V2G Dynamic Headroom Control

Implementation Risk Modelling Results

Version 1.0

19th December 2025

Prepared by:

Andrew Urquhart and Murray Thomson

CREST, Loughborough University

Prepared for:

Liza Troshka

National Grid Electricity Distribution

Contents

1. Introduction

This report describes modelling activities for work package 3 of the V2G Dynamic Headroom Control project.

Work packages 1 and 2 of this project demonstrated the benefits of controlling exports to the grid from electric vehicle (EV) charge points with Vehicle-to-Grid (V2G) capability. There are risks associated with the implementation of the algorithms proposed for V2G control and these may be explored in a future trial of V2G control methods. Ahead of any trials, and in preparation, simulation methods have been used to explore some of these risks in more detail. The results of these simulations may now inform the design of a hardware trial and would help to avoid any obviously foreseeable problems.

Work package 3 deliverable 1 outlined a set of potential risks and methods that could be used to address these. Where possible, the modelling described there has been developed and results of these simulations are contained here. In some cases, these take a slightly different approach to that which was anticipated.

The work here covers the following operational aspects of V2G control:

Voltage management at transformers

On-load tap changers (OLTC) could be used to manage the substation busbar voltage, reducing variations in the busbar voltage due to voltage drops along the HV feeder, but also possibly managing the substation voltage to increase the margin for voltage rise or drop on the LV feeders according to the anticipated demand. This could reduce the requirement for voltage management by V2G control, potentially avoiding the need for V2G control completely. Even with V2G control, voltage management using OLTCs could reduce the frequency with which the exported power is constrained.

There is also a risk that the dynamic control of voltage from OLTCs could interact negatively with the real-time operation of the V2G control.

Within the scope of the simulation models available, it has been possible to address the potential for voltage rise caused by V2G to be mitigated through the use of OLTCs, instead of by V2G control, and allowing for the response time of the OLTCs to be slow relative to that of the V2G power electronics. These results are presented in Section 2.

The analysis also considers whether changes to fixed tap settings could resolve V2G voltage constraints.

Partial compliance with standards

The analysis in previous work packages has assumed that any techniques for voltage-based control of V2G inverters would be implemented with full compliance, and that the same control method would be employed for every V2G connection.

However, it is possible that some V2G inverters would not follow the standardised method. There is an incentive to ignore the control signals as this could lead to greater revenue for the customer, despite causing problems for the network operator.

The simulation models in Section 3 consider a case where there is partial compliance with the volt-watt control standard. It is also recognised that the smart meter data used previously to characterise the voltage profiles due to existing demand and generation, also include the voltage rise effects of solar PV. This solar PV generation also represents a partial compliance with the volt-watt control standard, as there is no requirement (in the UK) to implement any form of volt-watt control on solar PV exports.

Control algorithm instability

The V2G control techniques assume that individual inverters would modulate their active or reactive power based on a measurement of voltage. The control actions are also intended to affect this voltage, either by consuming reactive power or by reducing active power. Clearly this raises a risk that a control action to reduce voltage rise, for example by reducing active power exports, will then also reduce the requirement for control, allowing exports to increase again, and therefore causing instability. This instability could involve multiple V2G inverters, or just one.

Previous modelling in work package 2 has used a simplified iteration sequence where the V2G inverters initially export at full power and then each take one control action. This has been intended to demonstrate that successful operation is possible, and also to manage the complexity of the simulation models. The analysis described here in Section 4considers just one substation but extends the iteration sequence so that instabilities can occur.

Control method with improved fairness

The simulations in work package 2 have demonstrated that customers connected at the ends of the feeders will have greater likelihood of their exports being constrained than those connected nearer to the substations. Section 5 considers a dynamic control method where each customer would base their control actions on the maximum voltage seen by any device on the LV feeder. Since all devices see the same control signal, and if they have an equal probability of responding, then on average the export constraints are expected to be equal. This could enable a fairer sharing of the export constraints within the feeder, although additional communication between devices would be required to implement this approach.

2. Voltage management at transformers

2.1 Fixed tap settings

Although the settings of taps at distribution transformers are not directly a risk to the operation of the V2G control algorithms, there is a risk that these algorithms may involve a more complicated process than necessary, for example if a simple change to the tap settings would have resolved voltage range concerns.

Clearly, a change to the tap settings will not resolve thermal loading constraints, and could make the problem worse if the nominal voltage is lowered to allow for voltage rise, such that the energy transfers required by customer devices will operate with higher current amplitude.

Simulation results for 82 LV substations have been assessed to determine whether a tap setting could theoretically be found such that the maximum and minimum voltages for all feeders would remain with the required limits. This requires that the voltages for all customers, on all LV feeders, and for all time samples, would be between the upper and lower limits. No adjustment of the tap setting occurs, so voltage must remain within range for the case where there are simultaneous exports from each customer with an EV charger, and also for the case of simultaneous imports.

The voltages with V2G operation are modelled by combining the existing voltages from smart meter data with the voltage deviations due to either the V2G exports or imports.

Of the 82 substations included in the model, 64 substations would have voltages out of permitted limits of 230 +10% or -6%.when V2G operation is combined with the existing smart meter voltage readings.

This is reduced to 57 substations if the voltage limits are widened to be consistent with the European EN50160 standard of 230 ± 10%. The increase in permitted limits, when considered alone, therefore only resolves 11% of the substations with voltages out of range.

Many of the substations where voltages fall outside of the permitted range may have unused headroom for voltage drops, such that the tap setting could be adjusted so the LV busbar voltage is lower, allowing more margin for voltage rise. It is also possible, though less likely, that there is unused margin for voltage rise but insufficient margin for voltage drops.

Based on existing regulations, and neglecting for now the constraint that tap settings have discrete intervals, an improved tap setting could maintain the substation within range if the difference between the maximum and minimum voltage is less than 16% of 230 V. Using these idealised tap settings, only 48 substations would have voltages out of range. With the wider range of 20% of 230 V, only 40 substations would have voltages out of range.

A combination of widened voltage ranges to 230 ±10%, and a modified tap setting, could therefore resolve the voltage constraint at 30% substations where limits would currently be exceeded. These substations would not require any further voltage-based control.

This calculation involves an assumption that the voltage rise and voltage drop would remain the same if the substation busbar voltage were to be varied. This is an approximation since, for example, voltage drops would be increased if the busbar voltage were lowered, since the current would increase to maintain the same power supplied to the EV.

An implementation of V2G dynamic headroom control could therefore simply exclude EV chargers from any voltage-based control requirements if they are connected to one of these substations where no voltage problems occur. This assessment would need to be reviewed periodically, as the number of V2G installations increases, or possibly if the range of substation busbar voltages widens due to other changes at the primary substation or on the HV feeder network.

A more granular approach would allow for voltage-based control to be enabled only on specific feeders, if other feeders would remain within range once a change to the tap settings had been implemented.

2.2 Distribution transformer on-load tap changers

There are proposals with NGED and also within ENWL to introduce on-load tap changers (OLTC) at distribution substations.

This raises two questions in relation to the voltage-based control considered in this project:

- Will the voltage-based control of V2G installations interact negatively with the control of the tap-changer, for example by causing instability.
- Does the OLTC function remove the need for the voltage-based control of V2G installations

It has not been possible to address the first of these questions in this project as there is insufficient information available to specify how the real-time control of the voltage-based control would operate. It is assumed that the V2G control would respond rapidly, but that the tap-changers at the transformer would have a response in the order of minutes. The risk of instability would also depend on where and how the voltage sensing for the OLTC was configured.

The second question can be addressed here using a similar approach as for the fixed tap setting modifications described in the section above. An OLTC could resolve voltage constraints if the maximum and minimum voltages remain within either the current 16% range or a possible future 20% of 230 V. The requirement differs from the case with a fixed tap settings as it is only necessary for the difference between maximum and minimum voltages at any one time to be within this range, rather than the difference between the maximum over all time samples and the minimum over all time samples.

To determine whether an OLTC could theoretically have resolved voltage constraints at a substation, the simulation results have been analysed to find the maximum voltage range when all customers are exporting, and when all customers are importing. The OLTC would need to accommodate both cases so, as a minimum, the range required is the maximum of the ranges for exports and for imports. This approach assumes that if any combination of simultaneous importing and exporting will have a range that is less than these two extremes. If this were to happen, the exports would supply the imports, and the voltage impacts would cancel. Alternatively, if this were coordinated by a common supplier or aggregator, it could be arranged that an import power matching the required export power could simply be avoided.

Hypothetically, it is also possible that all customers on one feeder could be exporting at the same time as all customers on another feeder are exporting. In this case the voltage rise or drop on the individual LV feeders would not cancel and the OLTC range would need to allow for both cases. Although this scenario seems improbable at a large substation, it is quite likely to occur at smaller transformers, which are more heavily represented in the site selection for this project, where there could be just a few customers on each feeder.

A risk may arise if, when the export event has completed, all the V2G chargers then switch to a charging mode to recover the battery state after the export event. This could cause a significant voltage drop if all installations recharge at the same time. This would occur when the OLTC has set an LV side voltage at the lower end of the operating range and so there is a risk that the voltage drop will cause an under-voltage for customers. Based on the standards seen so far, charging mode is not subject to volt-watt control and there would be no mitigation of this from the V2G chargers. If the OLTC response is slower than the electronics in the V2G inverters, one tap setting may therefore need to allow for both exporting and importing on the same feeder.

In summary, the voltage ranges required for OLTC are therefore calculated as the difference between the maximum voltage while all customers are exporting, and the minimum voltage while all customers are importing, calculated over all customers and all feeders.

On this basis, there are 37 substations where the existing 16% range of 230 V would still be exceeded, and 33 substations where a possible wider 20% range would be exceeded. Assuming the wider range applies, voltage constraints are therefore resolved by installing an OLTC for 42% of the 57 substations where the voltage range would otherwise be exceeded. However, if fitment of an OLTC is only considered where modifications to the fixed tap settings does not resolve the voltage constraint, then only 12% of those substations would require an OLTC.

As a caveat to this discussion, it should be noted that the substations used in this project study were selected due to having an increased risk of incurring voltage or thermal constraints, and many of these have pole-mounted transformers. The analysis in this section has assumed that OLTC could be applied equally to pole-mounted transformers as to ground-mounted transformers, but this may not be the case in practice. Even where adding an OTLC to a pole-mounted transformer is technically feasible, it may not be cost-effective at substations where there are a relatively small number of customers.

OLTCs therefore appear to be a valuable technology in resolving voltage constraints, especially if a wider 20% voltage range is permitted, and a combination of fixed tap changes and OLTCs would resolve voltage issues at 48% of the substations where they occur. The remaining 52% of substations have voltage ranges within the LV feeders that remain outside of the permitted ranges. Optimising tap settings and deploying OLTCs would presumably increase compliance with voltage limits, but some form of control would still be needed if the substations are to support the uptake of V2G that has been modelled here.

2.3 Summary

The discussion above is summarised in the diagrams below, in Figure 1 for the existing 216.2 V to 253 V range and in Figure 2 for a possible future 207 V to 253 V range.

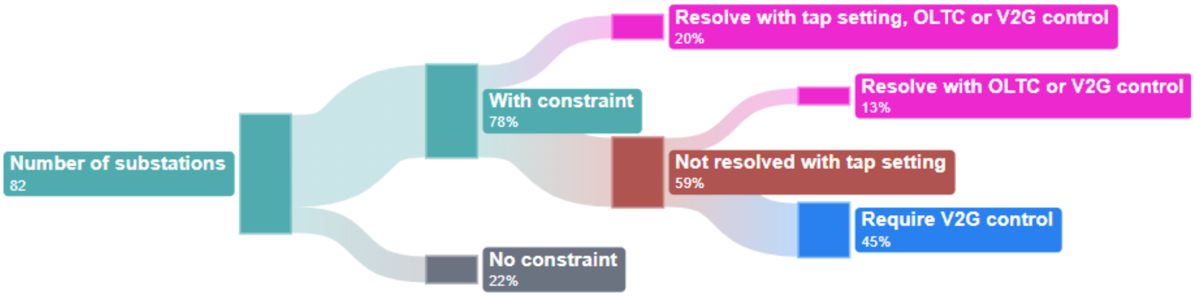

Figure 1: Constraint resolution using transformer options with existing 16% voltage range

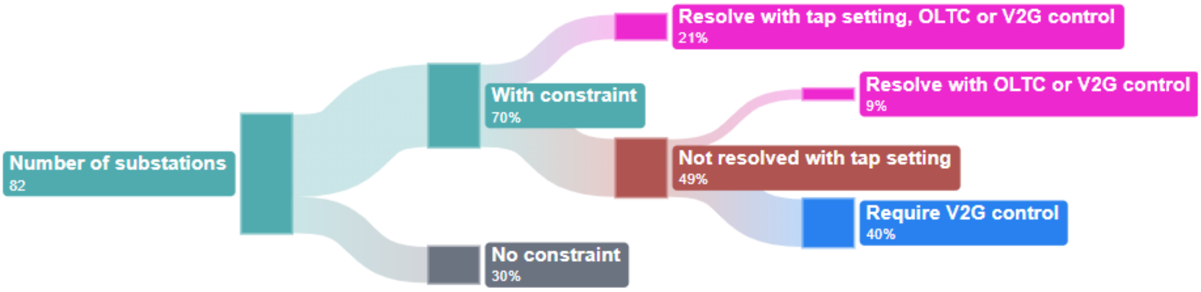

Figure 2: Constraint resolution using transformer options with possible 20% voltage range

3. Partial compliance with standards

There is a risk that not all V2G chargers will follow the desired standards. This reduces the effectiveness of the control action taken by those that implement the standard, increasing the constraint to their active power exports.

Although exports from solar PV have not been added in the control algorithm modelling, there are already customers on the selected feeders with PV and this can occasionally lead to voltage rise in the existing smart meter data.

The results that have already been presented in WP2 of this project, where the impact of V2G operation is superimposed onto the voltage profiles at smart meter customers, therefore already include the effects of exports that are not managed by voltage control.

Rather than modelling further instances where the compliance with voltage control algorithms is reduced, results are presented here for a scenario where the existing smart meter voltages are set to a constant value, removing the effect of any existing voltage rise due to solar PV that has no control standard, and recognising that the effect of this non-compliance was already included in previous results.

Results with the existing voltage rise removed, and with varying volt-watt control threshold ranges, are shown below. The trends here can be compared to Figures 35, 37 and 38 in the WP2 D3 report, where the smart meter data offsets were included. Specific values cannot be compared directly as removal of the smart meter data offsets changes the implied total demand.

In Figure 3, here without the smart meter offset and associated possibility of voltage rise, the probability of over-voltage with volt-watt control reduces to zero for voltages below the upper voltage limit. There is no requirement to set a lower threshold than this to achieve compliance with the voltage constraints. This differs to the results in WP2 D3, where even with lower thresholds, there was still some residual risk that of over-voltages.

With volt-watt-var control, some risk of over-voltage remains even when the active power cut-off threshold is significantly below the upper voltage limit. This possibly arises due to V2G inverters nearer to the substation having a lower voltage due to their reactive power consumption, and therefore not being required to reduce active power exports, causing other connections elsewhere and on different mains or service cable branches to have a voltage higher than the upper limit.

In Figure 4 the probability of over-current is also zero with volt-watt control for thresholds below the upper voltage limit. This may not always be the case as current overloads can occur even when the voltages remain within range, as has been demonstrated in Figure 37 of WP2 D3. As with previous results, there is a risk of overloads with the volt-watt-var control due to additional current amplitudes associated with the reactive power.

The mean export power shown in Figure 5 reduces as the control thresholds are lowered. The lower thresholds increase the probability that active power exports will be constrained, although there is no benefit from lower thresholds once the risk of over-voltages or current overloads has been avoided.

The results suggest that, in the absence of voltage rise due to appliances without active power control, the thresholds can be set around the upper voltage limit and there is no further need for any dynamic variation.

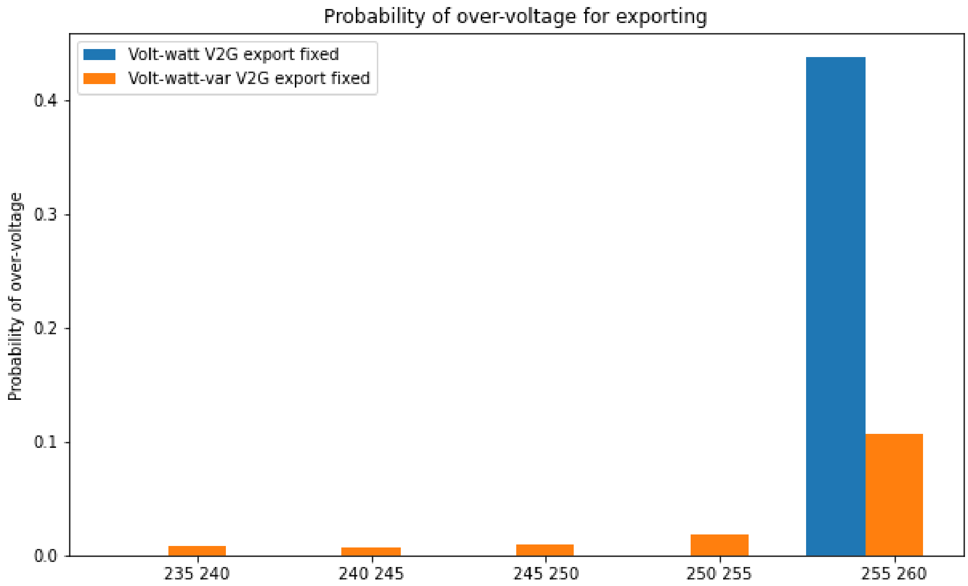

Figure 3: Probability of over-voltage with controlled V2G exports, varying ramp thresholds

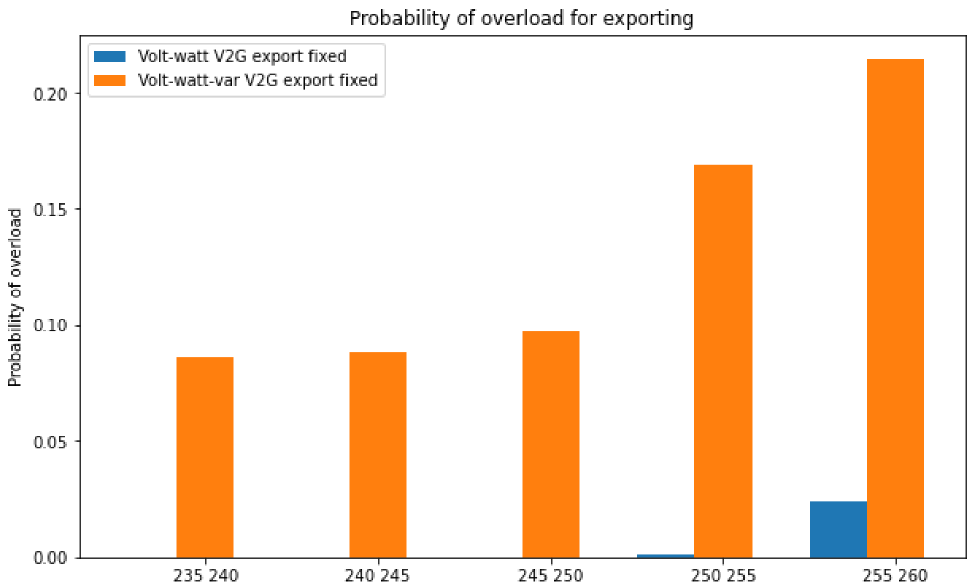

Figure 4: Probability of over-load with controlled V2G exports, varying ramp thresholds

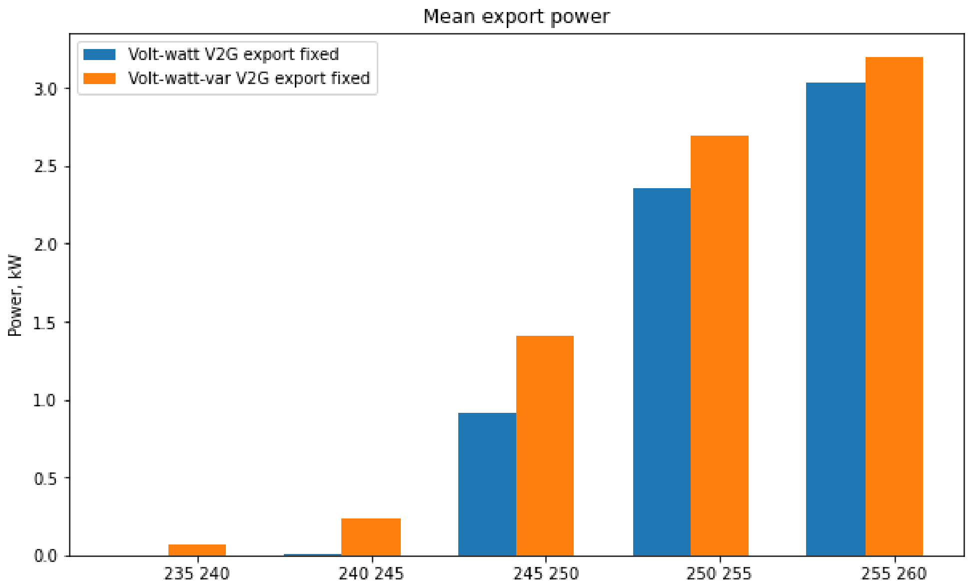

Figure 5: Mean V2G export power, varying ramp thresholds

The effect has also been modelled using an alternative approach where a small proportion of the V2G installations do not follow the control strategy. As in the results presented above, these simulations also do not use voltage offsets from the smart meter data, such that existing voltage rise is not included, but differ as the export power from the un-constrained V2G exports is included in the model. In the plots below, this is denoted as volt-watt control where the offsets are ‘fixed.

The simulation has been configured so that 90% of V2G installations follow the volt-watt control standard, and 10% of V2G installations export at full power regardless of the voltage. These are compared to the case with full compliance (in blue) in the plots below.

Figure 6 shows the probability of over-voltages, which is non-zero even for threshold well below the upper voltage limit for the model with 90% compliance, whereas with full compliance the threshold can be set at the upper voltage limit and need not be set below this. The risk of current overloads in Figure 7 is non-zero for threshold ranges between 250 V and 255 V, for both models, but higher where there is only 90% compliance with the standard. With 90% compliance there are also a few feeders where there is a risk of current overloads even with thresholds in the range 245 V to 250 V.

As would be expected if some customers do not follow the volt-watt control standards, the mean export powers are higher in the results with 90% compliance. The consequence of these increased exports are as described above.

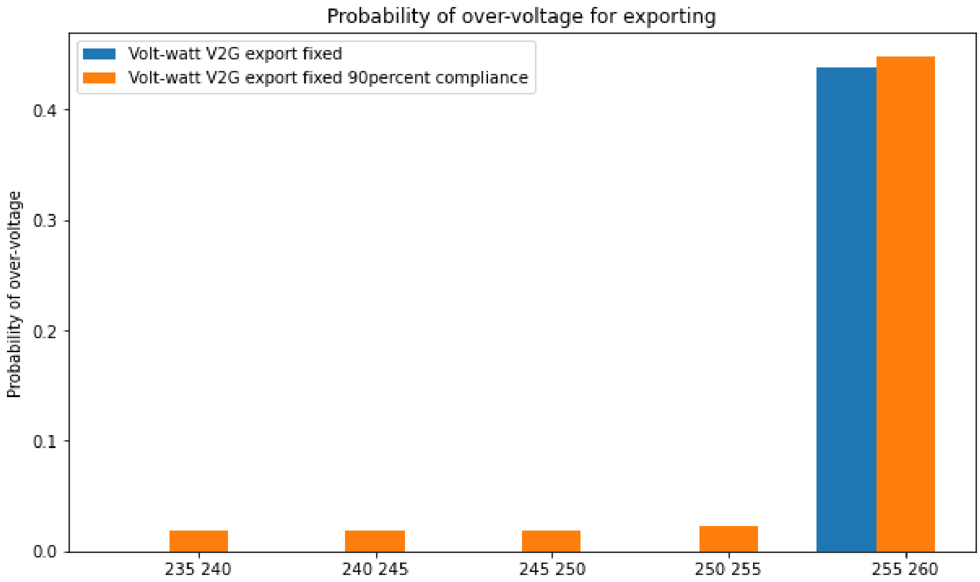

Figure 6: Probability of over-voltage with controlled V2G exports, volt-watt control with 90% compliance

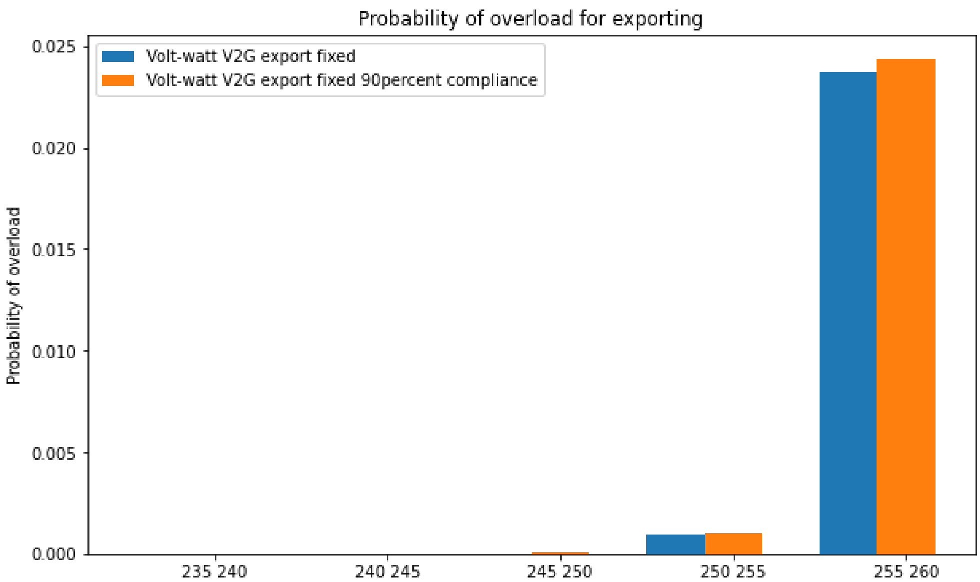

Figure 7: Probability of over-load with controlled V2G exports, volt-watt control with 90% compliance

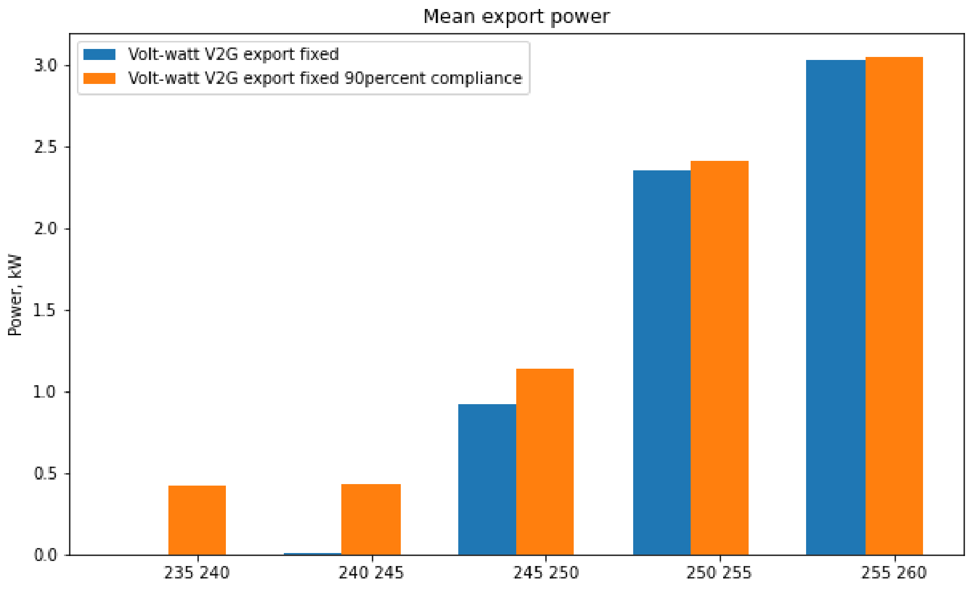

Figure 8: Mean V2G export power, volt-watt control with 90% compliance

4. Control algorithm instability

There is a possible scenario that one V2G charger will take action to mitigate voltage rise, either importing reactive power or reducing exports, thereby reducing voltage for other V2G chargers on the same LV feeder, such that these other V2G chargers then see a lower voltage and increase exports.

This could lead to the first V2G charger then taking further action to mitigate voltage rise and further increases in exports from the other V2G chargers. This mechanism could introduce a high level of unbalance between the V2G chargers with one or more having exports severely impacted and other having no impact.

Conversely, multiple V2G chargers may respond together to a high voltage and then take action to reduce voltage rise. Subsequently, when the voltage is reduced, it is possible to either reduce imports of reactive power or increase exports. This then raises the voltage again, returning to the initial condition. This mechanism could lead to the V2G chargers constantly switching exports on and off, causing a voltage flicker problem, and also failing to limit the peak voltage rise, even though average voltages would be reduced.

Two approaches have been taken to model this effect.

- Extended simulations using power-flow analysis where the volt-watt control is iterated over a longer sequence
- Simplified model to investigate steady state error

These two approaches are described below.

4.1 Extended iteration sequence

Modelling results presented in Work package 2 have used an iterative process for volt-watt and volt-var control that is intended to demonstrate a best possible outcome from the control technique. The simulation method was developed and enhanced throughout work package 2, and in the final deliverable D3, the simulations utilised the full time series available from smart meter data.

The method begins with unconstrained V2G, typically resulting in very high voltages towards the ends of the LV feeders. The export or import powers are then updated in response to the voltage, with each V2G installation being updated in turn. Voltages are re-calculated between updating each V2G installation so that power reductions from one customer will reduce the need for further customers to reduce their power. To avoid the risk of results being unrepresentative due to the V2G installations being updated in a particular order, a different sequence is used for samples from each day in the time-series data.

In this iteration method, the powers from each V2G installation are only updated once. Since the simulation begins with each V2G at full export or import power, when the powers are updated, they can either reduce or remain the same, and there is no opportunity for them to increase. This avoids the risk of instability in the simulation results as powers can only ever reduce.

To investigate longer-term effects, the simulation has been modified so that the iteration process is repeated for 10 cycles, in each case re-using the same sequence of power updates. This means that V2G installations can either increase or decrease power, in response to the voltages seen after the previous update.

Results are shown below in Figure 9 and Figure 10 for one substation 892320 feeder 1, with all V2G installations exporting. There are 92 V2G installations on the feeder, and so the 10 cycles in which the iteration is repeated gives 920 individual updates. The model uses no smart meter offsets and so voltage rise due to any existing solar PV is not included. The volt-watt control thresholds were set such that full active power could be exported below 252 V and zero active power would be exported above 253 V.

Figure 9 shows the maximum voltage, calculated over all the V2G installations on the feeder and over all the time-samples in the data set. Even though no smart meter voltage offsets were included, the simulation has the same duration as if they were. Since there is no underlying voltage rise, the iteration process would be expected to reduce the maximum voltage to 253 V or below, and his is achieved at the end of the first iterations when each V2G installation has updated power once. Subsequently, the maximum voltage is seen to rise again, and settles between 255 V and 256 V.

The mean export power, shown in Figure 10, has an oscillatory trend, although may be stabilising slightly towards the end of the sequence.

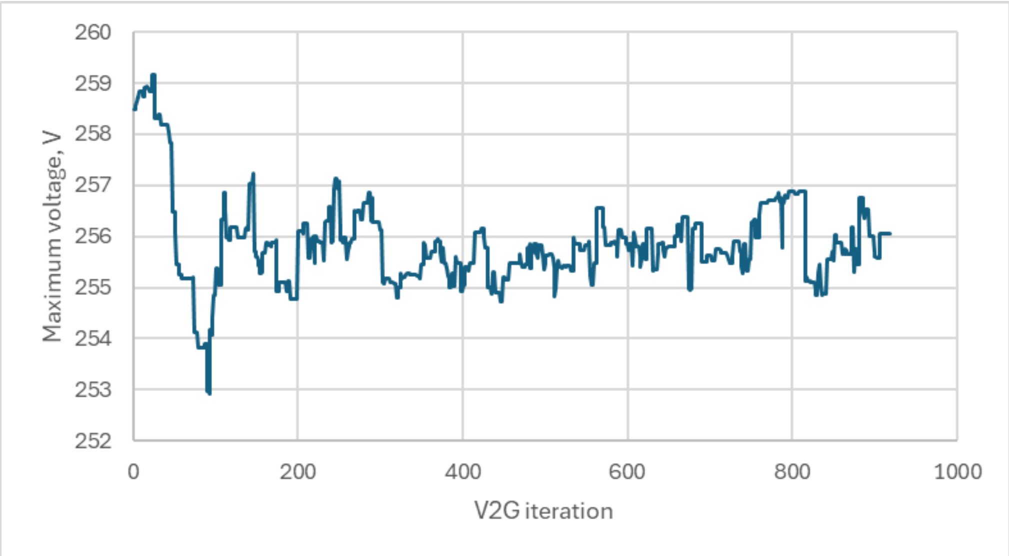

Figure 9: Substation 892320 feeder 1 volt-watt with V2G export, maximum voltage with volt-watt control iteration

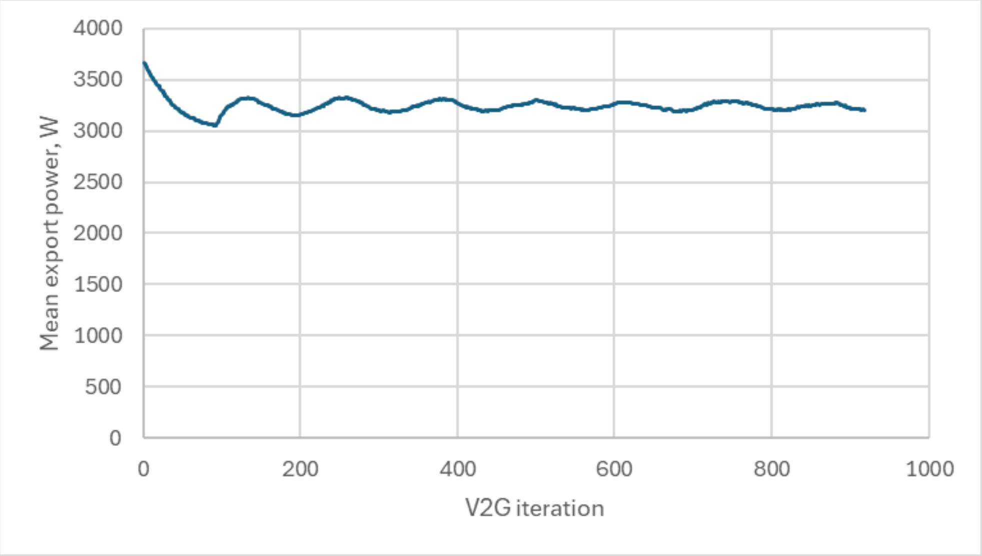

Figure 10: Substation 892320 feeder 1 volt-watt with V2G export, mean export power with volt-watt control iteration

V2G operation with constantly fluctuating export powers is not necessarily a problem from the network perspective, and the impacts on the EV batteries are outside of the scope of this project. However, even if the total exported power from the feeder remains similar, if the locations from which this power are delivered are varying continuously, there is potential for a voltage flicker effect. Variations in the exported power could also cause repeated tap operations if an OLTC was installed at the distribution substation.

The offset between the desired maximum voltage of 253 V and the achieved maximum of around 256 V is also a cause for concern, and would initially suggest that the thresholds would need to be set lower than 253 V so that this maximum voltage might be achieved.

4.2 Steady state error

An additional model was created to investigate this offset between the desired and achieved maximum voltage. This model simplifies the simulation method and the feeder topology so that the control effects can be more easily understood.

A hypothetical LV feeder is considered, with 12 customers, each having a V2G installation, connected in a string topology, as in Figure 11.

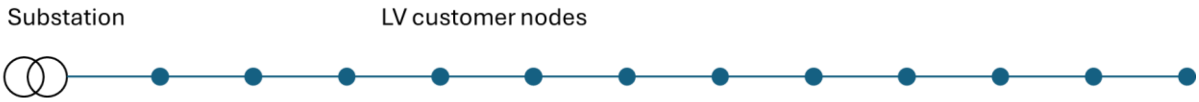

Figure 11: Simplified LV feeder model

In this model the LV customers connect directly to the mains cable, with no service cables. The mains cable branches all have the same length and impedance, defined as such that the current from one V2G export of 3.675 kW at 245 V will produce a voltage rise of 0.25 V. Each customer has V2G exports with a current of 15 A, defined based on a power of 3.675 kW at 245 V. This current varies with volt-watt control, but the model in this simplified model there is no dependency on the voltage otherwise. The substation voltage is set at 245 V.

The model has no reactive power in the V2G exports and no reactance within the cable branches.

If all V2G customers export at full power, there will be a voltage rise of 3 V between the substation and the first customer, progressively reducing to 0.25 V between the penultimate customer and the end of the feeder. This initial state is shown in Figure 12, where the voltage along the feeder is plotted for branches numbered 0 to 11, starting from the substation. This also shows an equal 15 A current from each V2G customer.

Volt-watt control is then applied to the customers in a randomised sequence. The control uses a single threshold of 253 V, with full export current for voltage at or below this, and zero current for voltages above. There is no constraint that all customers must be updated before further returning to a previous customer. In the first iteration, customer number 8 updates power and switches off. The voltage rise along the feeder is correspondingly reduced, as in Figure 13.

If the first three customers are requested to update power, no change will be made as their voltages are below the 253 V threshold, but other customers further along the feeder will switch off when updated. This process continues for a number of iterations, at one point giving the profile shown in Figure 14. The maximum voltage on the feeder is now reduced to 255 V, but customers 9 and 11 have not yet been updated and remain on full export current. If updated, they would switch off as their voltage is above the threshold. However, customers 0 to 5 now have voltages below 253 V, and two of these are switched off. If these two customers update their power they will switch back on.

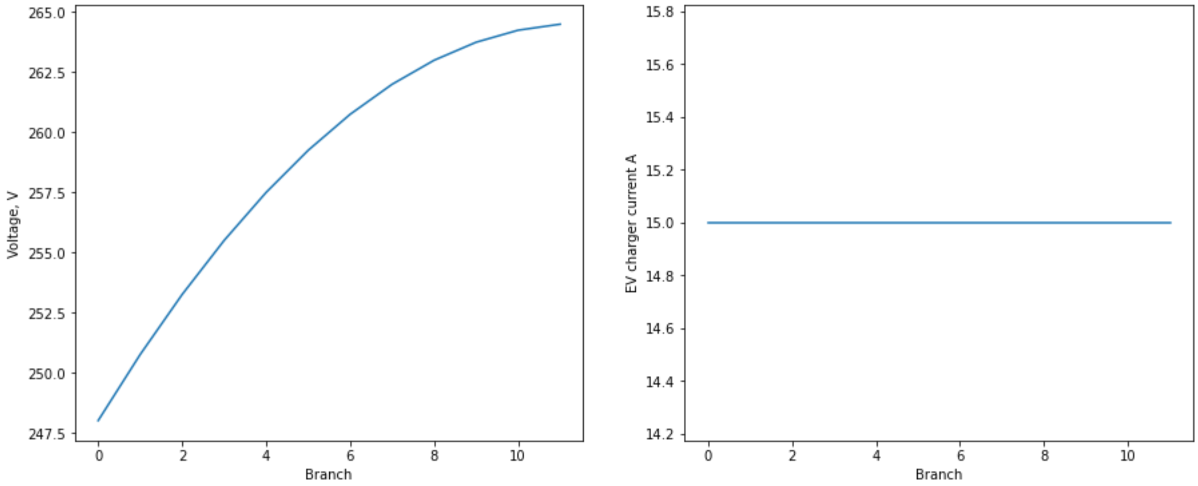

Figure 12: Initial voltage and current state with V2G unconstrained

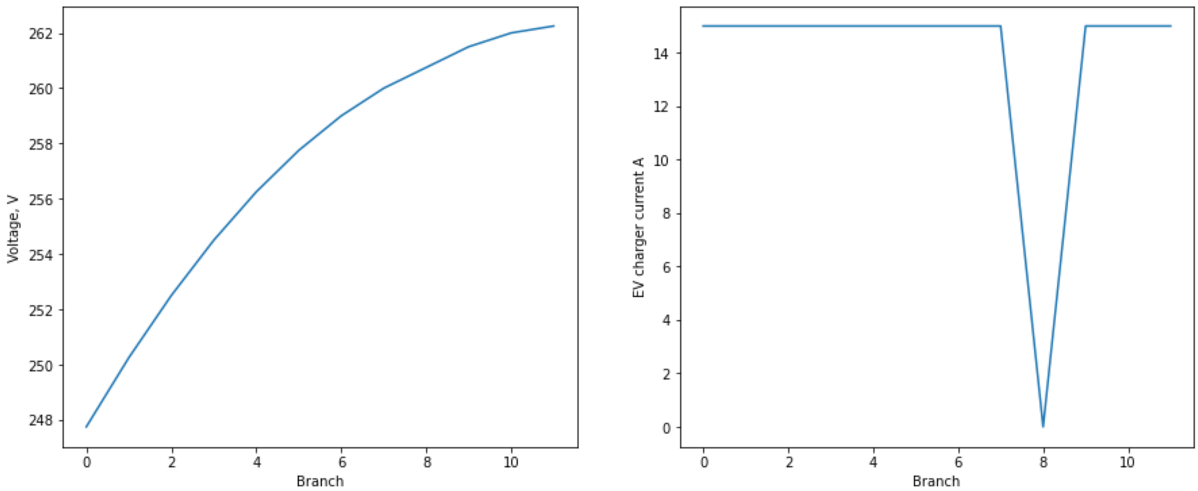

Figure 13: Voltages and currents after one V2G power update

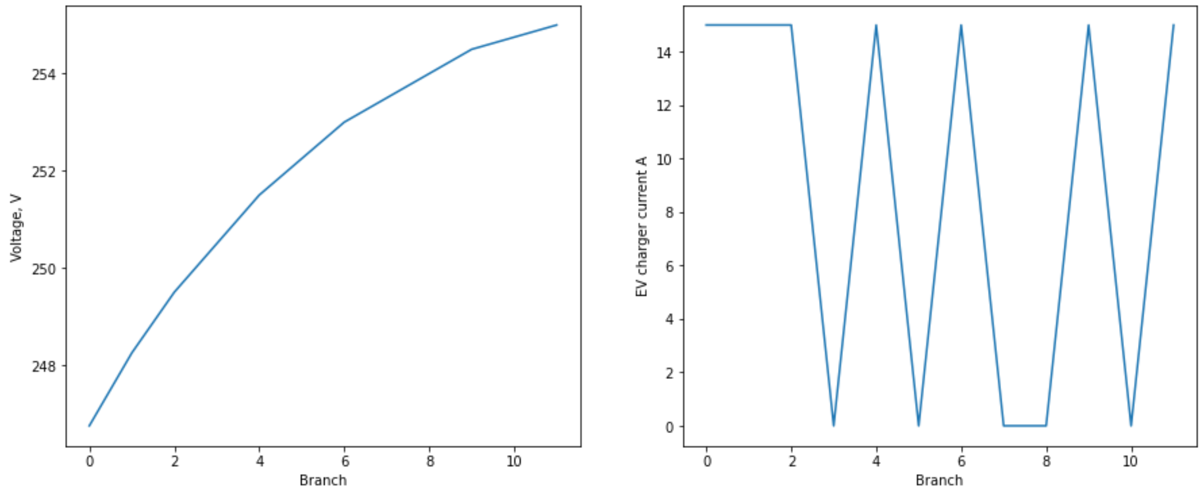

Figure 14: Voltages and currents after a further V2G power updates

This process has been continued for 200 iterations, and the maximum voltage plotted after each iteration, as in Figure 15. As in the results from the section above, there is an offset between the desired threshold of 253 V and the achieved maximum voltage, in this case around 254 V. In this 'steady state', if the maximum voltage at the feeder end is 254 V, then the majority of the customers will have a voltage below 253 V, and will switch on when their power is updated. Only a few customers at the end of the feeder will have voltage above 253 V, and these will turn off when power is updated. Even though there are fewer customers likely to switch off than those likely to switch on, the impact on voltage of the customers at the end of the feeder is much greater as their current travels through a greater number of branches.

This effect is well known in control theory and is referred to as a steady state error. The error arises since there is a possibility that control updates will cause the voltage to exceed the threshold, and since the counteracting control updates to reduce the voltage can only take place when the threshold is exceeded.

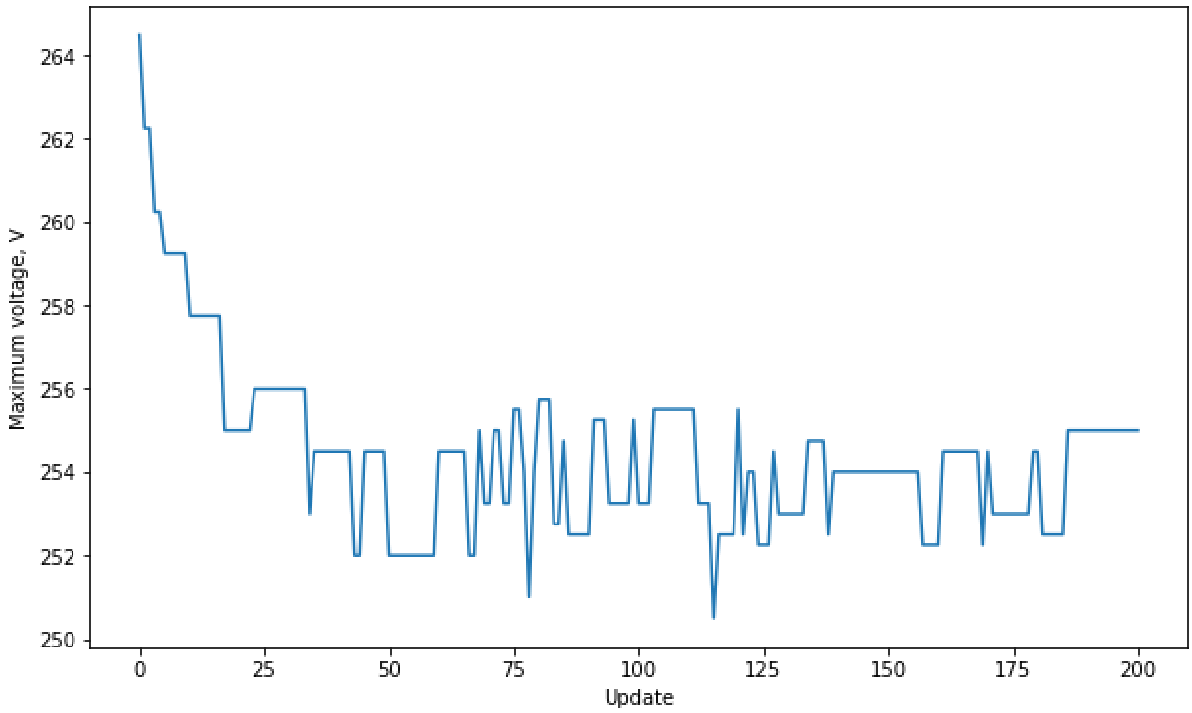

Figure 15: Voltages and currents after 200 V2G power updates

There are also well-established methods to reduce this steady state error, although it cannot be eliminated completely other than by introducing a compensating bias into the system, in effect by setting a lower threshold. To demonstrate one possible method, a modified volt-watt control method has been tested in which, as before, the export current is switched off when the voltage is above 253 V, but only switched on if the voltage has remained at or below 253 V for 10 successive updates. Results in Figure 16 demonstrate that this method can lower the maximum voltage.

Although the maximum voltage in Figure 16 is below 253 V at the end of the iteration sequence, there will be periods within a much longer iteration when the voltage increases above the threshold. This particular modification introduces a probabilistic integration effect where the likelihood of voltages being above the threshold is reduced, but the extent of the deviations above this remains the same as in Figure 15.

The aim here is not to design an improved control strategy, as any realistic implementation would need to take account of the sensor response times, frequency of updates, and any averaging or hysteresis in the real-time control. However, the results demonstrate that methods to achieve a maximum voltage that is close to the desired threshold can theoretically exist. These could be investigated further in a trial using real V2G equipment.

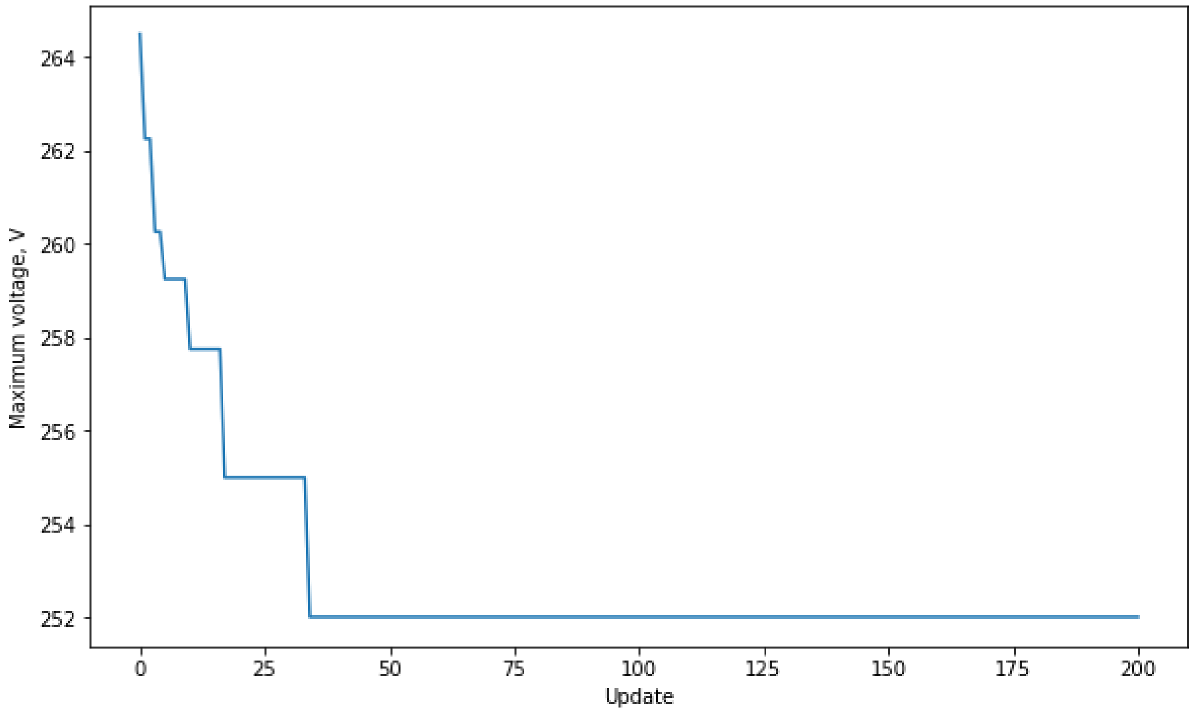

Figure 16: Voltages and currents after 200 V2G power updates with modified control algorithm

5. Control method with improved fairness

The results in work package 2 highlighted the increased likelihood of customers located near the ends of LV feeders having their export power constrained, as these customers see the greatest effects of voltage rise. This is caused by their own exports, in combination with voltage rise due to exports from all the other customers nearer to the substation.

As noted in work package 2 deliverable 3, customers at the feeder ends are constrained more, and customers near to the substation are constrained less. Figure 17 repeats these results from work package 2 deliverable 3, showing a cumulative distribution of the mean export power for V2G customers on feeders with 20 or more customers. This limit was included as the power constraints on feeders with fewer customers are not statistically homogenous with those on longer feeders.

In work package 2 deliverable 2 a method was tested to set individual volt-watt control thresholds for each customer, with the objective of improving the fairness of export constraints. This method would set lower thresholds for customers nearer to the substation so that their exports would reduce, thereby lowering the voltage rise seen by customers at the feeder ends, and enabling the customers at the feeder ends to have less constrained exports. The results presented in WP2 D2 were not successful at improving the fairness and would be complex to implement.

A simpler approach, though also with complicated implementation, would be for each V2G customer on a feeder to make volt-watt control responses to the maximum voltages at any customer connection. Since there is only one maximum value, if the probability of responding to this is the same for all customers, then the expected power constraint would also be equal for all customers.

This approach has been modelled using volt-watt control where every customer is assumed to have real-time knowledge of the maximum voltage seen at any of the customers on the feeder. The control actions follow a randomised sequence, as before, so that on each day in the smart meter data time series, customers have power updates in a different order.

This method achieves the same compliance with voltage and thermal capacity limits as before when customers only could respond to their own voltage measurements. However, the probability distributions of mean power exports per customer are the same regardless of whether customers are at near to the substation or at the feeder ends, as shown in Figure 18.

Although the traces in Figure 18 for customers at the feeder ends and near to the substation do not overlap entirely, they are much closer than in Figure 17 and the remaining differences are likely to the random sampling and the reduced number of customers at feeder ends compared to the number elsewhere on the feeders.

Figure 18 also shows that the mean exports where the control response is based on the maximum customer voltage is less than in Figure 17 where the control is based on the voltage at each customer connection. In effect, there is a disadvantage associated with introducing fairness as, on average, the mean exported power is reduced. Exports are lower with the fairer approach as the constraints are shared equally between customers, rather than being mostly at the feeder ends. More power is exported from the feeder ends and less from customers nearer to the substation. For the same total exported power, voltage rise will therefore be higher. In other words, for the same voltage rise limits, the mean exported power will be less.

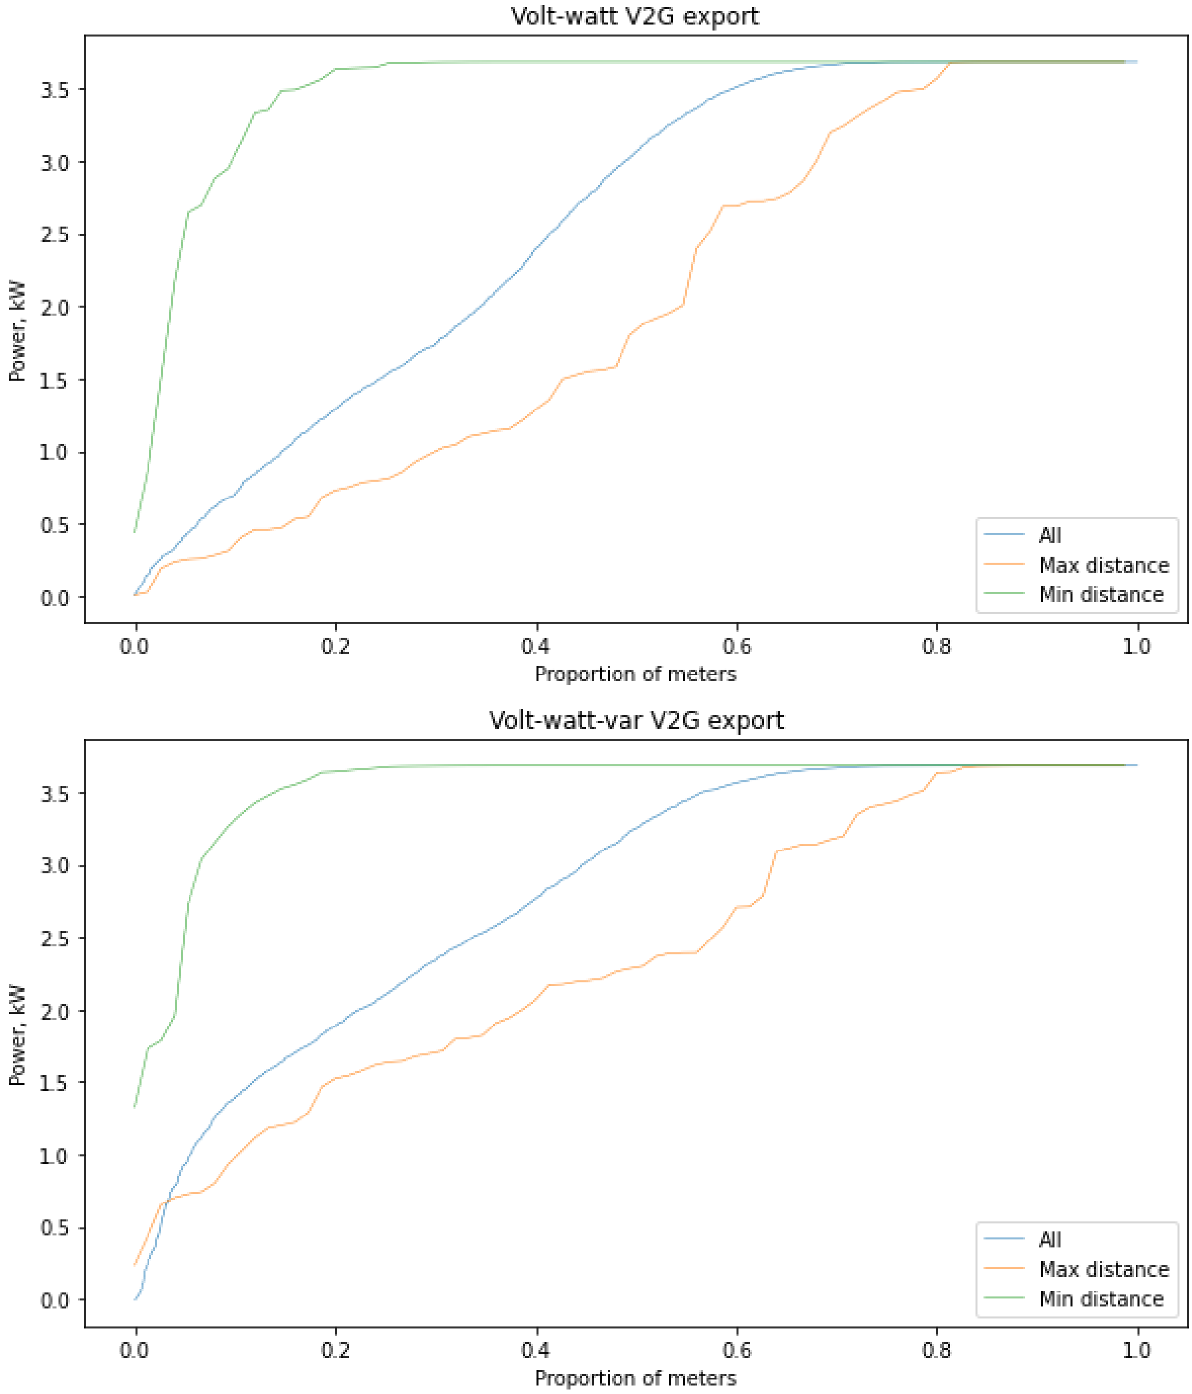

Figure 17 Mean V2G export power, minimum and maximum distances, volt-watt ramp 252 V to 253 V, with control using voltages measured at the customer connection

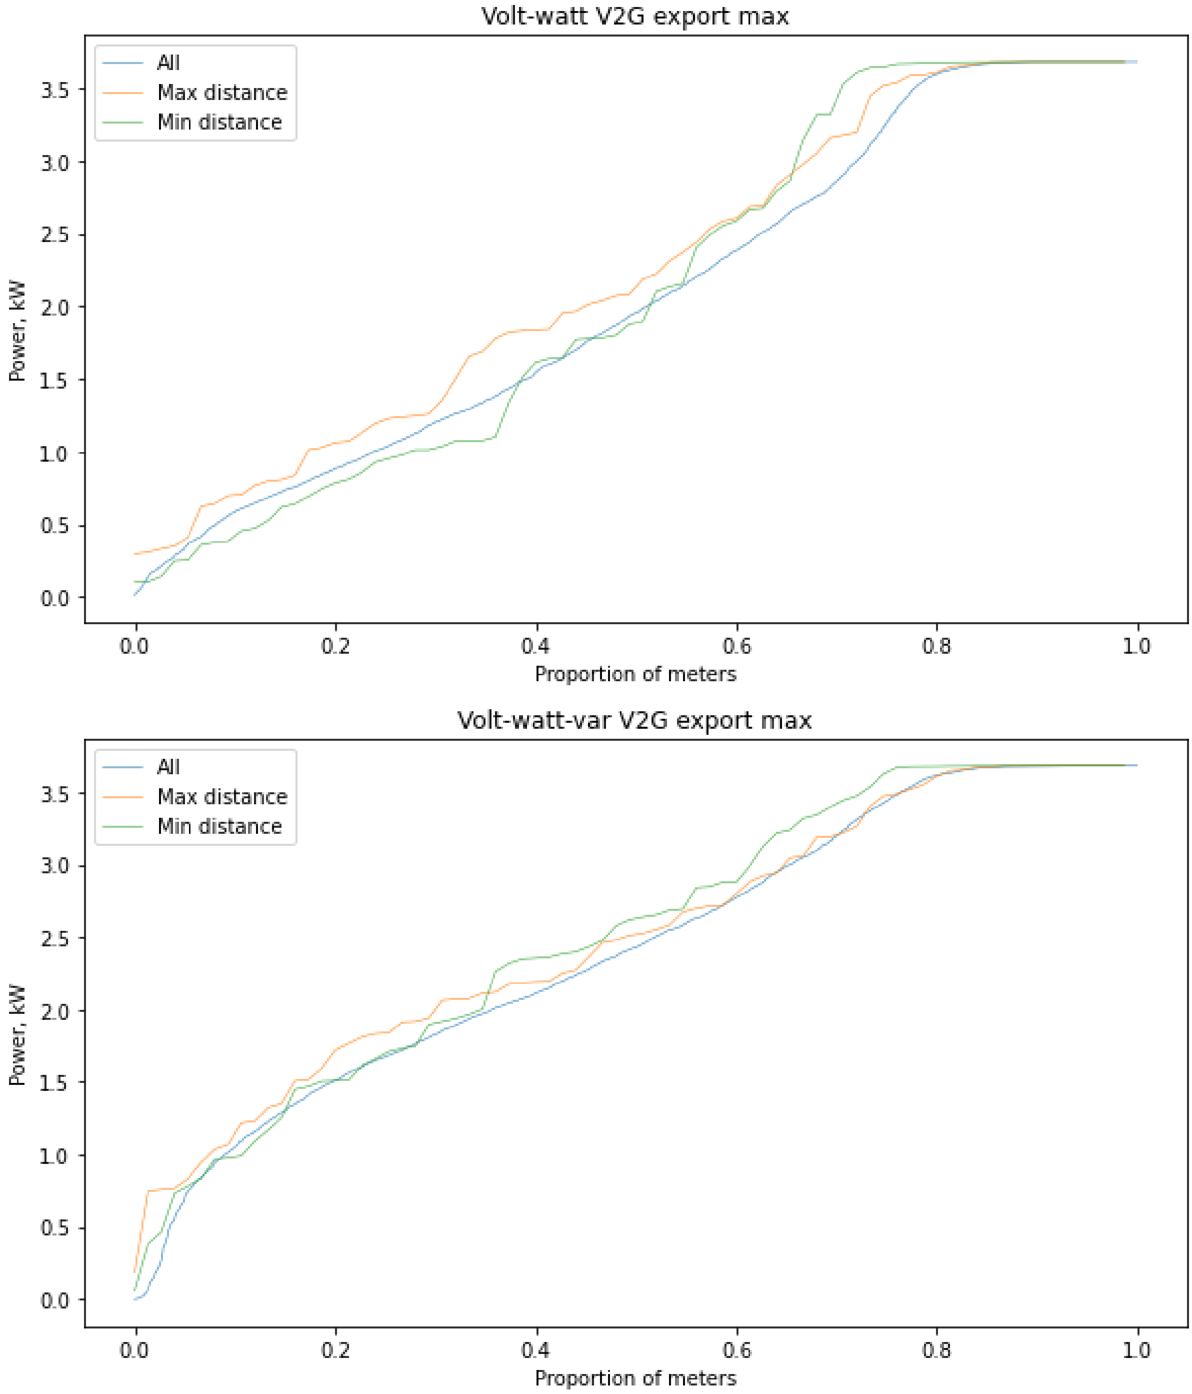

Figure 18 Mean V2G export power, minimum and maximum distances, volt-watt ramp 252 V to 253 V, with control using maximum voltages for any customer on the same LV feeder

6. Conclusions

This work package has included additional modelling to investigate possible risks with the implementation of V2G control. The simulations presented here extend the previous modelling work in work package 2, for which the most comprehensive results were presented in deliverable WP2 D3.

The modelling of volt-watt control has addressed a question arising from previous modelling that showed that thresholds to constrain active power exports, for some feeders, needed to be set much lower than the upper voltage limit. In some cases, this is required as the voltage-based control does not address thermal overloading constraints, but the simulations demonstrated that lower thresholds were needed even to avoid the voltage thresholds being exceeded.

Modelling in this work package has shown that this effect relates to occasional periods when the smart meter demand already contains voltage rise due to existing solar PV. If this existing voltage rise is removed by removing the voltage offsets derived from the existing smart meter data, then setting the volt-watt control thresholds at 253 V will result in full compliance with the voltage limits. Further simulations, also with the existing smart meter data removed, and where 10% of the V2G devices did not follow the control standard, indicated a similar effect as with the voltage rise due to the existing solar PV.

This suggests that, if voltage limits were the only concern, then no dynamic variation of the thresholds would be needed and a common factory-configured threshold for all V2G devices would be sufficient. However, there are some feeders where thermal limits also cause a constraint. Voltage rise is not necessarily an ideal signal to ensure that feeders remain within a maximum current limit, but this can be achieved by setting a lower voltage threshold, above which active power exports would be constrained. A lower limit is also required where there are exports that do not respond to the same volt-watt control as required for V2G. In effect, if some devices do not follow the control standard, then those that do follow the standard need to be constrained at a lower threshold.

This is a useful conclusion from the perspective of dynamic V2G as it only becomes necessary to deviate from the default threshold settings for a relatively small number of feeders where exports would be limited by thermal constraints rather than voltage rise limits. Lower limits would also be needed where there is sufficient existing generation to cause voltage rise along the feeders. Initially this is likely to be widespread, but if the same control standards were applied to all new devices capable of exporting, then the need for this would diminish over time.

There may also be a need for lower voltage thresholds to be applied on feeders where there is significant unbalance, so that high demands on one phase could cause voltage rise on other phases. Any V2G exports on phases with low demand could add to this voltage rise effect and could cause customers located towards the feeder ends to have voltage above the upper limit.

The modelling here also considered whether the impacts of voltage rise due to V2G exports could be resolved by adopting wider voltage limits of 230 V ±10% or with modifications at the transformer, either changes to fixed tap settings or by installing on-load tap changers (OLTC).

Clearly, the wider voltage limits are helpful in cases. Assuming these limits apply, 70% of substations have a voltage constraint in the worst-case scenario modelled here where approximately 80% of customers have simultaneous V2G exports of 3.7 kW. All of these constraints could be resolved by volt-watt control of the V2G, but for 30% of substations it would also be possible to resolve these constraints using OLTCs, and for 21% a change to the fixed tap setting could be sufficient. This modelling has assumed idealised tap ranges, and that perfect foresight could be available to select an optimum setting. In practice the discrete tap steps and uncertainty over future voltage deviations would limit the applicability of this approach. However, even where volt-watt control is employed to ensure compliance with voltage limits, a combination of V2G control and OLTC or tap changes could allow the export power constraints to be reduced, with greater benefit to national system balancing.

The simulations have also started to address concerns over instability where multiple V2G devices could take conflicting control actions, such that reductions to voltage rise when one inverter reduces exports could then allow another inverter to export a higher power. Previous simulations m work package 2 have intentionally avoided this effect to demonstrate that a successful control effect is possible, and to allow the modelling to represent many substations and feeders with a reasonable computation time.

Simulations for a selected substation have demonstrated that continued control iterations of the export power, based on voltage measured by each V2G inverter, can cause ongoing variations in the maximum voltage and the aggregated export power. This could arise in practice where the control software makes continued adjustments to the export power, with each inverter operating asynchronously to the others, and potentially with short update periods of a few seconds. The resulting variations in feeder voltage may cause undesirable voltage flicker effects.

The modelling also demonstrates that the control action can cause a steady-state error, such that the maximum voltage, on average, remains several volts above the volt-watt control threshold. Achieving a maximum voltage no higher than the upper voltage limit would then need the volt-watt control threshold to be set lower than this. When the maximum voltage at the feeder ends is at the upper voltage limit, inverters nearer to the substation will see a lower voltage than this and may increase their export power, causing the voltage at the feeder ends to rise above the limit. The voltage rise can be mitigated by other inverters reducing their export power, but this will only occur when signalled by a voltage that is above the threshold. In an asynchronous system where all inverters are continuously adjusting voltages, any reductions in voltage will be offset by simultaneous increases elsewhere.

These are well-known characteristics of control systems, and the impact can be moderated using integration or hysteresis effects. The work here has not aimed to design a real-time control system for V2G but simple simulations have been included where export power is immediately reduced for a voltage above the threshold, but where export power is only increased when the voltage has remained below the threshold for at least 10 update cycles. This is not proposed as an optimal solution but serves to demonstrate that the steady-state error and the voltage variations can be significantly reduced.

The effects of control instability are not readily simulated with a simulation approach based on sequential time-steps. These asynchronous effects would be better addressed with representative V2G devices in a real-world trial.

Finally, a dynamic mode of operation has been modelled to demonstrate that a more even distribution of export constraints could hypothetically be achieved, thereby avoiding the unfair characteristic of volt-watt control where customers at the feeder ends are most likely to have export power constrained. This effect has been demonstrated with volt-watt control in work package 2 deliverable 3. An alternative implementation has been modelled here where the control action responds to the maximum voltage seen by any customer on the feeder, rather than the voltage at each individual connection. If all devices respond to the same maximum voltage, and have an equal probability of responding to this voltage, then it would be expected that the mean export power would be the same for all customers. This could be considered a fairer approach as all customers on one feeder would have the same average constraints.

A practical method would require real-time communication so that all V2G devices on the feeder, and on the same phase, would share voltage measurements with each other. The voltage readings could be collected and re-shared by an energy supplier or aggregator, using existing internet connections to the devices for communications. This would represent a much higher bandwidth use of the communications system than previously discussed, where the operation of voltage control on V2G devices may need only occasional updates to the threshold parameters with the real-time operation otherwise being autonomous. A method of sharing voltage readings between devices also requires a bi-directional communication, so either all devices would need to communicate with the same aggregator, or there would need to be additional data sharing between aggregators. Possibly the

DNO may coordinate the voltage reading communication. The control communications implied by this method is therefore significantly more complex, although not infeasible.

A consequence of the improved fairness of export constraints within the LV feeder is that the overall mean export power is reduced, relative to the ‘unfair’ approach where the volt-watt control responds to the voltage at each connection. This arises because power exported at the feeder ends causes greater voltage rise than power exported nearer to the substation, since the associated currents pass through the higher impedance of a longer cable length. The unfair approach tends to constrain exports at the feeder ends where their impact would be worse. In the fairer method where constraints are equalised between customers, the proportion of exports at the feeder ends is increased, and so exports must be reduced overall to maintain the same maximum voltage rise.

It could be argued that a method to equalise mean export powers within an LV feeder only partially addresses the issue of fairness. The inherent differences between LV feeders and between substations would remain.

While the simulations in this work have demonstrated a significant difference in the mean export power between customers at feeder ends and those nearer to the substation, the modelling here is only representing a worst-case scenario that all V2G customers export simultaneously. While this worst-case scenario is important from the perspective of the network operator, in a more typical scenario only some EVs would be connected and so peak export powers would be less. Although the control would still have the same unfairness towards customers at feeder ends, the loss of export power would be less.