knowledge is being assumed (over the event duration) as this layer is executed, a second reactive scheduling layer is utilised to refine the schedule in real-time. A third layer is used to apply a separate set of scheduling rules in periods where no events take place.

This strategy was presented around an example application in which a V2G network is used to support the power demands of nearby electric rail traffic. In this context events were defined as the departure or arrival of electric trains at a station resulting in a sudden surge in power demand under train acceleration or excess power under train deceleration enabled by regenerative braking.

The algorithms presented were developed around a modular aggregator control strategy that uses indirect communication via a mutual network-wide SQL database. To reduce computational cost of the scheduling process the suitability of EVs to serve each type of events was pre-determined within the

database (CW and DCW).

The assessment criteria for EV suitability greatly depend on the specific application of V2G and constraints on the EV population. Scheduling during events (first layer) was resolved in one second time steps - a sub-second resolution could be employed however at greater computational cost. Reactive scheduling (second layer) and smart charging (third layer) were not bound by fixed time steps.

It was shown that the multi-layer scheduling approach can lead to a quasi instantaneous system responsiveness as long as power demands can be accurately predicted. Lags in system response are still present if uncertainty leads to a mismatch between expected and actual power demands, but the length of such lags depends on the magnitude of the mismatch (determining how many schedule adjustments need to be made by the reactive scheduling layer).

ACKNOWLEDGMENT

The authors would like to thank the Engineering and Physical Sciences Research Council (EPSRC) Centre for Doctoral Training in "Energy Storage and Its Applications" (EP/L016818/1) and the "Transenergy – Road to Rail Energy Exchange (R2REE)" (EP/N022289/1) projects for supporting the work presented.

REFERENCES

- [1] International Energy Agency, "Global ev outlook 2019," *IEA publications*, pp. 09–10, 2019.
- [2] P. Lienert, "Global carmakers to invest at least \$90 billion in electric vehicles," *Reuters Business News*, 15 Jan 2018.
- [3] M. Rocco, "Volkswagen to invest \$800m in us electric vehicle production," *Financial Times*, 14 Jan 2019.
- [4] P. McGee, "Porsche accelerates race for electric market share," *Financial Times*, 22 Feb 2019.
- [5] S. Falahati, S. A. Taher, and M. Shahidepour, "A new smart charging method for evs for frequency control of smart grid," *International Journal of Electrical Power & Energy Systems*, vol. 83, pp. 458 – 469, 2016.
- [6] B. Yagcitekcin and M. Uzunoglu, "A double-layer smart charging strategy of electric vehicles taking routing and charge scheduling into account," *Applied Energy*, vol. 167, pp. 407 – 419, 2016.
- [7] C. O. Adika and L. Wang, "Smart charging and appliance scheduling approaches to demand side management," *International Journal of Electrical Power & Energy Systems*, vol. 57, pp. 232 – 240, 2014.
- [8] W. Kempton and J. Tomić, "Vehicle-to-grid power fundamentals: Calculating capacity and net revenue," *Journal of Power Sources*, vol. 144, no. 1, pp. 268 – 279, 2005.
- [9] J. Tomić 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] J. Geske and D. Schumann, "Willing to participate in vehicle-to-grid (v2g)? why not!," *Energy Policy*, vol. 120, pp. 392 – 401, 2018.
- [11] H. Krueger and A. Cruden, "Modular strategy for aggregator control and data exchange in large scale vehicle-to-grid (v2g) applications," *Energy Procedia*, vol. 151, pp. 7 – 11, 2018, 3rd Annual Conference in Energy Storage and Its Applications, 3rd CDT-ESA-AC, 11–12 September 2018, The University of Sheffield, UK.
- [12] M. Paredes-Valverde, G. Alor-Hernández, A. Rodr'iguez-González, and G. Hernández-Chan, "Developing social networks mashups: An overview of rest-based apis," *Procedia Technology*, vol. 3, p. 205–213, 12 2012.
- [13] R. Battle and E. Benson, "Bridging the semantic web and web 2.0 with representational state transfer (rest)," *Web Semantics: Science, Services and Agents on the World Wide Web*, vol. 6, pp. 61–69, 02 2008.


- [14] R. Shi, Y. Yang, L. Shi, and K. Y. Lee, "Bi-level day ahead optimization of v2g dispatch strategy based on the dynamic discharging electricity price," *IFAC-PapersOnLine*, vol. 51, no. 28, pp. 462 – 467, 2018. 10th IFAC Symposium on Control of Power and Energy Systems CPES 2018.
- [15] L. Wang, S. Sharkh, A. Chipperfield, and A. Cruden, "Dispatch of vehicle-to-grid battery storage using an analytic hierarchy process," *IEEE Transactions on Vehicular Technology*, vol. 66, pp. 2952–2965, April 2017.
- [16] E. Sortomme and M. A. El-Sharkawi, "Optimal scheduling of vehicle-to-grid energy and ancillary services," *IEEE Transactions on Smart Grid*, vol. 3, pp. 351–359, March 2012.
- [17] S. Faddel, A. Aldeek, A. T. Al-Awami, E. Sortomme, and Z. Al-Hamouz, "Ancillary services bidding for uncertain bidirectional v2g using fuzzy linear programming," *Energy*, vol. 160, pp. 986 – 995, 2018.
- [18] Y. He, B. Venkatesh, and L. Guan, "Optimal scheduling for charging and discharging of electric vehicles," *IEEE Transactions on Smart Grid*, vol. 3, pp. 1095–1105, Sep. 2012.
- [19] H. Turker and S. Bacha, "Optimal minimization of plug-in electric vehicle charging cost with vehicle-to-home and vehicle-to-grid concepts," *IEEE Transactions on Vehicular Technology*, vol. 67, pp. 10281–10292, Nov 2018.
- [20] A. Ghosh and V. Aggarwal, "Menu-based pricing for charging of electric vehicles with vehicle-to-grid service," *IEEE Transactions on Vehicular Technology*, vol. 67, pp. 10268–10280, Nov 2018.
- [21] S. Gao, K. T. Chau, C. Liu, D. Wu, and C. C. Chan, "Integrated energy management of plug-in electric vehicles in power grid with renewables," *IEEE Transactions on Vehicular Technology*, vol. 63, pp. 3019–3027, Sep. 2014.
- [22] J. J. Escudero-Garzas, A. Garcia-Armada, and G. Seco-Granados, "Fair design of plug-in electric vehicles aggregator for v2g regulation," *IEEE Transactions on Vehicular Technology*, vol. 61, pp. 3406–3419, Oct 2012.
- [23] Energy Storage Association, "Frequency regulation - executive summary," *Energy Storage Association, Online Resources*, 2019. url: <http://energystorage.org/energy-storage/technology-applications/frequency-regulation>, accessed on 04.03.2019.
- [24] D. Greenwood, K. Lim, C. Patsios, P. Lyons, Y. Lim, and P. Taylor, "Frequency response services designed for energy storage," *Applied Energy*, vol. 203, pp. 115 – 127, 2017.
- [25] A. Abdallah and X. S. Shen, "Lightweight authentication and privacy-preserving scheme for v2g connections," *IEEE Transactions on Vehicular Technology*, vol. 66, pp. 2615–2629, March 2017.
- [26] C. Q. Adamsen, A. Møller, R. Karim, M. Sridharan, F. Tip, and K. Sen, "Repairing event race errors by controlling nondeterminism," *Proceedings of the 39th International Conference on Software Engineering*, pp. 289–299, 2017.
- [27] D. Stone, "Details of grant, transenergy - road to rail energy exchange (r2ree)," *Engineering and Physical Sciences Research Council (EPSRC) Online Resources*, 2015. url: <http://gow.epsrc.ac.uk/NGBOViewGrant.aspx?GrantRef=EP/N022289/1>, accessed on 27.02.2018.
- [28] M. Ceraolo, G. Lutzemberger, E. Meli, L. Pugi, A. Rindi, and G. Pancari, "Energy storage systems to exploit regenerative braking in dc railway systems: Different approaches to improve efficiency of modern high-speed trains," *Journal of Energy Storage*, vol. 16, pp. 269 – 279, 2018.
- [29] F. Lin, S. Liu, Z. Yang, Y. Zhao, Z. Yang, and H. Sun, "Multi-train energy saving for maximum usage of regenerative energy by dwell time optimization in urban rail transit using genetic algorithm," *Energies*, vol. 9, p. 208, 03 2016.
- [30] Z. Zhang, X. Liu, and K. Holt, "Positive train control (ptc) for railway safety in the united states: Policy developments and critical issues," *Utilities Policy*, vol. 51, pp. 33 – 40, 2018.
- [31] V. Behrends, M. Haunschild, and N. Galonske, "Smart telematics enabling efficient rail transport – development of the viwas research and development project," *Transportation Research Procedia*, vol. 14, pp. 4430 – 4439, 2016. Transport Research Arena TRA2016.
- [32] R. I. Rajkumar, P. E. Sankaranarayanan, and G. Sundari, "Gps and ethernet based real time train tracking system," in *2013 International Conference on Advanced Electronic Systems (ICAES)*, pp. 282–286, Sep. 2013.
- [33] S. Fazel, S. Firouzian, and B. Khalkhali Shandiz, "Energy-efficient emplacement of reversible dc traction power substations in urban rail-transport through regenerative energy recovery," *International Journal of Railway Research*, vol. 1, no. 2, 2014.
- [34] A. González-Gil, R. Palacin, P. Batty, and J. Powell, "A systems approach to reduce urban rail energy consumption," *Energy Conversion and Management*, vol. 80, pp. 509–524, 2014.



Hannes Krueger received the MEng. degree in Mechanical Engineering from the University of Southampton, Southampton, U.K. in 2015.

He is currently working towards the Ph.D. degree with the Faculty of Engineering and Physical Sciences at the University of Southampton as part of the EPSRC Centre for Doctoral Training in Energy Storage and Its Applications.

His research interests include electric vehicles, Vehicle-to-Grid, electric rail systems, smart grids, communication technologies and data management.



Andrew Cruden received the B.Eng., M.Sc., and Ph.D. degrees from the University of Strathclyde, Glasgow, U.K., in 1989, 1990, and 1998, respectively, all in electronic and electrical engineering.

He is currently the Head of the Energy Technologies Research Group and a Professor of Energy Technologies with the University of Southampton. He has significant experience in the field of renewable energy, particularly in fuel cell technology and condition monitoring of wind turbines.

His interest in fuel cell technology led, indirectly, to an interest in the field of electric vehicles and, subsequently, to research activity in energy storage and traction drives.