V2G Dynamic Headroom Control

Extended Simulations with Smart Meter Data

Version 1.0

20th January 2026

Prepared by:

Andrew Urquhart and Murray Thomson

CREST, Loughborough University

Prepared for:

Liza Troshka

National Grid Electricity Distribution

Contents

1. Introduction

Project background

This report describes simulation modelling to investigate the operation of vehicle-to-grid (V2G) power exports from electric vehicles (EVs) and their impact on LV feeder capacity.

V2G power exports are expected to cause voltage rise on LV feeders such that the customer voltages may exceed the permitted maximum and currents may exceed the rated cable capacities. To mitigate this risk, it is suggested that the V2G exports could be controlled by either volt-watt control, where exports are reduced for higher voltage; volt-var control, where reactive power is imported when voltages are high and exported when voltages are low; or a combination of both these techniques. These methods respond to voltage and so do not necessarily ensure that the feeder currents remain within limits, and their operation on representative UK LV feeders has not yet been assessed.

This project aims to investigate the use of these controls on typical UK LV feeders and also to consider whether a more targeted and granular customisation of the voltage threshold settings would help to minimise constraints on the exported power while ensuring that LV feeders remain within capacity limits. This 'dynamic headroom control' approach assumes that smart meter data would be used to characterise the voltage rise and thermal loading due to the existing demand.

Purpose of this report

The work presented here builds on previous simulations from Work Package 1 deliverable WP1 D3 and earlier deliverables in Work Package 2.

In Work Package 1, initial simulations modelled with V2G power exports on a set of selected LV feeders. These simulations used representative network and impedance data for the LV feeders but the demand data was synthesized and so did not account for real-world variations in demand or allow for the impact of variations in the HV feeder voltage and corresponding changes to the LV substation busbar voltage.

Further simulations in Work Package 2 D2 then incorporated real world smart meter data into the modelling, taking into account the voltage variations seen on LV feeders due to the existing demand and also tap changes from primary substations that determine the voltage at distribution substation busbars.

This deliverable Work Package 2 D3 extends the results from Work Package 2 D2, making use of smart meter data captured over a longer time period, specifically to include variations over the summer season, where voltage rise might be expected to be higher due to solar PV exports and heating demand would be lower.

The simulation method adopted here has remained similar to that in Work Package 2 D2 and so this report has been prepared by updating the previous deliverable, rather than by providing an addendum, such that the results are fully described in on up-to-date document.

Simulation approach

The initial simulation models using synthesized demands have been revised to incorporate smart meter data as a representation of the existing demand. Although the existing demand for each customer cannot be obtained from smart meter data so that customer privacy is maintained, the impact of these demands in terms of customer voltages can be determined from smart meter voltage data. The smart meter demand data aggregated over all the customers on an LV feeder can be used to estimate the current in the first mains branches from the substation, typically the cable branches that are closest to their capacity limits.

The smart meter data can also be used to characterise the existing demand and determine the headroom available for the addition of V2G power exports. At many times of the day, small scale power exports would be beneficial and offset the power imports of the existing demand. However, there is potential for exports to occur, coordinated between all the V2G customers on an LV feeder, and at times when the voltages are already at their maximum. Section 2 considers this existing headroom and the selection of a specific set of worst-case conditions that could combine with the addition of V2G exports. There is also a risk that the associated recharging will also occur when voltages are at their minimum.

The simulations summarised in this work package may be the first within NGED to combine measured smart meter data with hypothetical future demand models. They also demonstrate a new approach to power-flow analysis that overcomes the limitation that the demand is not known for every customer, by considering the new V2G demand in super-position with the measured voltages and aggregated demand to represent the impacts of the existing demand. Section 3 of this report describes this new simulation method.

The report presents results from these simulations, initially with no control methods in section 4, and then with various control approaches in section 5. These controls are in some cases ‘static’ where the control thresholds are constant for all V2G installations, as if they were factory-programmed, or ‘dynamic’ where they are adjusted for each substation, feeder or customer. This mode of operation implies a background control process where revised threshold settings are determined by the DNO, and also requires a communication process in which the thresholds are updated periodically for each V2G customer.

2. Existing headroom

2.1 Voltage ranges

The existing maximum and minimum ranges, as measured by the smart meter voltage data, have been assessed to determine the headroom available with the existing demand and uptake of low carbon technologies.

Voltage data has been obtained for the smart meters at each of the selected substations defined in the WP1 D1 Site Selection Report [1]. The voltage data has 30-minute resolution and has been used previously to determine the phases of single-phase connections [2].

An analysis method has been developed using Python software to determine the margins available at both the upper and lower end of the permitted voltage ranges. This analysis uses the following steps to import and process the smart meter voltage data:

1. Data from each smart meter is imported from an SQL database containing voltages for each smart meter device id.

2. Zero values and any spurious high values are removed. From previous experience, errors in the smart meter data processing typically result in very high values over 1e9 so this threshold has been used here. This leaves voltages that are outside of the permitted ranges (over 253 V or under 217 V) are therefore otherwise included.

3. The smart meter samples are captured at approximately 30-minute intervals but there appears to be a clock drift such that the 30-minute periods do not always begin at 0 and 30 minutes after the hour. The offsets differ for each smart meter. A common time-frame is needed to compare the meter voltages against each other and so readings are interpolated to a 1-minute basis, and then re-sampled to 30-minute periods starting at 0 and 30 minutes past the hour.

4. The voltage data is stored in a tabular format for each substation and LV feeder, where the table rows correspond to successive half-hour periods and the columns hold the data for each smart meter connection. Three columns are created for the individual phase readings from three-phase smart meters.

5. These voltage data tables are populated where readings exist for each smart meter, extending from 1st July 2024 to 31st August 2025. This length of data history is only available for a few meters, in some cases because the meters have only been installed more recently. However, there are also many gaps in the data where readings have been unavailable.

Figure 1 shows the maximum voltages on each LV feeder, where the maximum is calculated over each of the connected smart meters and separately for each half-hour period. The maximum voltages for each half-hour period are presented as a cumulative probability density function (CDF). Dashed line shave also been added to show the upper range limit of 253 V and the lower range limit of 217 V.

For example, taking the line for the feeder with highest voltages in cyan blue, the maximum voltage over all the smart meters is below 253 V for around 30% of time samples, and therefore out of range for 70% of time samples. Several feeders have maximum voltages that are above 253 V for a significant proportion of time.

These very high voltages could potentially only occur at a very small number of meters on the LV feeder. This is examined in Figure 2 where the maximum voltage is calculated over all the available half-hour time samples and plotted as a CDF relative to the proportion of meters included. The lines on this plot are straight, indicating that there are many feeders where all the meters record voltage that are over the threshold of 253 V, even though this may occur for only a small proportion of time. Figure 1 shows this effect as the upward pointing tail on the right of the plot where many feeders have

voltages over 253 V for a very small proportion of time (or, in other words, the maximum voltage is below 253 V for nearly 100% of time samples).

It seems likely that these infrequent high voltages, experienced by all the meters on the feeder, are due to occasional high voltages at the substation busbar.

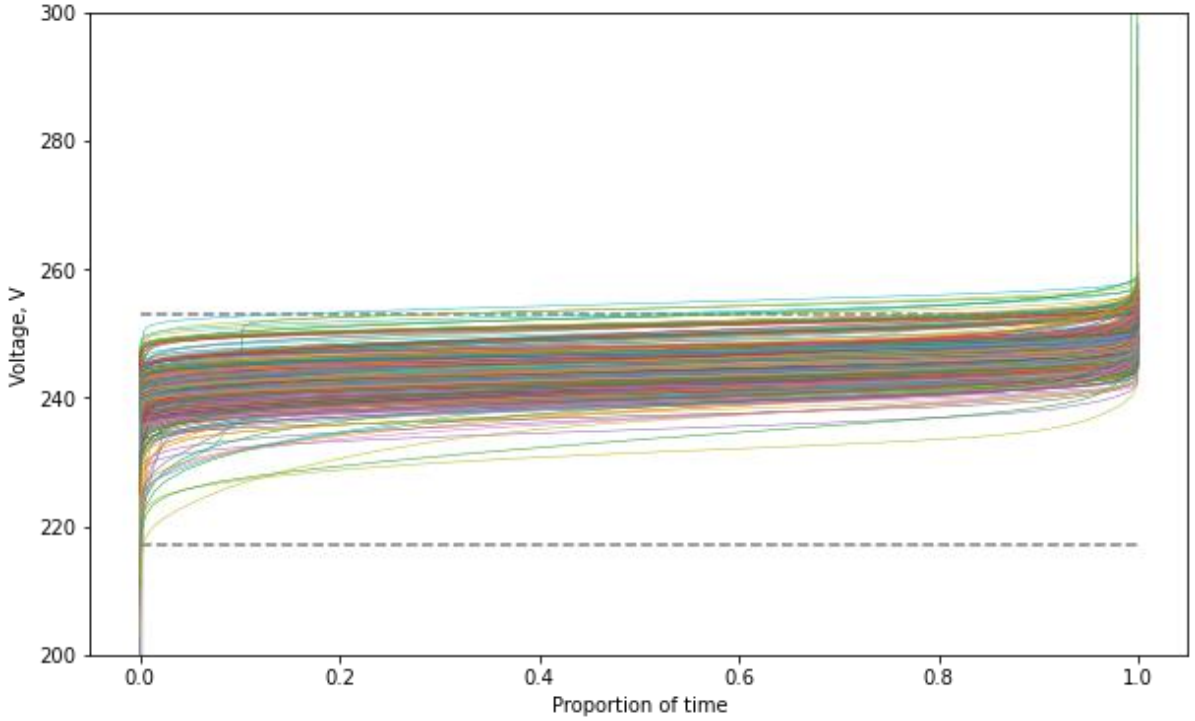

Figure 1: Maximum voltage over all meters for each LV feeder

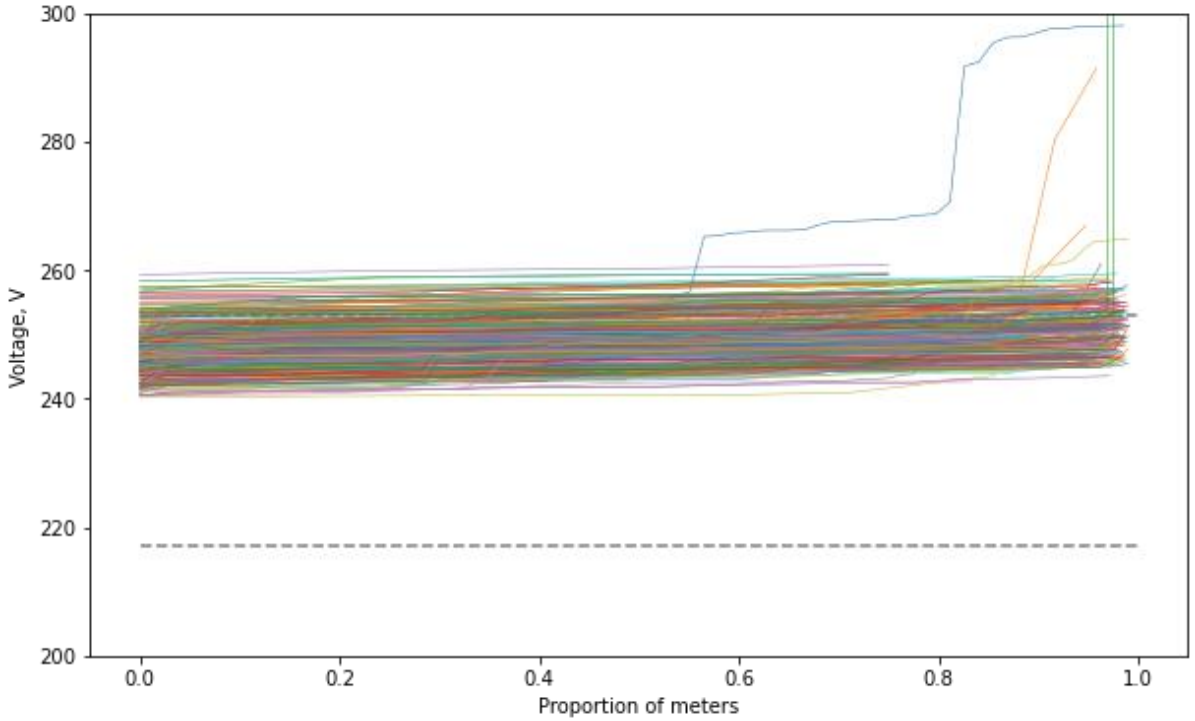

Figure 2: Maximum voltage over all time samples for each LV feeder

The probability of feeders being above the maximum voltage threshold is summarised in Figure 3. This plot shows the proportion of feeders for which the maximum voltage is above the limit range for the given proportion of time. For example, 1% of feeders are out of range for 50% of time samples, and around 4% of feeders are out of range for 10% of the time.

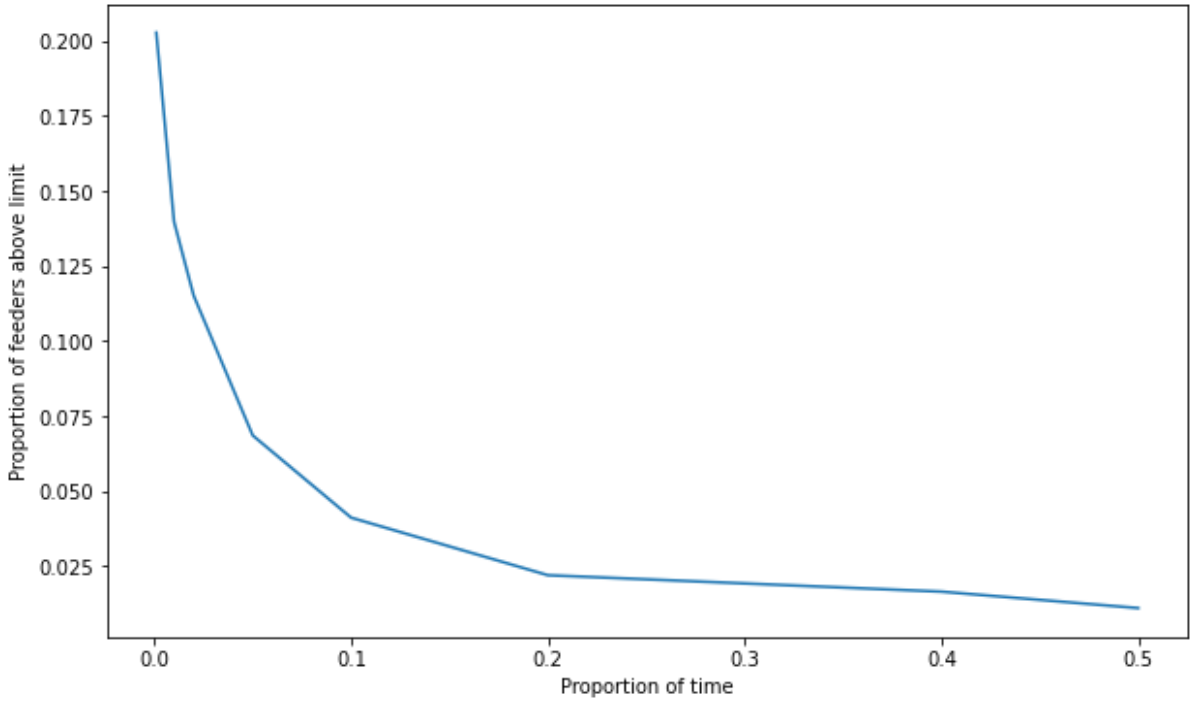

Figure 3: Proportion of feeders above maximum voltage range

A similar analysis for the minimum voltages is presented in Figure 4 and Figure 5.With the exception of two feeders, the minimum voltages within the required range for most of time. However, two feeders (in pink trace colour) have a minimum voltage below 217 V for nearly 5% of time. The two exceptions have voltages near 25 V for a significant proportion of time, and this is assumed to be a spurious data reading. Otherwise, feeders have voltages below 217 V only very infrequently.

This highlights a difficulty in interpreting the data as it is not possible here to distinguish between a ‘normal’ operation and occasional low voltages due to a power cut. Zero values have been excluded from the analysis but very low voltages that might occur as power is lost or restored remain in the data.

Figure 5 shows the minimum voltages calculated over all the time samples and plotted as a CDF of the proportion of meters on the feeder.

The likelihood of feeders being below limits is shown in Figure 6. Apart from one feeder that has readings near zero for a significant length of time, it is less likely to find low voltages below the minimum limit than to find high voltages above the maximum limit. There are around 1% of feeders that have a minimum voltage below the limit for 1% of time.

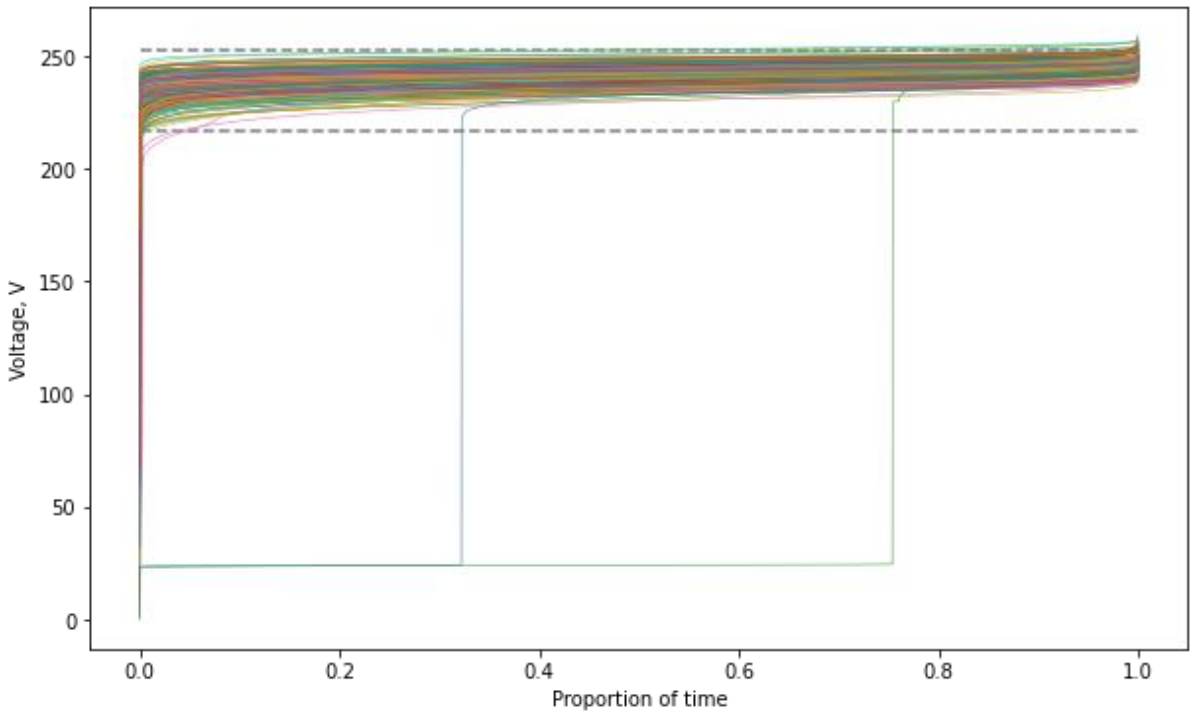

Figure 4: Minimum voltage over all meters for each LV feeder

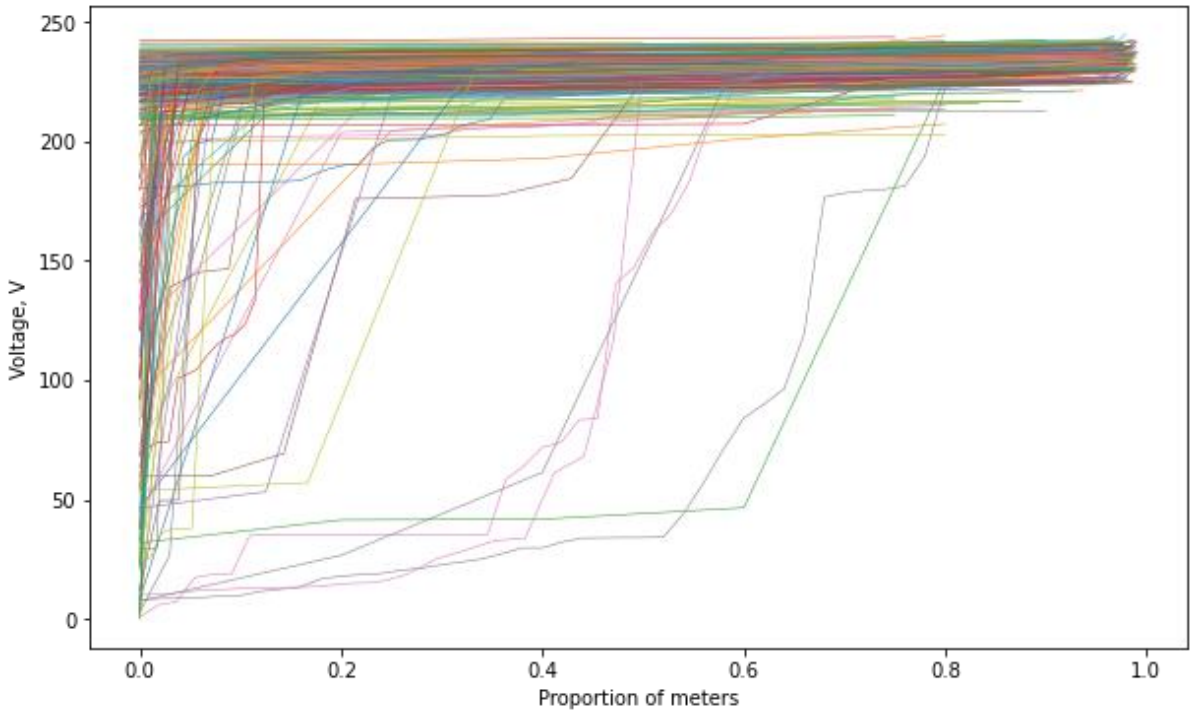

Figure 5: Minimum voltage over all time samples for each LV feeder

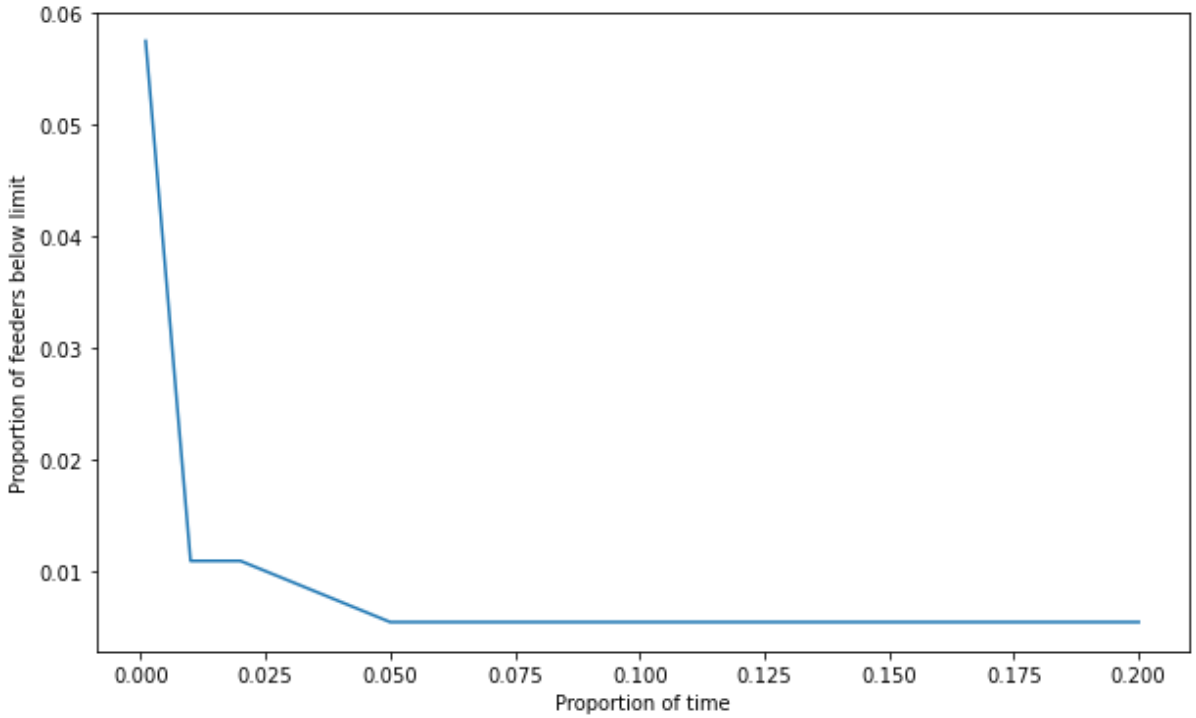

Figure 6: Proportion of feeders below minimum voltage range

2.2 Voltage headroom

An upper voltage headroom can be calculated as the difference between the limit of 253 V and a maximum voltage value found from the smart meter measurements.

In Figure 7 the maximum voltage is calculated for each feeder based on the maximum over all smart meters and over all available time samples. As expected from the plots shown above, some feeders have no headroom available, but others have up to 12 V available for additional voltage rise, such as would be generated by V2G.

In Figure 8 the headroom is based on the maximum voltage over all feeders that is not exceeded for 99% of time samples. This avoids degrading the headroom due to rare transient events, but raises a risk that these could coincide with voltage rise due to V2G and so make the extreme voltages worse. This 99% confidence approach gives an increased capacity for V2G, such that there are fewer feeders with no headroom at all, and many feeders have between 5 V and 10V available, with up to 15 V in one case.

Equivalent plots are shown in Figure 9 and Figure 10 for the lower voltage ‘headroom’. There is a much greater margin available for increased voltage drop than for voltage rise, either with absolute minimum voltages of a 99% confidence level over the time samples.

Voltage ranges with 99% confidence limits have been adopted for the purposes of the simulation modelling as this gives a more reliable estimate of the headroom typically available. There is a risk with using the absolute maximum and minimum values that the headroom available would reduce as longer time periods of data are used in the analysis. There is a greater chance of finding an extreme high or low voltage if longer periods of smart meter data are assessed.

Selecting a 99% confidence level does not imply that deviations from permitted voltage ranges can be accepted in 1% of cases, but allows the analysis to move forward without the available headroom being degraded by occasional deviations from normal operating conditions that are not the main concern of the V2G control methods.

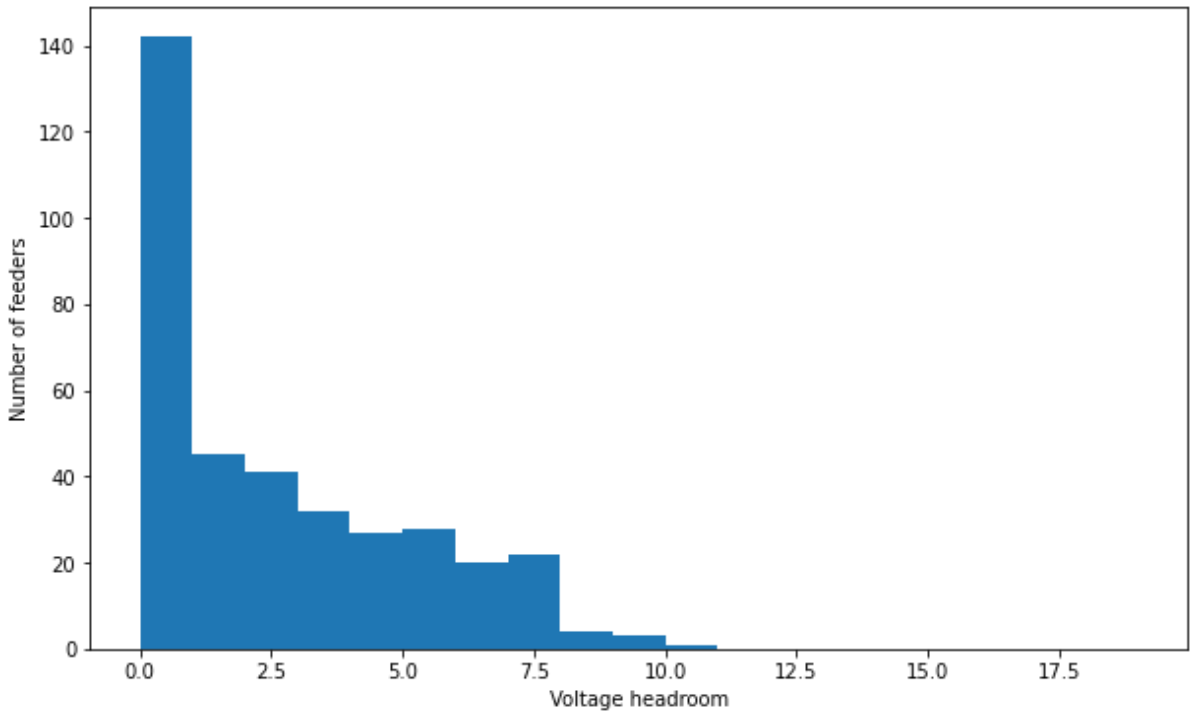

Figure 7: Upper voltage headroom for maximum voltages over all meters and time samples

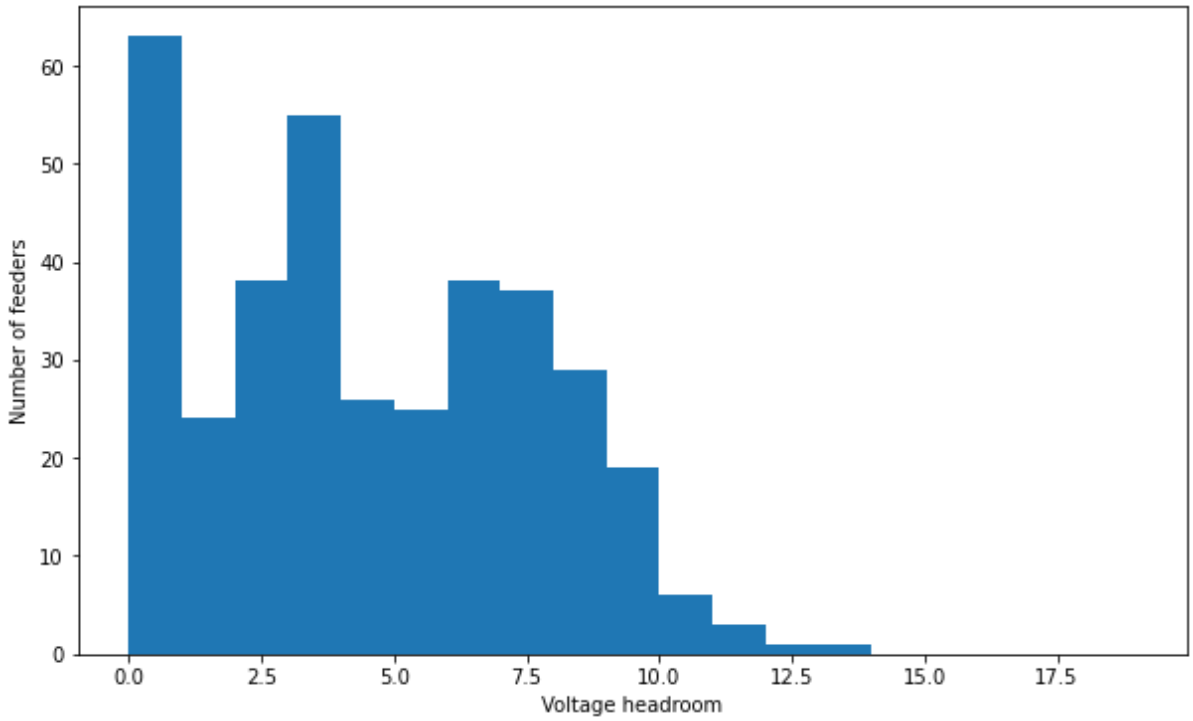

Figure 8: Upper voltage headroom for maximum voltages over all meters and 99% of time samples

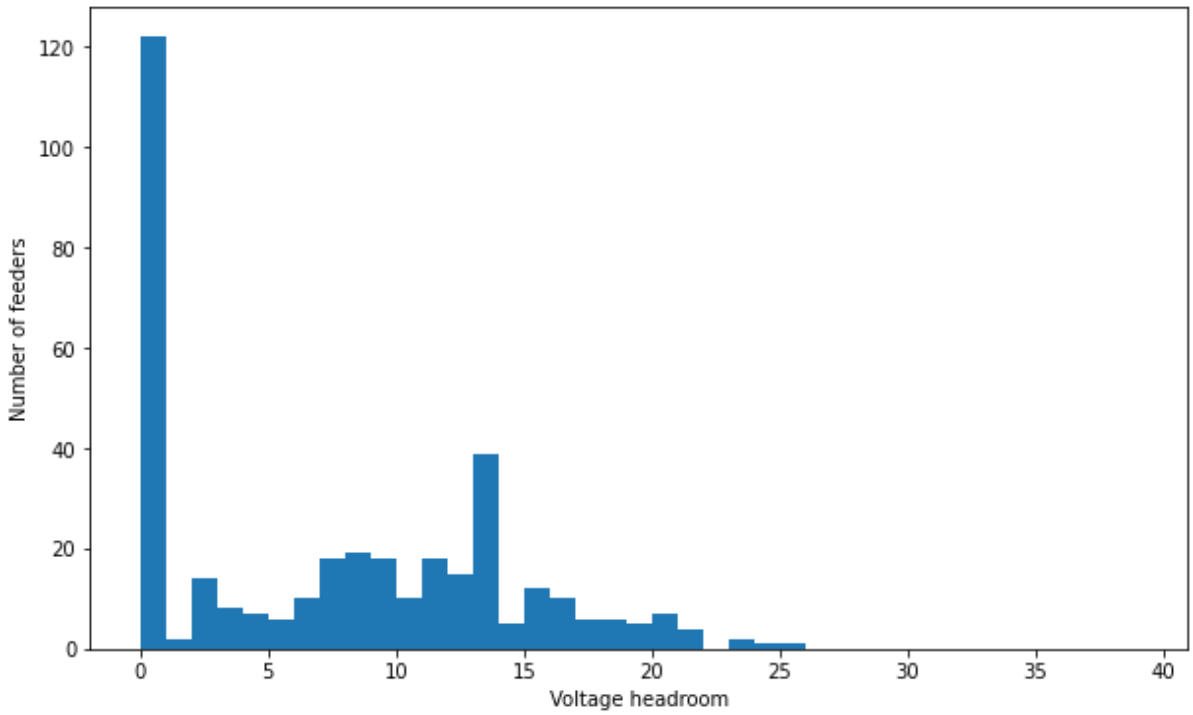

Figure 9: Lower voltage headroom for minimum voltages over all meters and time samples

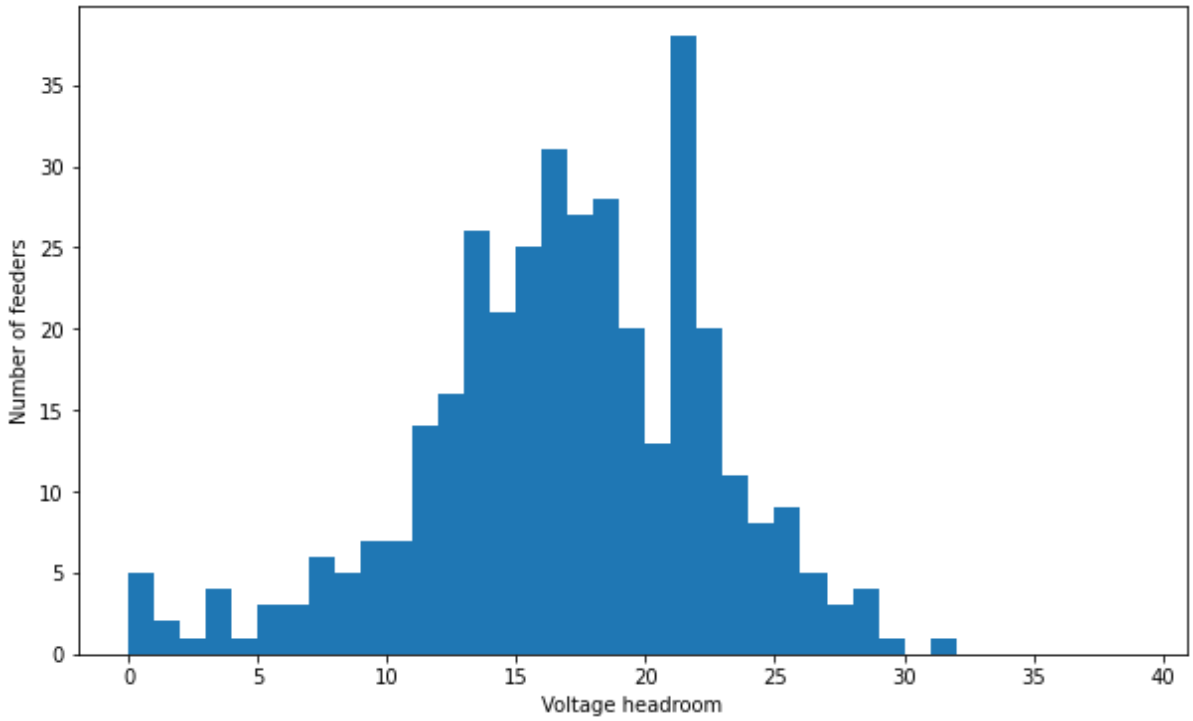

Figure 10: Lower voltage headroom for minimum voltages over all meters and 99% of time samples

2.3 Thermal capacity headroom

Typically, the bottleneck on an LV feeder in terms of thermal capacity is the mains cable that runs from the substation to the first junction, as this section carries the aggregated current from all the connected customers. There may be other constraints within the LV feeder mains, for example of the conductor sizes are tapered and if the capacity is reduced by a greater proportion than the reduction

in aggregated current, but for a first approximation, the thermal headroom can be considered on the basis of the first mains branch from the substation.

Smart meter demand data cannot be provided for individual domestic customers due to privacy constraints and is instead recorded as an aggregated demand for each LV feeder. This is therefore an appropriate metric for comparison with the rated capacity of the LV feeder cables.

This comparison also needs to take into account the additional customers without smart meters, and for which the only demand metric available is an estimated annual consumption (EAC).

The demand data for each LV feeder is imported and a list of customer meter point administration numbers (MPANs) for the customer connections is defined. Demand data has been obtained for the period from 1st July 2024 to 31st July 2025, though with less complete coverage before October 2024.

Aggregated demand data samples are excluded for values over 1×10^6.

The smart meter demand data provides four power quadrant values, active import P_{AI} and export P_{AI}, and reactive import Q_{RI} and export Q_{RE}. The mean complex power S_{SMETS} for any half-hour period is therefore given as

$$S_{\mathrm{SMETS}} = (P_{\mathrm{AI}} - P_{\mathrm{AE}}) + j \cdot (Q_{\mathrm{RI}} - Q_{\mathrm{RE}})$$

The data is assumed to be recorded in units of watt-hours and on a half-hourly basis. Powers in watts are therefore found by multiplying the recorded complex power by a factor of 2.

There is no phase information within the aggregated demand data and so an approximation is made that the demand can be partitioned into phases in proportion to the EAC for each customer. The method also needs to allow for the variable number of smart meters providing a reading, in some cases where there are fewer smart meters at an earlier time as some had yet to be installed, but also allowing for instances where not all smart meters provided a reading. The complex power for the feeder s_{FEEDER} is then calculated as

$$S_{\mathrm{FEEDER},t,p} = S_{\mathrm{SMETS},t} \cdot \left(\frac{N_{\mathrm{SMETS}}}{r_{\mathrm{SMETS},t}}\right) \cdot \frac{\sum_{i \in N_{\mathrm{SMETS},p}} \mathrm{EAC_i} + \frac{1}{3}\sum_{i \in N_{\mathrm{non-SMETS}}} \mathrm{EAC_i}}{\sum_{i \in N_{\mathrm{SMETS},p}} \mathrm{EAC_i}}$$

where $s_{\mathrm{SMETS},t}$ is the aggregated smart meter demand power for sample t. The aggregated demand is scaled by the ratio of the total number of smart meters N_{SMETS} and $r_{\mathrm{SMETS},t}$ is the number of meters providing a reading at this time. The reading is also scaled for each phase by the ratio of the total EACs for the number $N_{\mathrm{SMETS},p}$ of smart meters on this phase, and a third of the total EACs for the number $N_{\mathrm{non-SMETS}}$ of non-smart meters on the feeder for which the phase is unknown.

The maximum conductor current amplitude for each half-hour period $I_{\mathrm{max},t}$ is calculated assuming a nominal voltage of 230 V and taking the maximum over the complex powers for each phase

$$I_{\mathrm{max},t} = \max_p \frac{\left|S_{\mathrm{FEEDER},t,p}\right|}{230}$$

The maximum current can be calculated over the available time samples and compared with the minimum current rating of the feeder main from the substation. This current rating is derived from the initial simulation results for the selected feeders [3] where network data was imported with Connect LV cable type impedances associated with each cable impedance line code. The simulation software finds the minimum capacity rating, where this is known, for all cable branch sections from the substation to the first service cable junction or the first mains junction.

Figure 11 shows the distribution of maximums current for LV feeders, below 300 A for most feeders but with a few outliers where the maximum current can be close to 500 A. The current rating differs for

each feeder so Figure 12 shows the utilisation ratio for each feeder, defined here as the ratio of the maximum current to the 'winter cyclic' cable capacity rating. This ratio is mostly below 1.0, but for a few outlier LV feeders there is a utilisation greater than 1.0. As with the voltage headroom analysis, these maximums can occur for a short period of time. With a few exceptions, feeders have a utilisation of up to 0.7 times the rated cable capacities.

A current capacity headroom is calculated for each LV feeder by taking the difference between the maximum utilisation and a nominal utilisation of 0.7, as shown in Figure 13. Feeders with an existing utilisation of 0.7 or above have a headroom of zero. So, for example, there are several histogram bars in Figure 16 for utilisation around up to 0.2, and these map to 0.5 or over in Figure 17.

Although there are a few feeders where the exiting demand appears close to the cable rating, most feeders have spare capacity. It is possible that on those where the utilisation appears close to the rating, the high estimated current is due to the various approximations relating to the non-smart meters, and so this result does not necessarily indicate that the feeders are over capacity in practice.

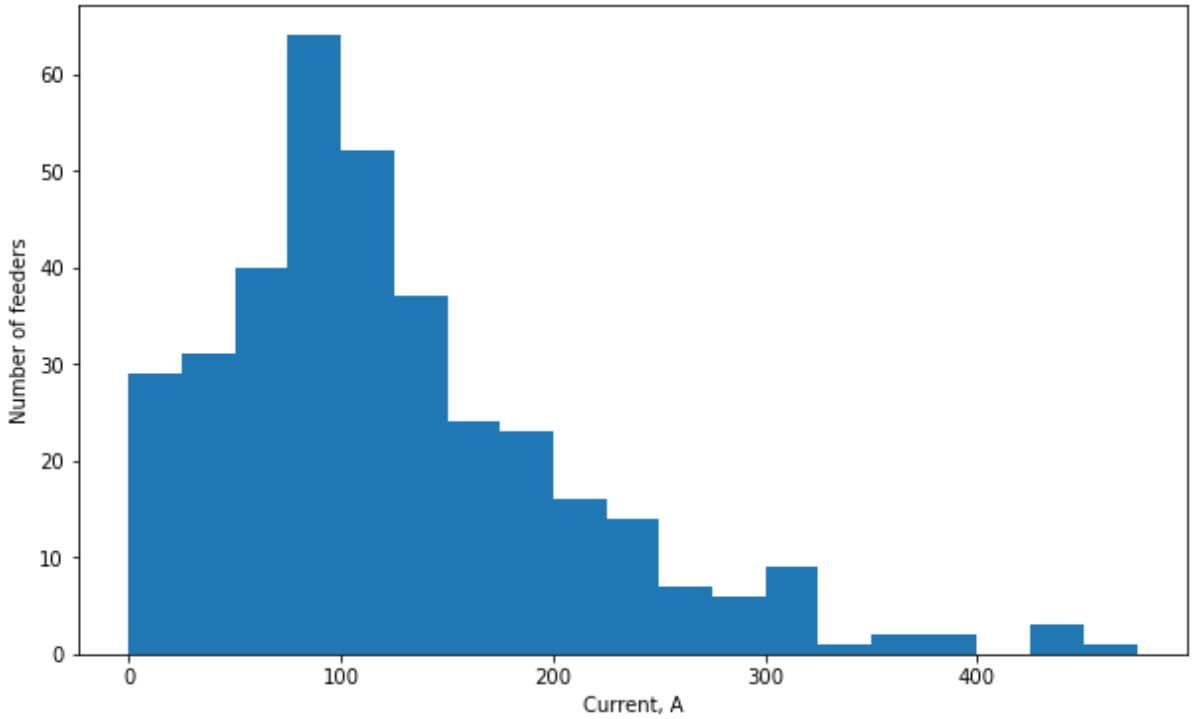

Figure 11: Maximum currents for each LV feeder

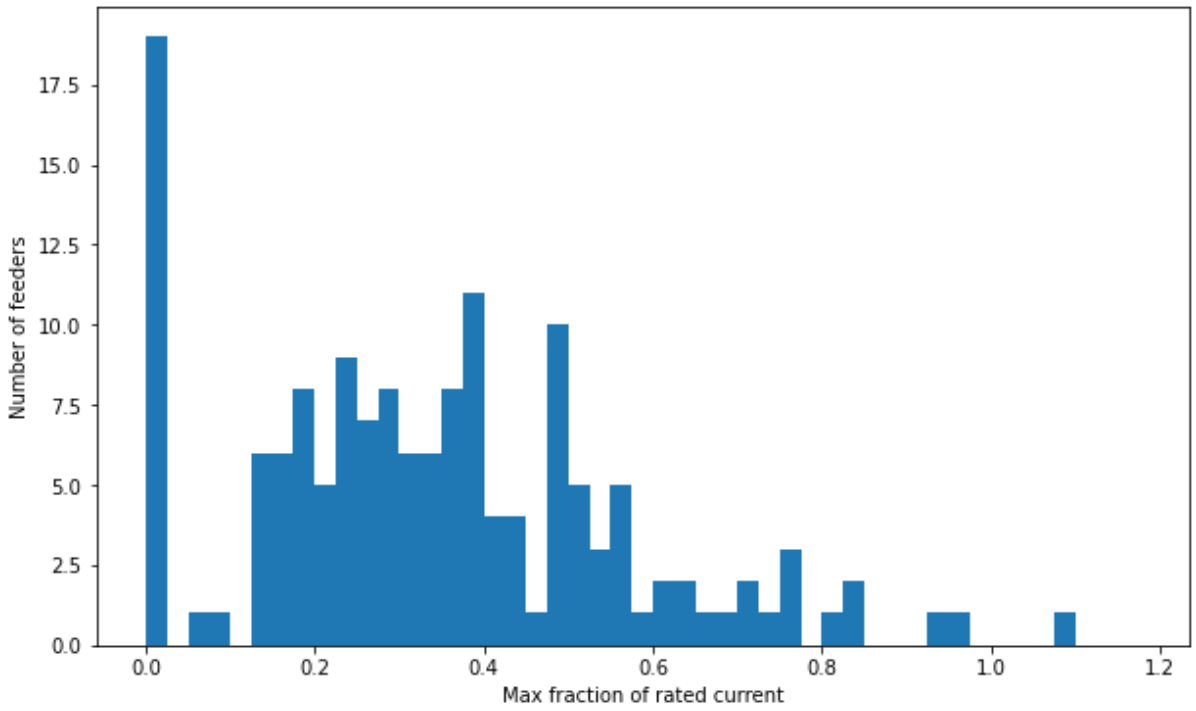

Figure 12: Maximum capacity utilisation for each LV feeder

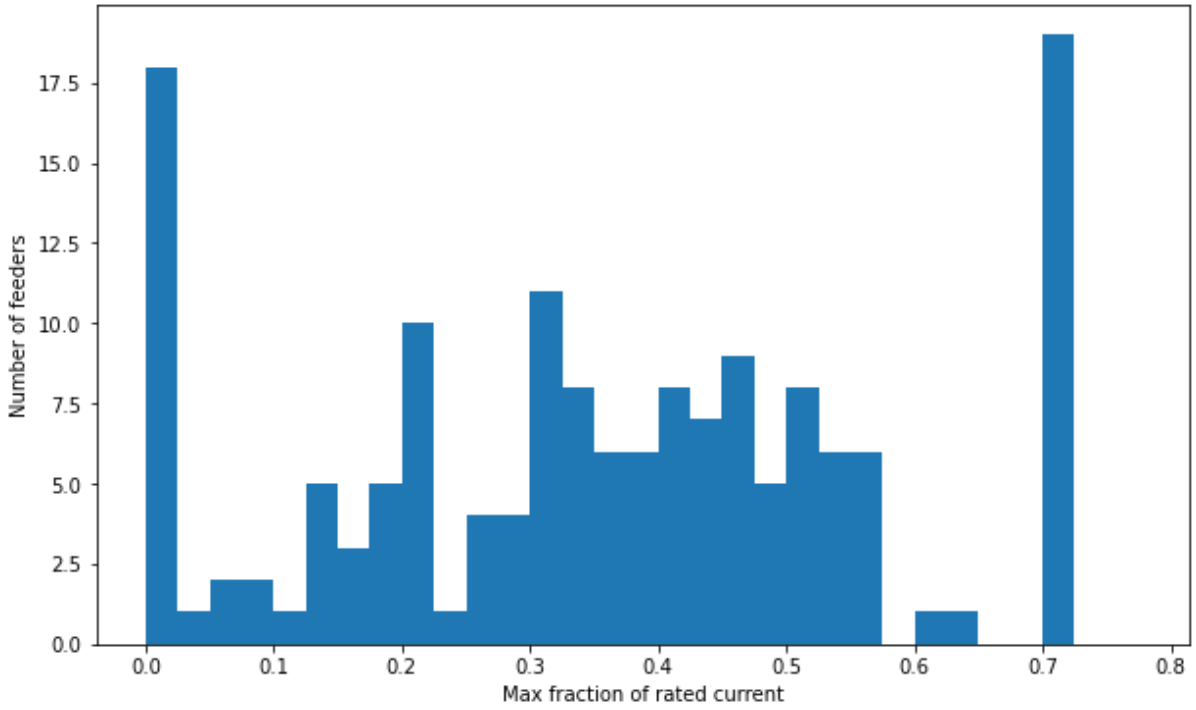

Figure 13: Current capacity headroom

3. Simulation method

The simulations presented in this report follow the outline plan from the WP1 D1.2 Control Algorithm Definitions.

The simulations described here develop several new steps in the power-flow analysis method for LV network modelling.

1. Customer connections are configured to use their detected phase groups
2. Smart meter voltage data is integrated into the power-flow analysis so that the model aligns with real-world data
3. A super-position method has been developed to model the added impact of the V2G imports and exports as a deviation relative to the existing demand
4. The iterative power-flow solver, based on the forward-backward sweep method, now also operates iteratively to apply the volt-watt and volt-var control methods

These modifications have added to the complexity of the software and so it has been necessary to include a simplification such that the simulation operates only for worst-case time steps, rather than for all the time steps in the data series. The development of this simulation method has been a time-consuming process and the simplification relating to the time steps maintains the processing time for the models within a feasible scale such that it is still possible to include each of the selected substations and feeders and a variety of different configuration options.

3.1 Phase identification

WP2 D1 describes the phase identification process. This uses smart meter voltage data to assign customers with three phase groups. Since there is mostly no independent substation monitoring that can act as a reference for the three-phase voltages, the three customer phase groups cannot be associated with phase L1, L2, L3. However, since the key concern is to model the effect of unbalance, it is sufficient to assign the customers into the three phase groups and knowledge of the actual phase conductors would not change the results. The phase identification assigns the customers to phase groups A, B, C, which are arbitrarily associated with L1, L2, L3 in the model.

When constructing the network model, all customer MPANs are modelled here as single-phase connections. For the simulations described in WP1 D3 the phases were randomly assigned but now use the detected phases.

3.2 Smart meter data test cases

The model makes use of both the smart meter voltage data for individual customers and the aggregated demand data for each LV feeder. This data effectively provides a calibration offset to the voltages seen at each customer node and to the total demand power supplied into the feeder at the substation busbar.

Voltage data is imported for each customer and re-sampled to a common half-hour time-series, as described in section 2.1. Readings are excluded for each meter if the voltage falls outside a band of 200 V to 300 V.

Since the aim here is to assess the effectiveness of V2G control algorithms based on voltages measured at the customer connection, the simulation requires that a reading is available for each of the V2G devices in the model. Time steps are therefore removed from the set of available data if any of the smart meters on the feeder has not provided a reading.

In some cases, for example where a smart meter has only recently been installed, or where the data capture process is very intermittent, this filtering process can result in very few half-hour periods remaining valid. Smart meters are therefore dropped from the set if they have provided readings for less than 50% of half-hour periods in which any of the meters on the feeder have provided a reading. While this modification reduces the number of smart meters available to the model, by removing a few meters with data communication problems the overall information content available to the simulation process can be optimised.

The half-hour periods are then further selected to ensure that an aggregated demand reading is available.

Once all these steps have been completed, the remaining data set has a series of half-hour time samples, each of which has an aggregated demand reading, and a number of smart meters that each have a corresponding voltage reading. The aggregated demand reading is represented as a complex power $S_{\mathrm{SMETS},t}$, as described in section 2.3.

This time-series data, for which both demand and voltage readings are available, is used in the power-flow analysis as a baseline condition. The additional impact of V2G can be modelled in combination with this baseline case.

Earlier modelling in this work package selected specific test case half-hours, but the results presented here included all of the half-hours for which data is available, giving a more representative indication of the frequency and probability with which voltage or thermal constraints are exceeded, and of the expected export power from the V2G installations.

For each of these test cases, the voltages at each smart meter and the aggregated demand are stored for later use in the power-flow analysis, together with the time at which the test case conditions occurred.

In addition to the aggregated demand, the smart meter data also provides the number of meters that returned a reading and that are included in the aggregation. This number of readings is stored together with the voltage and demand data for each test case, and is needed to scale the total demand to approximate the additional demand from customers without smart meters.

3.3 Smart meter data scaling

The complex power of the aggregated demand data needs to be scaled to allow for the inclusion of smart meters that have not provided a reading in the selected half-hour period and for the non-smart meters.

Since the demand data is aggregated across all the smart meters on the feeder, rather than those on each phase of the feeder, a further method is needed to disaggregate the total demand between the three phases. In a business-as-usual implementation of the V2G control techniques, the smart meter groups could be customised so that a separate aggregated demand is calculated for each phase of the feeder, subject to the number of meters being above the minimum required for privacy requirements. This process could make use of phase identification techniques developed for the SMITN NIA project, or alternatively the modified methods developed within this project and described in WP2 Deliverable 2.1.

A further complication in this scaling process arises as the number of customer connections that are represented in the network data is not necessarily consistent with the expected number of meters on each feeder.

The scaling method therefore proceeds as follows:

1. A simulation model of the feeder cables is constructed, as described in WP1 Deliverable 1.3, using the network data from Electric Office and with customer connections defined by service

connections within this data, where they exist, or by adding MPANs at their locations recorded in CROWN and with a synthesized service cable to connect to the nearest point on the feeder main.

2. A set of radial feeders is identified within these cable branches. Mostly this set of feeders corresponds with the feeder identifiers listed in CROWN as having customer connections, but in some cases additional feeders are found (for example where an unused cable remains in the Electric Office data), and there are also examples where feeders listed in CROWN do not appear to have any associated cables in the network data.

3. Time-series demand data is created for each of the meters on an identified feeder, using the Elexon profiles corresponding to the profile classes defined in CROWN and scaled by the EAC for each individual customer. This demand data is created at half-hourly resolution for a time duration that matches the earliest and latest possible dates of the smart meter data.

4. For each test case listed in section 3.2, the simulation sample corresponding to the test case timestamp is identified. The simulated demand data based on Elexon profiles for this half-hour period is then used to estimate the proportion of the aggregated smart meter demand that should be allocated to each phase.

5. The aggregated demand data is scaled by the ratio of the number of meters on the modelled feeder, to the number of meters providing a reading in the aggregation. This gives an estimate of the total demand on the feeder.

6. The demand on each phase is calculated by finding the proportion of the total demand due to meters in each of the phase groups.

This process makes several assumptions:

- The demand for customers with non-smart meters follows the same temporal variations as the aggregated demand for customers with smart meters

- The proportion of demand for each customer within the total varies linearly with their annual demand based on the recorded EACs

There is also an implied approximation that the aggregated demand of the customers, after scaling to allow for the non-smart meters, equals the combined power supplied into the feeder. In reality this power would be slightly increased due to losses. A future development of this method could run a power-flow analysis on this baseline test case to determine the total power on each phase at the feeder input with losses taken into account.

After the process outline above has been completed, the test case data is updated to include the complex power on each phase at the feeder input, together with the voltage data readings as defined above.

3.4 Selection of customers with V2G

The previous simulation models described in WP1 Deliverable 1.3 used the Distribution Future Energy Scenarios from 2035 to define an anticipated proportion of 81% customers having an EV charger. These customers were selected randomly from the full set of customers in the model of feeders at each substation.

In the revised simulations described here, it is only possible to model an EV charger where smart meter data is available, since the smart meter voltage readings are required as an input to the volt-watt and volt-var control methods. On average the proportion of customers with smart meters is similar to this predicted uptake of EV chargers, around 75%, and so the model has been revised to place an EV charger at each of the customers with smart meters.

As in the previous simulations, a worst-case assumption has been applied such that every customer with an EV charger participates in the V2G exports and the subsequent re-charging.

3.5 Demand model assumptions

As noted above, the previous simulations considered a future scenario where customers would be expected to adopt solar PV, heat pumps and EV charging in addition to their participation in V2G export events.

A future control method for V2G would be based on smart meter available at this future time, such that this data would include the impacts of any other low carbon technologies that had been adopted. The demand for these low carbon technologies would not need to be synthesized – it would already be represented in the customer voltage and aggregated demand data.

The same approach has therefore been adopted for the simulations described here, such that the baseline case defined by the existing smart meter data is modelled with the addition of exports from each of the customers with a smart meter.

The existing headroom is low, or zero, on many feeders and so there is already a need for the V2G controls to be invoked without the further addition of exports from solar PV. The addition of further demand to the model due to heat pumps or EV charging would reduce the need for V2G controls to be applied to exports, as the exported power would offset an increased local demand, but the modelling would still need to consider the case that exports could occur when these appliances are not in use.

Although the exports are described in throughout this report as being associated with V2G, the modelling could equally be taken to represent a case where exports are due to domestic battery storage.

3.6 Super-position simulation method

The modelling described here use a new method that enables a simulated demand for the V2G (or battery storage) exports to be combined with real-world smart meter data representing the existing demand.

This makes an approximation that the voltage differences along the feeder can be calculated individually for each type of demand, and that these voltage differences can be added to find the total voltage difference for the combined demand.

The voltage drop ΔV along a feeder branch can be approximated as

$$\Delta V = \frac{PR + QX}{V}$$

where P and Q are the active and reactive power and V is the nominal feeder voltage, and R and X are the resistance and reactance of the feeder cable. When P and Q are negative, ΔV is negative and represents a voltage rise.

If the total demand $P + jQ$ is comprised of two separate demands $P_1 + jQ_1$ and $P_2 + jQ_2$ then

$$\Delta V = \frac{(P_1 + P_2)R + (Q_1 + Q_2)X}{V} = \frac{P_1R + Q_1X}{V} + \frac{P_2R + Q_2X}{V} = \Delta V_1 + \Delta V_2$$

Within the limits of the approximation in this voltage drop calculation, it is therefore reasonable to assume that the voltage impact of separate demands can be added.

Strictly, there are interactions between the individual loads, as the losses in each cable branch increase with the square of the current, and therefore approximately with the square of the power

delivered. The power $P + jQ$ in delivered by branches carrying both load currents therefore depends on the combined downstream losses and so does not increase linearly as in the simple approximation. However, LV feeder losses are typically less than 2% of the total demand and so these second-order effects can be neglected as being small in comparison to the more critical estimates needed about the number and power of the future V2G exports.

The following process has been used to combine the smart meter voltage data and the simulation data:

1. The simulations of the V2G operation include only the export powers or the imports for re-charging. The demand at any customer meter with no V2G (in effect those with no smart meter) is zero.

2. The impact of the existing demand is reflected in the smart meter voltages as measured. Voltage at customers without smart meters are unknown and the impact of V2G on these customers is not modelled.

3. The simulation model assumes a nominal voltage at the substation busbar of 245 V.

4. Power-flow analysis of the V2G operation calculates a voltage at each customer connection.

5. The total voltage at each customer connection is then calibrated by adding a voltage offset, given by the measured smart meter voltage minus the nominal substation busbar voltage of 245 V.

This considers that

$$V_{\text{substation,real}} = V_{\text{customer,real}} + \Delta V_1$$

$$V_{\text{substation,model}} = V_{\text{customer,model}} + \Delta V_2$$

where $V_{\text{substation,real}}$ is the unknown actual substation voltage and $V_{\text{customer,real}}$ is the measured customer voltage from smart meter data and ΔV_1 is the voltage drop due to the existing demand; and similarly $V_{\text{substation,model}}$ is the nominal voltage assumed for the substation in the model, $V_{\text{customer,model}}$ is the calculated voltage of the smart meter customer with voltage drop ΔV_2 due to V2G operation (or rise if this is negative).

We wish to find a revised customer demand $V_{\text{customer,revised}}$ where the substation voltage remains the same and the additional voltage drop ΔV_2 is added

$$V_{\text{substation,real}} = V_{\text{customer,revised}} + \Delta V_1 + \Delta V_2$$

This is determined as

$$V_{\text{customer,revised}} = V_{\text{substation,real}} - \Delta V_1 - \Delta V_2$$

$$V_{\text{customer,revised}} = V_{\text{customer,real}} + V_{\text{customer,model}} - V_{\text{substation,model}}$$

For each smart meter customer in the simulation, a voltage is therefore calculated where the customer voltage from the V2G simulation is calibrated by the additional of an offset given by the measured smart meter voltage minus the nominal simulation substation voltage.

Based on this approximation alone, the simulation could adopt any value for the substation voltage but in order to minimise the inaccuracies associated with second-order effects such as losses, as described above, it is preferable for the substation voltage to be close to the actual value. A voltage of 245 V has been found in previous work to be a reasonable estimate.

The voltage offsets are also required where the model investigates the impact of V2G with voltage control, since this would operate on the basis of the customer voltage with both demands taken into account. The control action is therefore determined by the simulated voltage plus the offsets derived from smart meter data.

A second calibration has been included so that the combined power from the existing demand and the additional V2G demand can be calculated. This uses the aggregated feeder demand from smart meter data, scaled as described in section 3.3 to allow for the proportion of smart meters that returned a reading, the additional demand from non-smart meters, and then shared between the phases on the basis of the Elexon profile demands for the selected time sample.

The scaled estimate of the existing complex demand at the feeder input for each phase is then added to the simulated demand for V2G, giving a total complex demand power on each phase. This addition is an approximation, since the combined demand would have increased losses, but the associated error will be small relative to the other uncertainties in the demand model. The combined complex power for each phase therefore allows for the potential for existing demand to be reduced by V2G exports.

The combined complex power is used in the analysis to calculate a current amplitude, assuming a nominal voltage of 245 V at the substation, such that cable utilisation relative to the current rating can be estimated.

3.7 Iterative solution method

The power-flow analysis using the forward-backward sweep method already uses an iterative approach. It is assumed that the power of customer loads remains constant, and that the current drawn by these loads can be calculated based on this defined power and the simulated voltage. The method begins by calculating the current for each of the customer loads while assuming that there are no voltage drops within the feeder. Working from the customer loads backwards the substation, the solver then calculates the aggregated current in each branch. Once these currents have been defined, it is then possible to work forwards from the substation towards the customer loads and calculate a voltage at each node allowing for the voltage drops due to the current in each branch. This gives a revised estimate of the voltage at the customer loads. Again, assuming that the customer load has a constant power, an updated current estimate can be calculated. This continues until the voltage at each node converges so that differences between iterations are below a defined limit.

A second iterative process is needed to implement the volt-watt or volt-var control methods. Power-flow analysis produces an estimate of the voltage for each node with V2G connected, and the control methods adjust the active or reactive power demand with respect to this voltage amplitude (adjusted by the offsets to allow for the existing demand, as in section 3.6). A reduction in active power at one V2G customer will affect the voltage rise seen by all other customers on the feeder. In effect, this creates a revised power-flow analysis problem which must be solved again using the process described above.

A method has been adopted here where the power-flow analysis initially considers V2G with no control methods applied. Where V2G is exporting, this will produce a high voltage rise at many customer connections.

The simulation method then iteratively applies the control at each V2G connection in turn, repeating the power-flow analysis after each iteration to calculate revised voltages throughout the network.

As this process progresses, voltage rise from exports is reduced and it is possible that no control action is required by the V2G connections that are considered later in the sequence. Those considered first will modify their active or reactive power to a greater extent. The impact on individual customers is likely to differ based on the arbitrary sequence in which the simulation requests the control actions. In a physical device these control decisions are likely to be updated repeatedly and

asynchronously by each of the V2G installations and so the modelling approach here is necessarily a simplification.

To reduce the impact of selecting any one arbitrary sequence for implementing the control decisions, the ordering is randomised for each LV feeder and for each half-hour within the available time-series data obtained from the smart meters.

Rather than running the control iteration separately for each half-hour, this is implemented by solving each of the half-hour samples in parallel. The power-flow analysis is still repeated one for each V2G connection, but the V2G connections are updated in a different randomisation sequence for each of the half-hours in the data.

Work package 3 of this project will address the potential instabilities of this iterative process in further detail. For now, a control action is requested once from each V2G installation and there are no further updates once each V2G installation has been checked.

3.8 Voltage control methods

Three voltage control methods have been modelled.

1. Volt-watt control, where the active power is progressively ramped down to zero for voltages increasing from a lower threshold setting to an upper threshold setting, as shown in section 2.3.7 of WP1 Deliverable 1.3 but with variable thresholds defined for these simulations.
2. Volt-var control, where the reactive power is progressively ramped up to a defined percentage of the V2G permitted export power. Reactive power is imported at high voltages and exported at low voltages, with the magnitude defined by a ramp between lower and upper thresholds. This is shown in section 2.3.8 of WP1 Deliverable 1.3.
3. Volt-watt-var control, where both the above techniques are applied. The volt-var control action is defined first, as this may set a limit to the available active power. Further reductions to the active power are then applied based on volt-watt control. Both control actions are taken at the same time, rather than for example by only reducing active power if importing reactive power has not achieved a sufficient voltage reduction.

Throughout the simulation results, the same volt-watt control settings apply when results are compared for volt-watt control individually or when used in combination with volt-var control.

3.9 V2G parameters

As in the previous WP1 Deliverable 1.3 simulation work, the EV and V2G parameters have been assumed as follows:

EV charger rated power	7 kVA
V2G maximum export power	3.7 kW
Volt-var reactive power imports	Ramping from 0% of charger rated power at 240 V to 60% of charger rated power at 258 V
Volt-var reactive power exports	Ramping from 0% of charger rated power at 220 V to 40% of charger rated power at 207 V

For the volt-var reactive power of 60% of the charger rated power, the inverter would retain capacity to provide 80% of this rated apparent power of 7 kVA, or 5.6 kW, as active power. Since the export power is in any case limited here to 3.7 kW the use of reactive power will not further reduce the available active power range for exports. However, any reactive power exported when re-charging will reduce the scope for active power to be imported.

The frequency of operation for V2G and the duration of the export periods is so far highly uncertain. Rather than model a speculative operational time for V2G, the simulations consider two extreme cases:

1. Where V2G is exporting continuously
2. Where V2G is importing continuously (as with re-charging)

This approach allows the worst-case current loading and voltage rise to be calculated, taking into account the possibility that V2G could be active at any time.

The metrics considering over-voltage and current over-loads are probabilistic. To calculate the probability of an overload happening at any time of the day, the model therefore needs to consider V2G exports and imports at any time of the day.

This method implies an assumption that exports and imports at any time of the day are equally likely.

3.10 Simulation metrics

Results in the following sections present five key metrics, as follows:

Maximum voltage

The maximum voltage at any customer meter, calculated over all the half-hour periods in the simulation time series

Probability of over-voltage or under-voltage

The number of half-hour periods where the voltage exceeds 253 V, or is below 217 V, divided by the number of half-hour periods in the simulation time series

Maximum loading

The maximum aggregated current in the first feeder cable from the substation, calculated over all the phases connected to the distribution transformer, and over all the half-hour periods in the simulation time series at any customer meter, expressed as a fraction of the rated capacity of the cable.

Probability of over-load

The number of half-hour periods where the aggregated current in the first feeder cable from the substation exceeds the rated capacity of the cable, divided by the number of half-hour periods in the simulation time series

Mean export power

The mean export power, averaged over the number of half-hour periods in the simulation time series

The probability and mean export power metrics are also shown as summary values where results are averaged over all the customers on all the feeders for the substations included in the simulation results.

4. Results with V2G

The first results presented here show the impact of V2G in the absence of any control methods. Each set of results include data for 134 LV feeders from the list of selected substations previously used in the WP1 Deliverable 1.3 simulations.

Figure 14, Figure 16 and Figure 17 show results for a baseline using time samples from the existing smart meter data, as described in section 3.2.

Figure 14 shows the maximum voltages at customer meters, with each point on the curve representing the maximum voltage over all half-hour time samples for one meter. The plot shows that V2G will cause over around 70% of customers to have a worst-case voltage rise exceeding the upper limit of 253 V, shown on the plot as a dashed line. As noted previously in section 2.1, even in the baseline case, calculated from the existing smart meter, a few customers will have worst-case voltages that exceed the limit.

Figure 15 shows the probability of over-voltages for each of the customer meters. For the baseline case with existing smart meter data, even though the maximum voltage may exceed the limit, this will be a rare event for more customers. However, over-voltages with V2G exports are much more likely and for around 30% of customers there is close to a 100% probability of having a voltage over the limit. Voltage rise with V2G exports is therefore expected to cause higher maximum voltage, in addition to causing a much higher probability of over-voltages.

Clearly this would be a concern, although it should be noted that this represents a case where all customers with V2G are exporting simultaneously. The customers included in the model here are also on feeders that have originally been selected as having a higher risk of voltage rise due to V2G in the absence of any further reinforcement work.

The aggregated currents shown in Figure 16 are also high. This plot shows the maximum current for each feeder, calculated over each of the half-hour samples and over each of the phases at the transformer busbar.

The corresponding cable utilisation is shown in Figure 17, where the aggregated current in each of the feeder mains is expressed as a ratio of the cable rating. Around 80% of feeders remain below 100% utilisation but 20% of feeders are above this, with some reaching twice the rated capacity. This confirms that the thermal limits, at least in the rare worst-case scenario that has been modelled, will also be a concern, although it may be acceptable for cables to run at higher than their rated capacity for very short periods.

Figure 18 shows the corresponding probability of current over-loads. In the baseline case with existing smart meter data, these over-loads have a probability of occurrence that is near to zero, confirming that the worst-case currents shown in Figure 17 are very rare. This is not the case with V2G exports and around 10% of feeders can be expected to be overloaded at any time that simultaneous V2G exports occur.

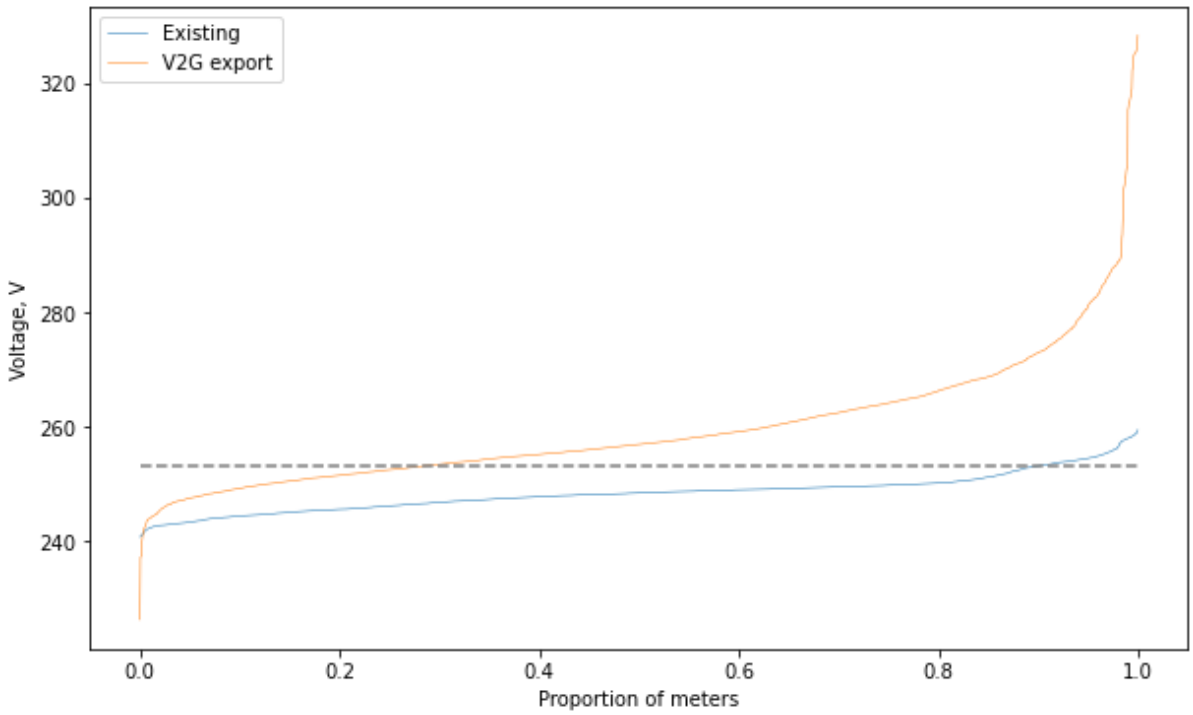

Figure 14: Maximum voltages with V2G exports

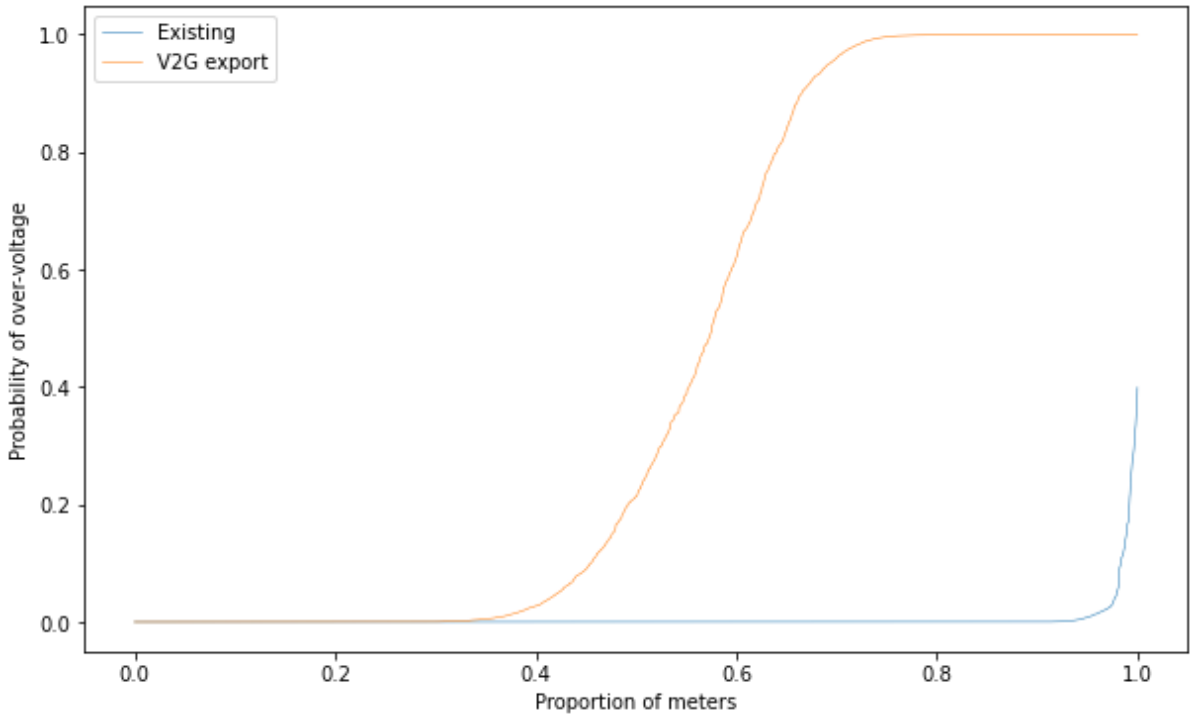

Figure 15: Probability of over-voltage with V2G exports

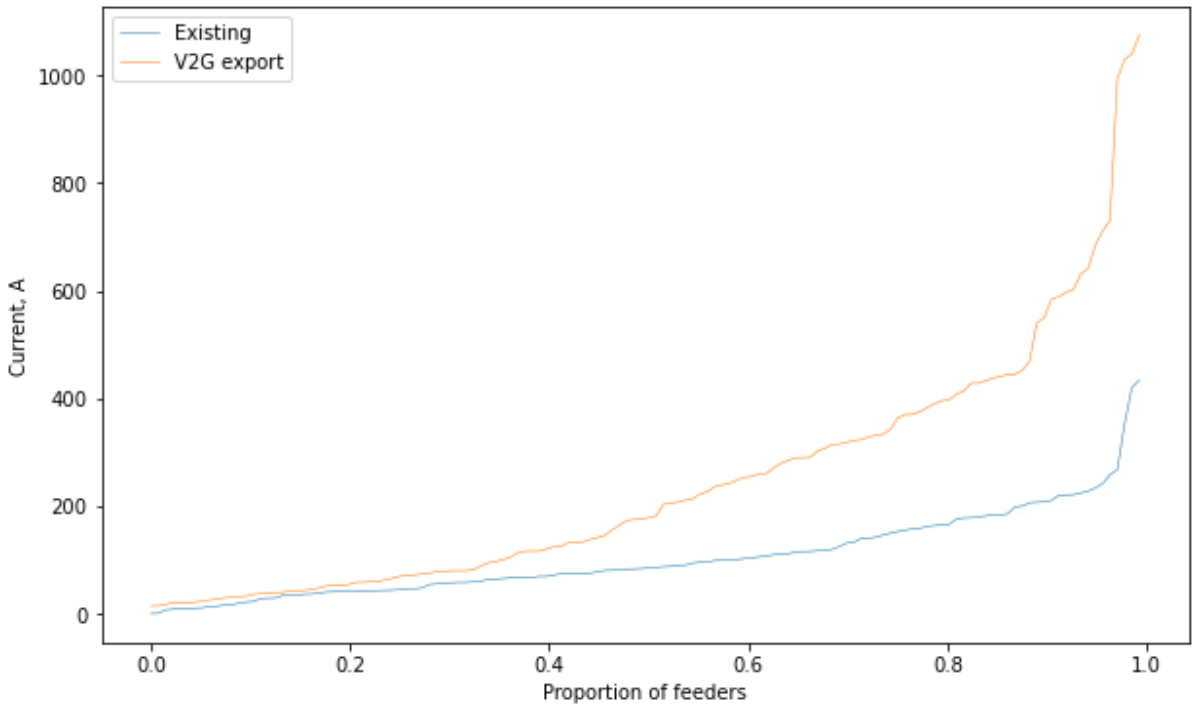

Figure 16: Maximum feeder currents with V2G exports

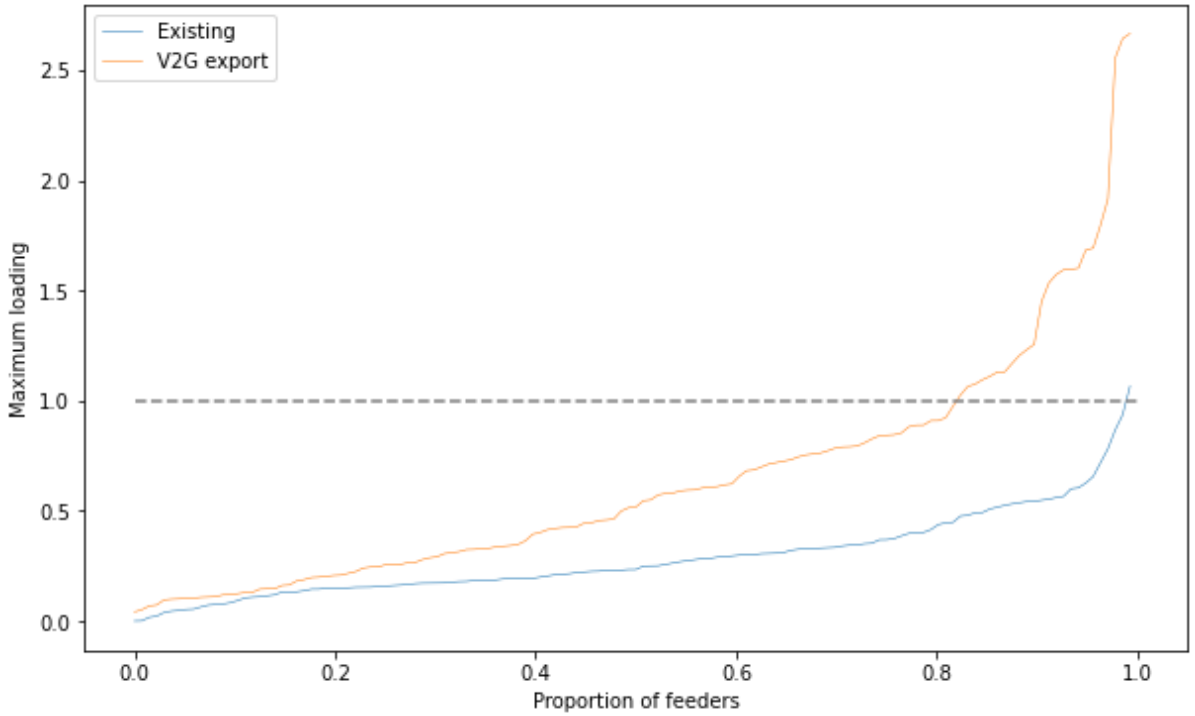

Figure 17: Maximum cable utilisation with V2G exports

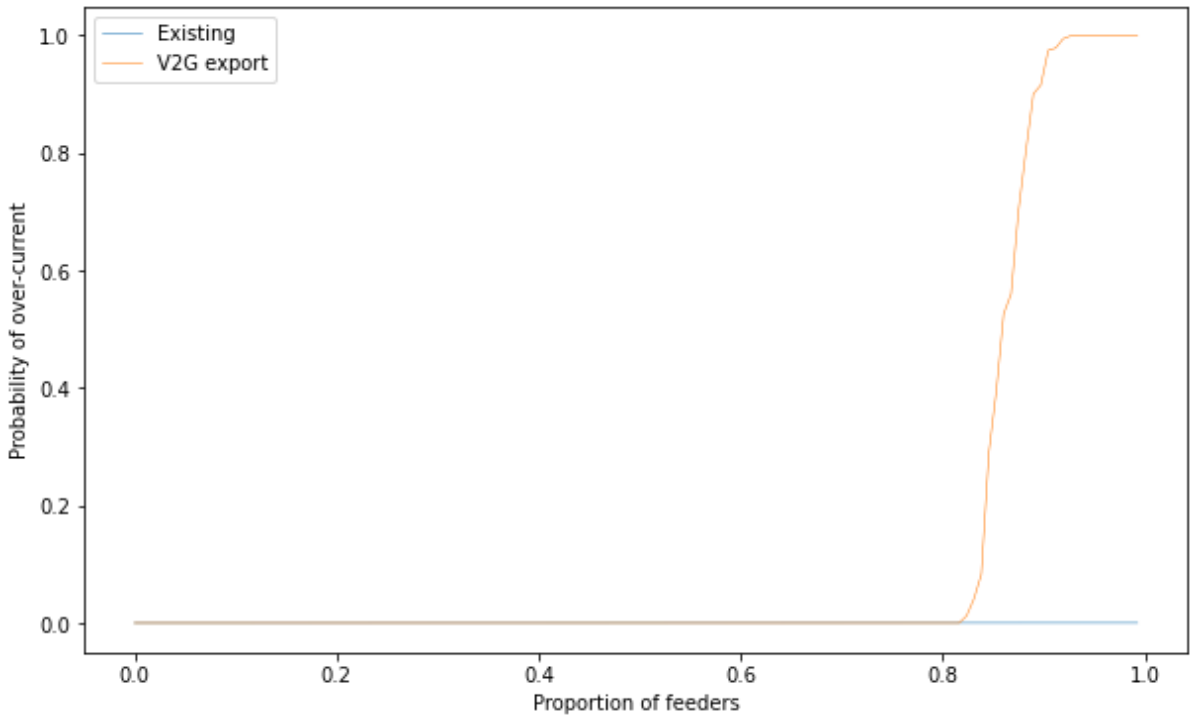

Figure 18: Probability of over-load with V2G exports

The plots below show the impact of coordinated V2G re-charging, again modelling the case that this could occur at any time over the simulated time series

Voltages shown in Figure 19 are out of range for over 40% of customer meters with V2G imports, but only a small percentage of customers would expect to have under-voltages based on existing smart meter data. The probability of under-voltages, shown in Figure 20 is near zero for the baseline case, again confirming that the existing problems with low voltages are very rare, but over 20% of customer meters would expect to have voltages below the limit with simultaneous V2G imports.

The maximum cable utilisations and probability of over-load shown in Figure 21 and Figure 22 are also extremely high, with 50% of feeders having a worst-case utilisation higher than the cable rating, and with 50% of feeders also expected to have over-loads whenever V2G imports occur.

The cable utilisations for importing are higher than those for exporting since the full 7 kW capacity of the EV charger can be used, rather than the permitted export of 3.7 kW. As noted in the WP1 Deliverable 1.3 report, some means of diversifying the recharging would be needed to avoid these very high currents. Customer may delay recharging until an off-peak period, for example when a low-rate tariff begins. Since the start time for this tariff is likely to be the same for all customers, it is plausible the recharging demand will be synchronised as modelled here.

A further risk is that the proposed volt-watt control so far only applies to voltage rise. There may also be a need for a control method to apply to imported power.

A caveat to these simulation results should be noted as there is a risk of double-counting the EV charging import power. It can safely be assumed that the baseline demand does not already include any V2G exports, although it is possible that the selected test cases may coincide with exports from battery storage systems, so it appears reasonable for the simulation to add exports from each customer. However, some customers on a feeder will already have EV chargers and it is likely that the minimum voltage test cases will occur when these EV charging is active. The impact of these EV chargers is therefore already represented in the baseline smart meter data and so the demand may be over-estimated adding the V2G recharging at each of the selected customers.

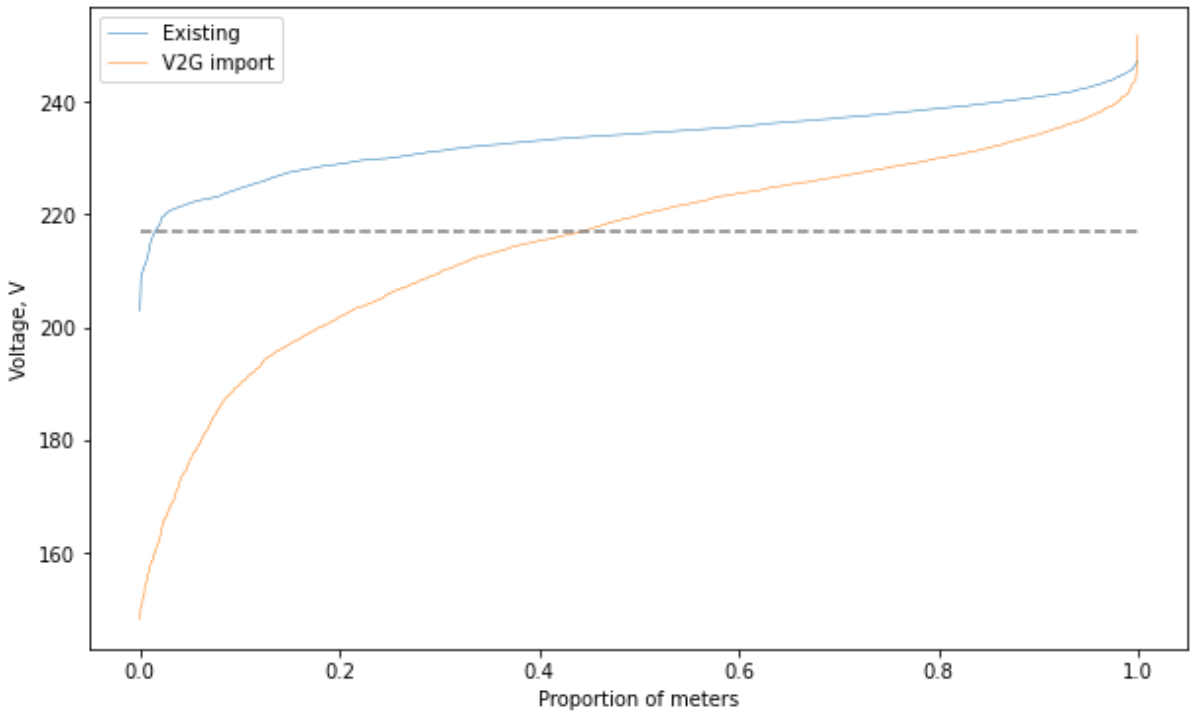

Figure 19: Minimum voltages with V2G imports

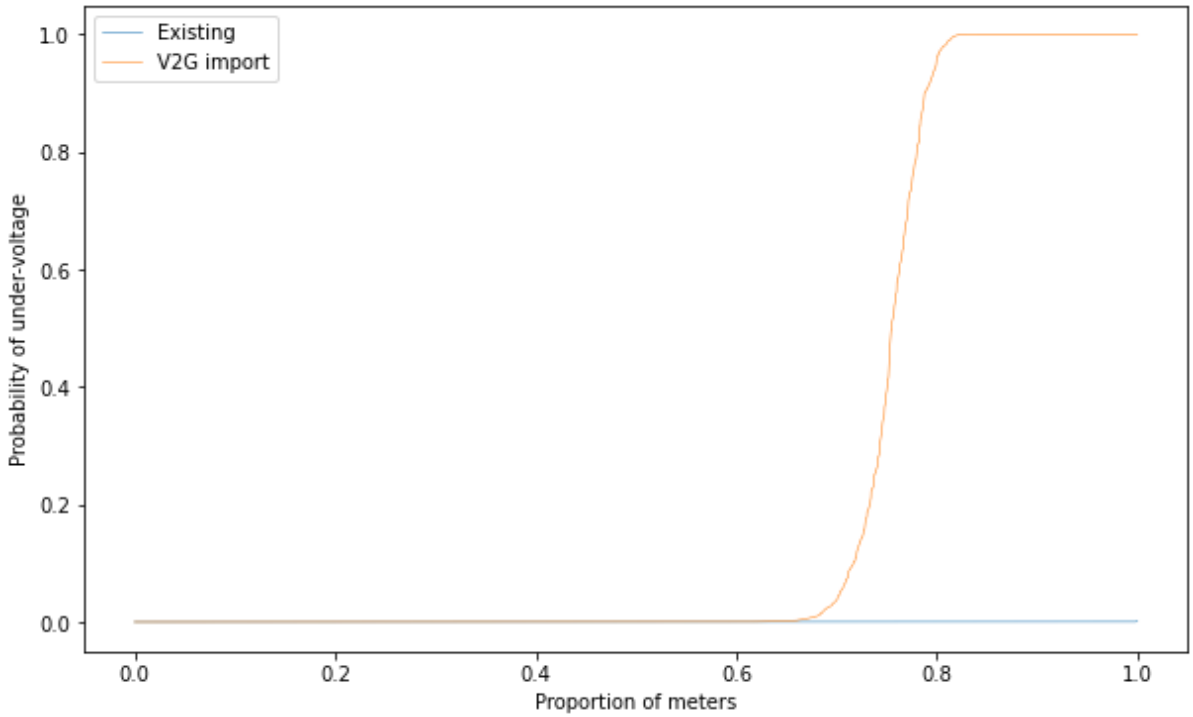

Figure 20: Probability over under-voltage with V2G imports

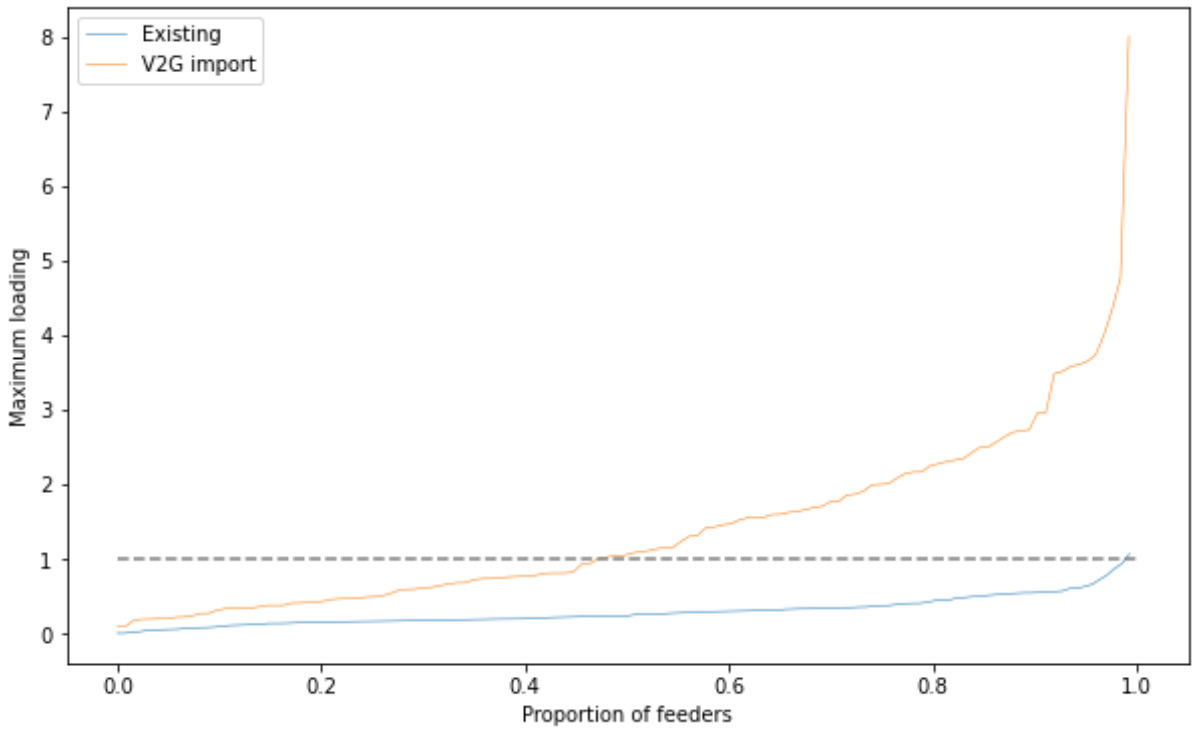

Figure 21: Maximum cable utilisation with V2G imports

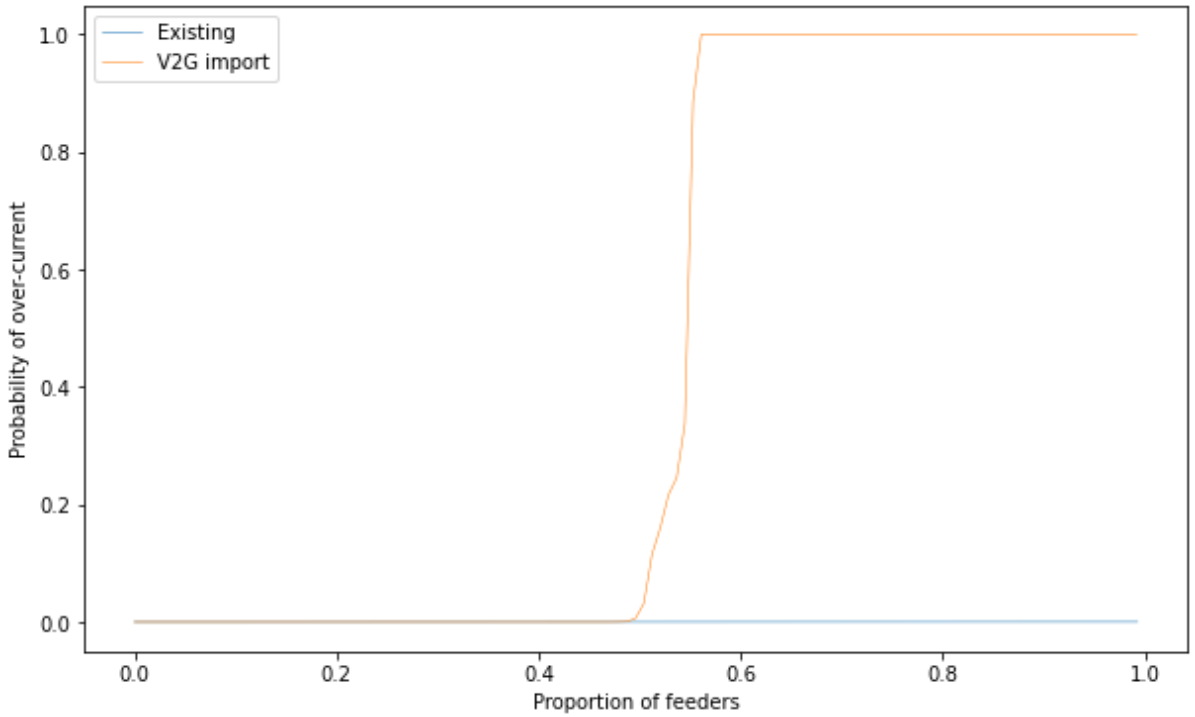

Figure 22: Probability of over-load with V2G imports

The results shown so far in this section assume that there no control has been applied to the V2G exports or imports. When exporting, the mean export power for V2G expected from the customer installations is therefore equal to 3.7 kW, which has been assumed as the rated export capacity.

Given that uncontrolled V2G exports are predicted to cause voltage and thermal overload issues on LV networks, one option would be to limit the number of V2G installations that have permission to

connect to the feeders. In an idealised scenario, the number of connections would be exactly chosen such that the worst-case voltage rise for any customer on the feeder, and the worst-case feeder current will be extended exactly to the permitted limits, but not beyond.

This scenario can be approximated using a pro-rata combination of the existing baseline case and the case with uncontrolled V2G at the rated export power for each customer. This approach maintains the same number and spatial distribution of V2G installations as in the case of uncontrolled V2G but finds a mean power level from each customer on the feeder that would just reach either the voltage rise limit, or the thermal overload limit, whichever is the more severe constraint. In practice, it is more likely that the export power from each V2G installation would be maintained at the modelled 3.7 kW, and the number of customer connections would be reduced. However, the impact can be modelled approximately by assuming that all V2G customers can connect, but at a lower maximum export power. This approximate approach avoids the statistical difficulty that would arise from the many different voltage rise results that would arise with many different subsets of the original set of V2G customers.

The export power for a V2G installation on a feeder is calculated as:

$$P_{\text{limited}} = P_{\text{rated}} \times \min(h_{\text{V}}, h_{\text{I}})$$

where P_{limited} is the export power that does not exceed limits, P_{rated} is the uncontrolled export power of 3.7 kW, and h_{V} and h_{I} are limiting factors relating to the voltage and current headroom respectively.

The voltage constraint factor h_{V} is zero if the maximum existing voltage $V_{\text{max,existing}}$ on the feeder is already above 253 V, but if headroom is available

$$h_{\text{V}} = \min_{t \in N_{\text{t}}} \left(\frac{253 - V_{\text{max,existing}}}{V_{\text{max,export}} - V_{\text{max,existing}}} \right)$$

where $V_{\text{max,export}}$ is the maximum voltage with uncontrolled V2G exports.

Similarly h_{I} is zero if the maximum existing cable current $I_{\text{max,existing}}$ is greater than the cable rating I_{rated}, but if headroom is available

$$h_{\text{I}} = \min_{t \in N_{\text{t}}} \left(\frac{I_{\text{rated}} - I_{\text{max,existing}}}{I_{\text{max,export}} - I_{\text{max,existing}}} \right)$$

where $I_{\text{max,export}}$ is the maximum current with uncontrolled V2G exports.

Headroom factors h_{V} and h_{I} are calculated for each feeder and then applied to all the meters on each feeder. This gives a cumulative probability distribution for the expected export power of V2G overall, as shown below in Figure 23. This shows that around 50% of meters would be entirely unable to export power and only around 5% of meters would be permitted to export at full power. Although this approach has the advantage that no feeder reinforcement would be required, from the perspective of V2G, but the potential contribution from V2G exports to system balancing on the grid would mostly be lost.

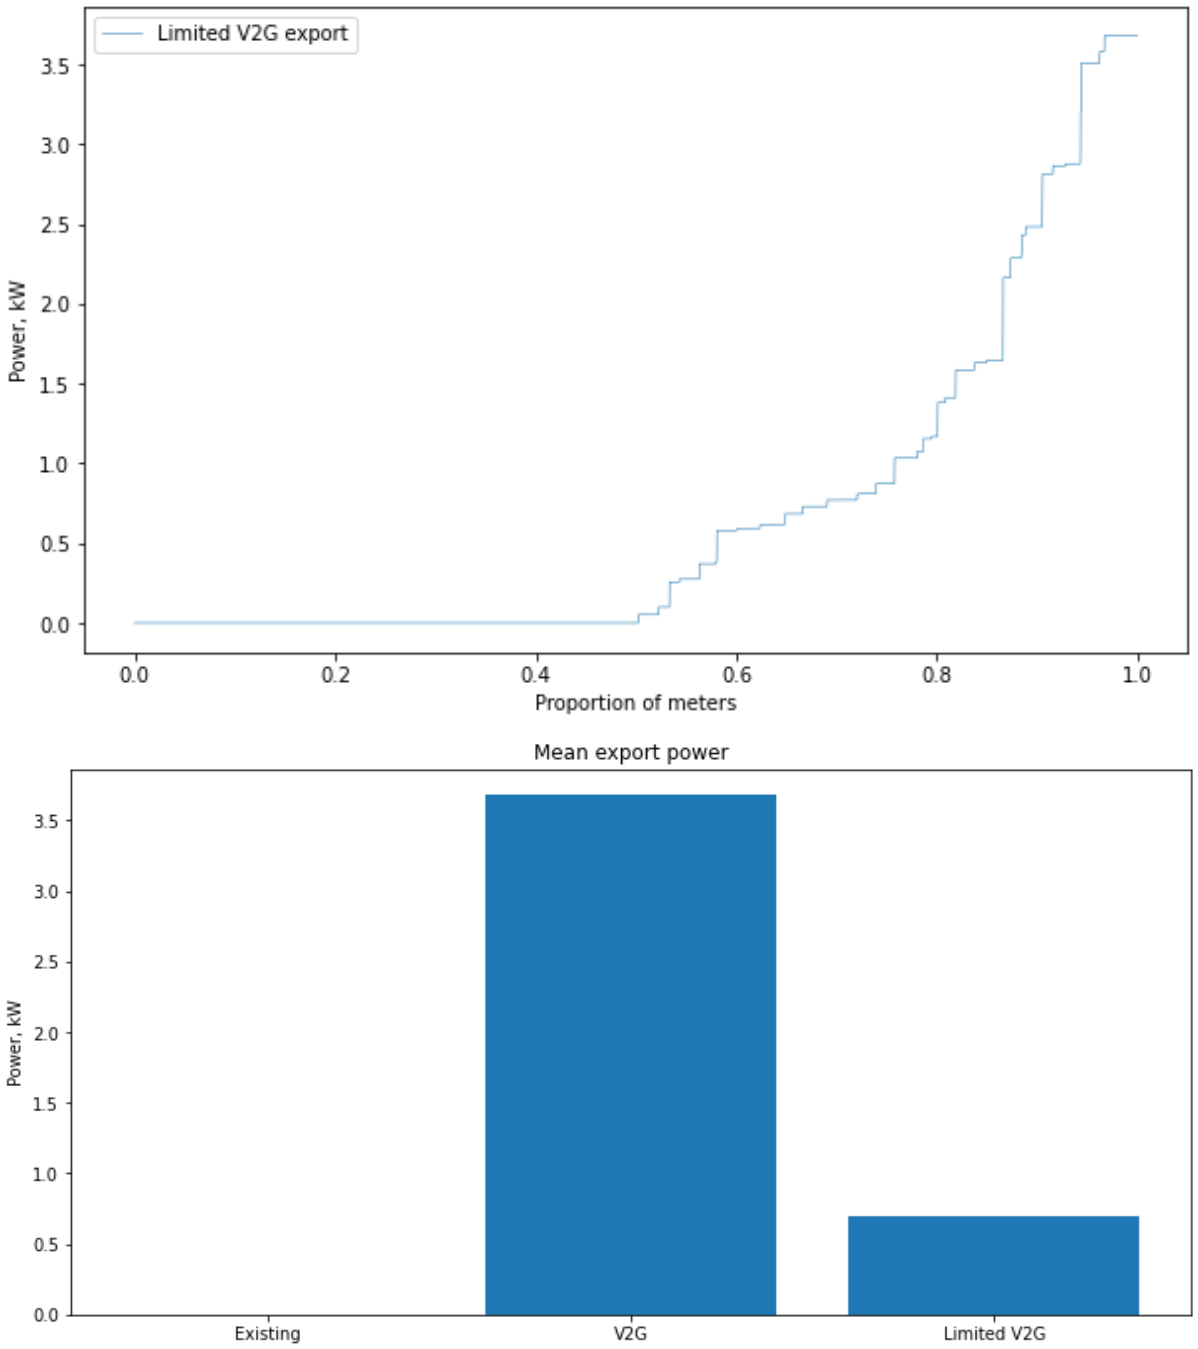

Figure 23: Mean export power for V2G for connections on feeders with voltage and current limits

5. Results with voltage control algorithms

5.1 Control algorithms for voltage rise

This section introduces the voltage control methods to mitigate the impacts of V2G.

The plots below show the impact on maximum feeder voltages and on the aggregated feeder currents in terms of utilisation of the rated capacity.

The volt-watt control method for these results uses a lower threshold of 252 V and an upper threshold of 253 V, in effect so that the control method is applied over a range that is close to the actual permitted limit.

As expected from the threshold settings, both the volt-watt and volt-var-watt control are mostly effective in limiting the voltage rise such that it remains below 253 V for most feeders. The maximum voltages in Figure 24 move closer to the limit, making use of the headroom where this is available. Figure 25 shows the corresponding probability of over-voltage with these control methods, also indicating that over-voltages are only slightly more likely with V2G using volt-watt control than in the existing baseline case.

The volt-var control method achieves a reduction in the maximum voltage but nearly 50% of feeders still exceed the limit. The probability of over-voltage with volt-var control is less than with uncontrolled V2G, but still very much greater than in the baseline case with existing smart meter voltages.

The voltage drop is approximately proportional to $PR + QX$, where P and Q are the active and reactive powers, and R and X are the cable resistance and reactance. Active power P is negative when there are net exports from the feeder cable, and in the absence of any reactive power Q the voltage drop will be negative, i.e. a voltage rise. For LV distribution cables $R \gg X$ and so high reactive (positive) values of reactive power are needed to introduce a useful difference to the voltage drop.

Volt-watt control is also effective at limiting the increase in cable capacity utilisation, as shown in Figure 26. The utilisation increases, as V2G exports can exploit the available voltage headroom, and the maximum feeder currents mostly remain within 100% of the rated capacity. However, a small proportion of feeder have currents higher than their rated capacity, demonstrating that a method based on voltage control does not guarantee compliance with thermal constraints.

For both the volt-var and volt-watt-var control methods the currents are significantly increased and are higher than with the V2G exports alone. This suggest that an additional control method is needed to ensure that the currents associated with reactive power, created to address the concern with voltage ranges, do not create a further problem with thermal capacity limits.

Figure 27 shows the corresponding probability of over-loads for the control methods, again showing that this risk is increased by volt-watt and volt-watt-var control, and is greater than with uncontrolled V2G. Volt-watt control is much more successful although there are still a few feeders where overloads are highly likely.

The reductions in extreme voltages with volt-watt and volt-watt-var control are achieved by limiting the exported power, as shown in Figure 28. In this figure, only the volt-watt and volt-watt-var control methods are shown. In the baseline case, the exports are of course zero, and with uncontrolled V2G and with volt-var control the exports are the full rated 3.7 kW.

Exports are reduced less for the volt-watt-var control than for volt-watt alone, although the volt-watt-var control is also slightly less effective in mitigating voltage rise.

On average, with volt-watt control, 40% of customers have no constraint, but the remaining 60% have between zero and 100% reduction in their exports. Mean export powers are higher with volt-watt-var control than with volt-watt control alone.

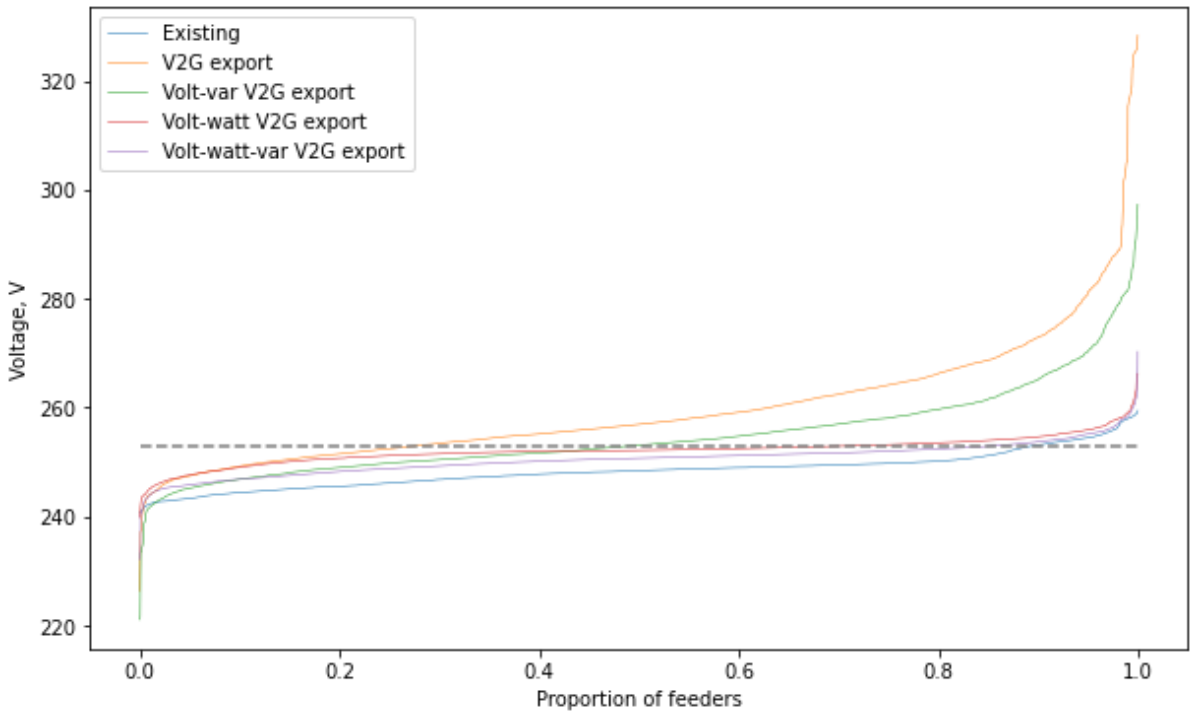

Figure 24: Maximum voltages with controlled V2G exports, volt-watt ramp 252 V to 253 V

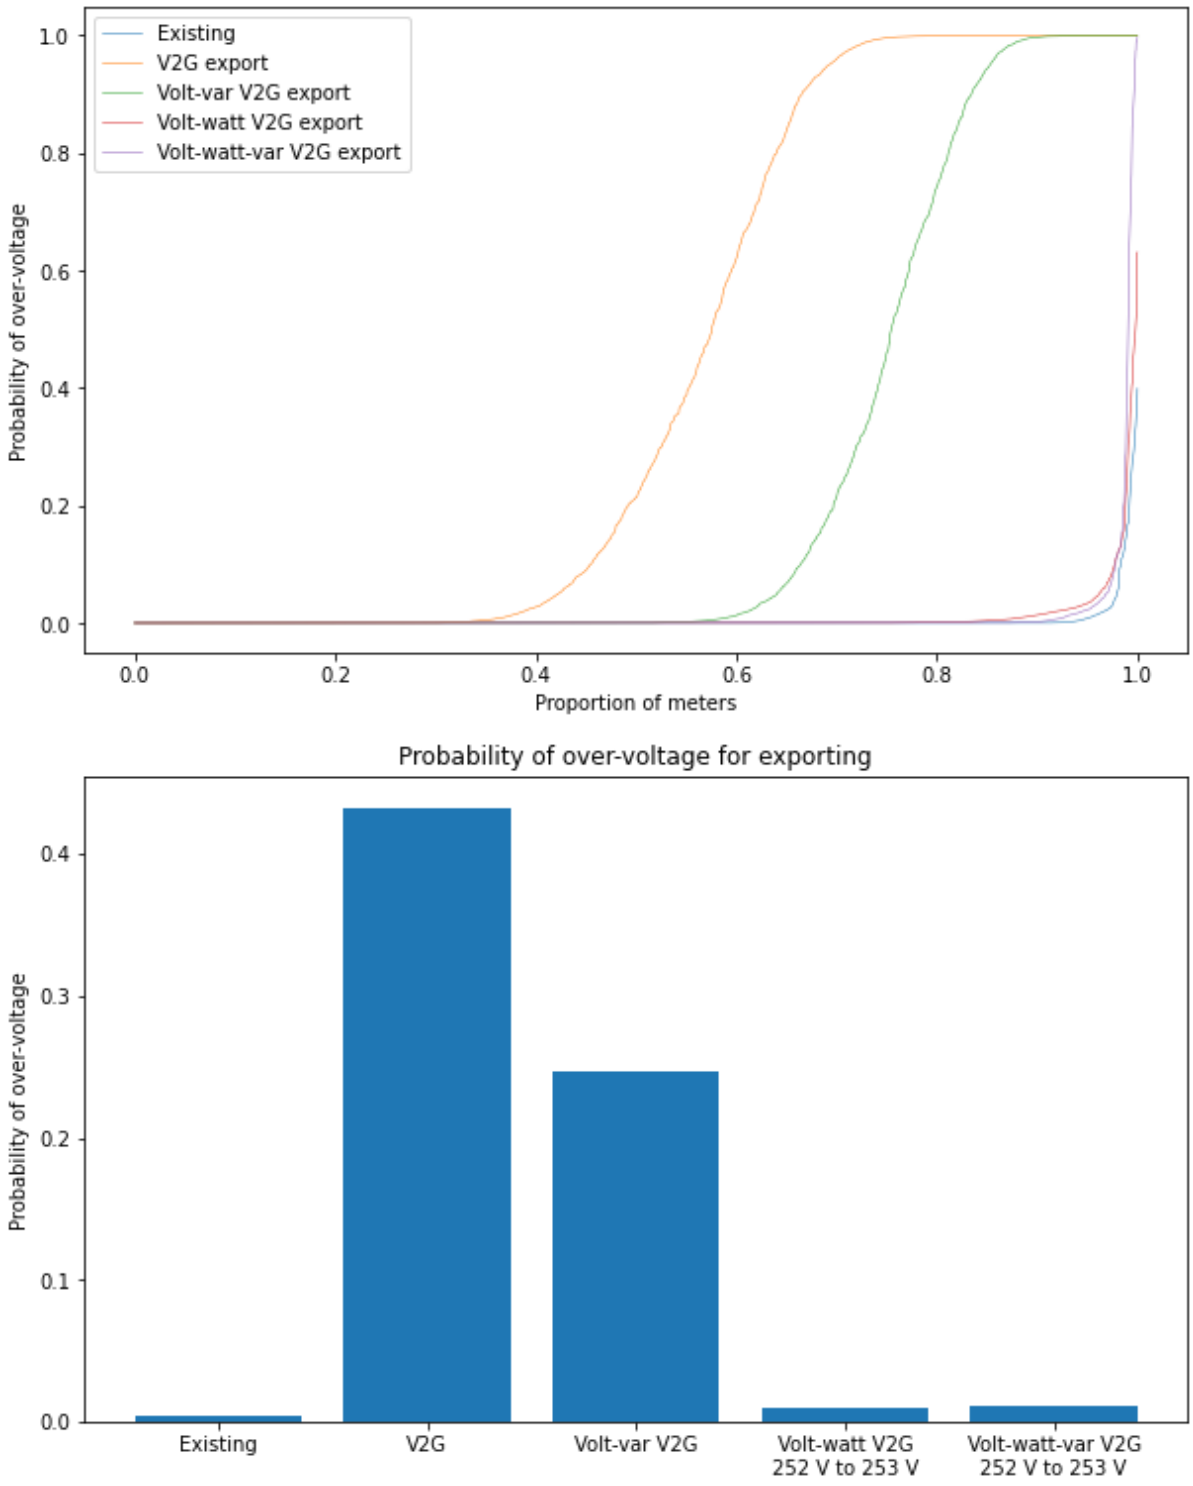

Figure 25: Probability of over-voltage with controlled V2G exports, volt-watt ramp 252 V to 253 V

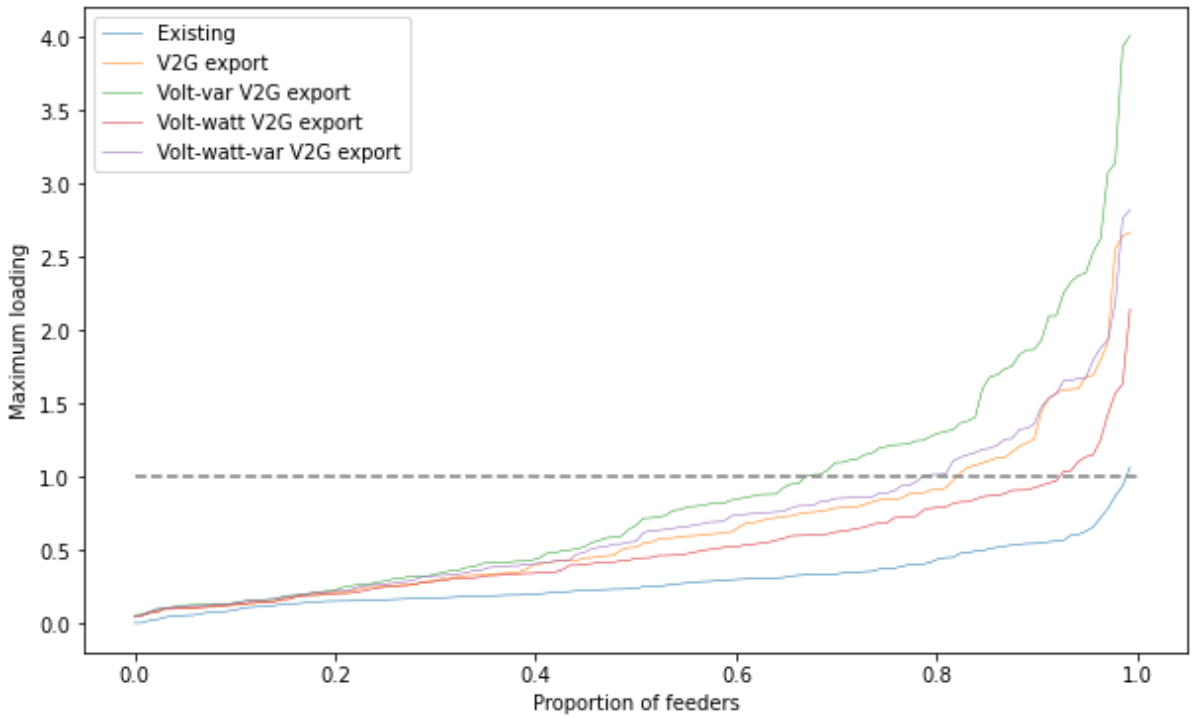

Figure 26: Maximum cable utilisation with controlled V2G exports, volt-watt ramp 252 V to 253 V

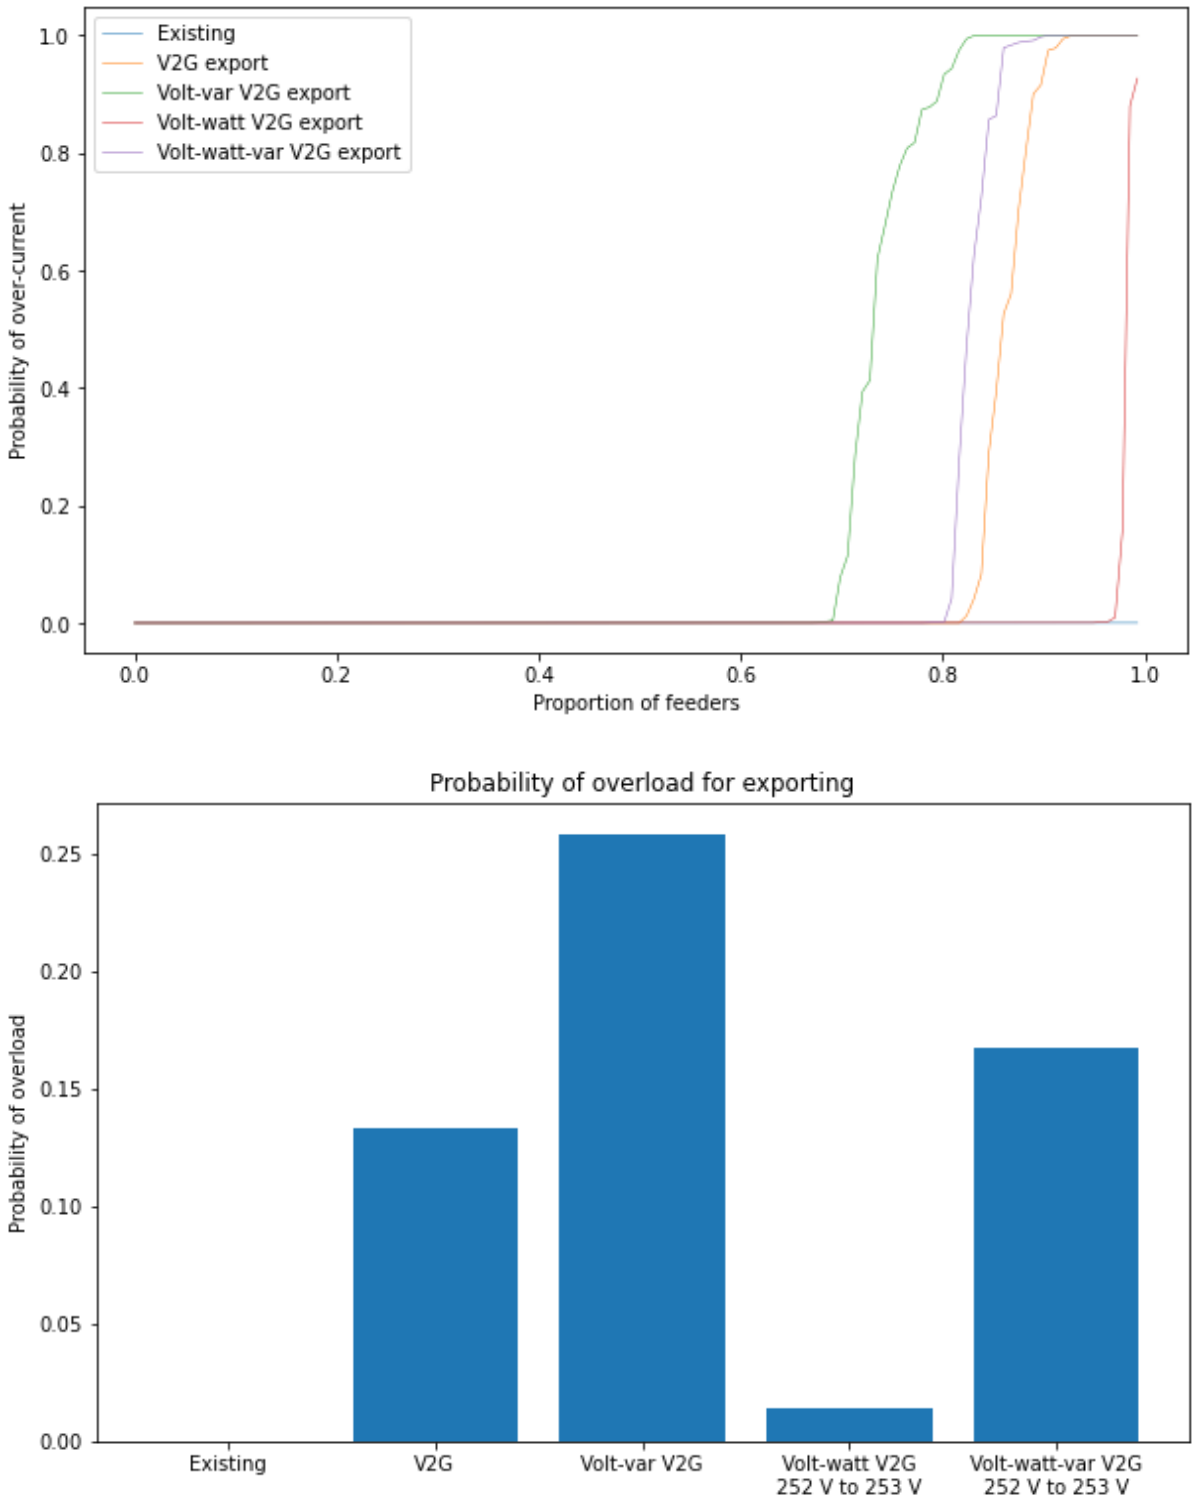

Figure 27: Probability of over-load with controlled V2G exports, volt-watt ramp 252 V to 253 V

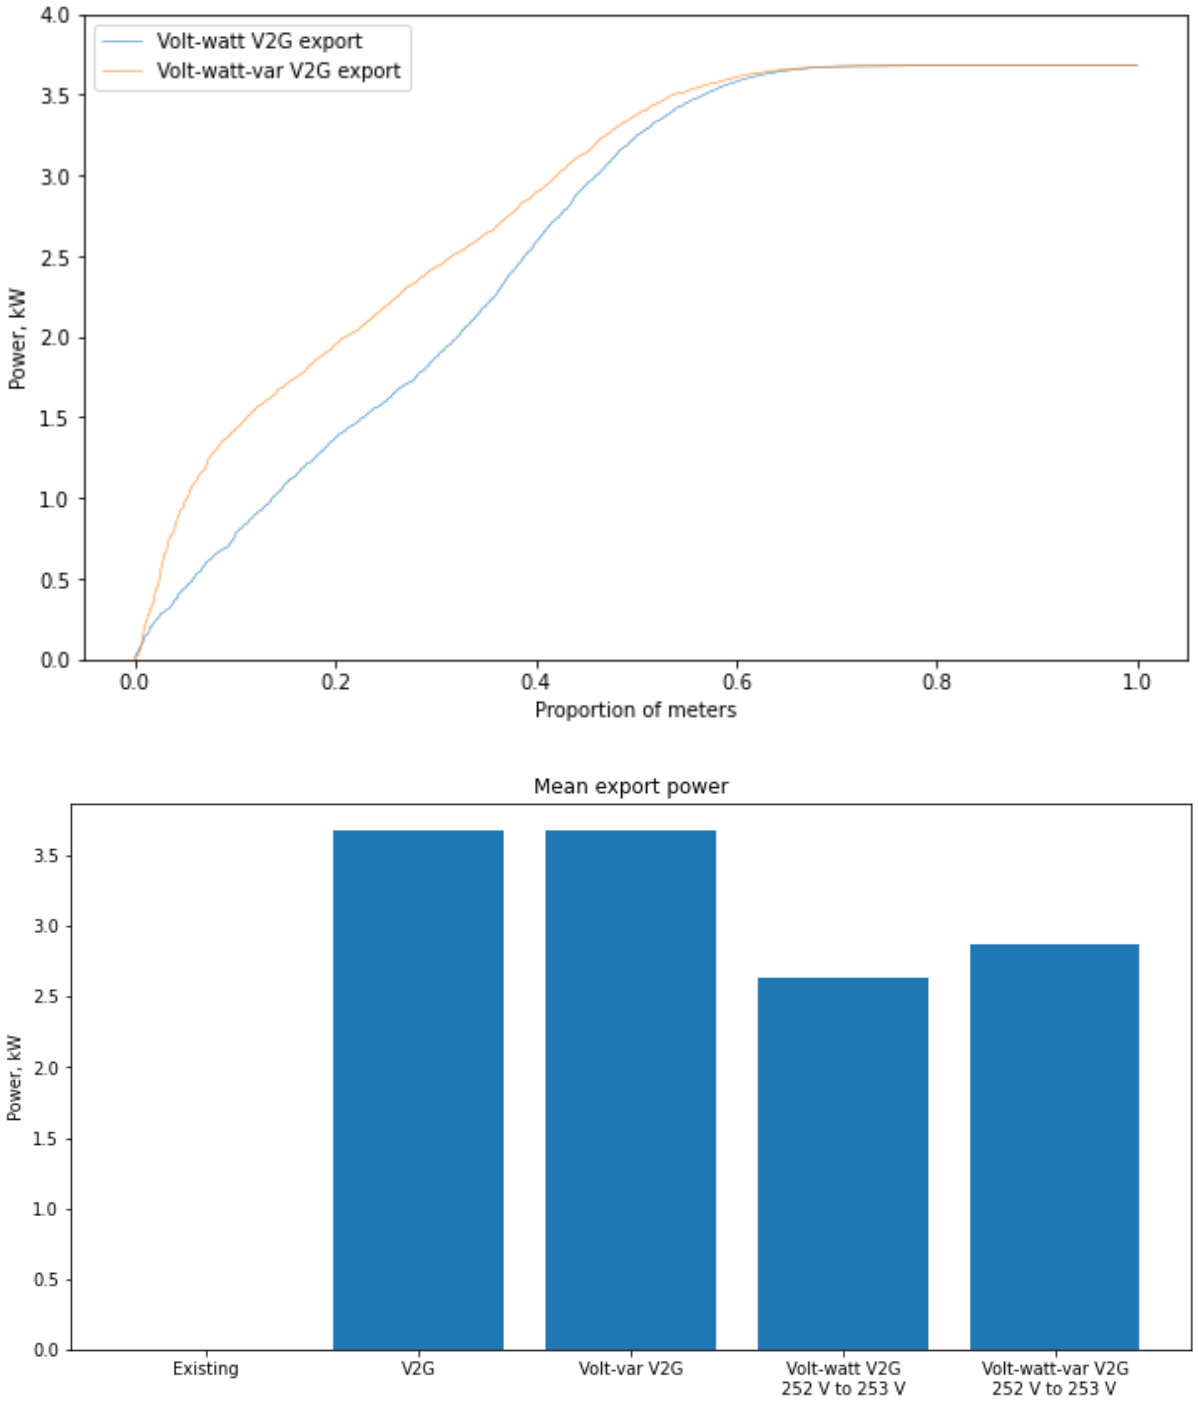

Figure 28: Mean V2G export power, volt-watt ramp between 252 V and 253 V

The impact on customers has been found to be unevenly distributed. Figure 29 shows the mean export power with additional traces plotted showing the CDFs for the customers that have the nearest distance on the feeder to the substation and for those that have the furthest distance. These CDF traces are less smooth as there are fewer meters included, but they show that customers that are furthest away from the substation are constrained more than those who are nearer.

These plots use results from a subset of feeders where there are 20 or more customers as the curves would be skewed by including feeders with very few customers. In an extreme case, if there is only one customer on a feeder, the customers at maximum and minimum distance from the substation would be the same. More generally, when data from all the customers is plotted, the points from each feeder contribute to the cumulative frequency distribution in proportion to the number of customers on

the feeder. However, when the data from the customers at maximum and minimum distance is plotted, there is on data point per feeder. Data points from feeders with fewer customers are then given greater weighting. Since the constraints on feeders with fewer customers have a different statistical distribution, this change in the weighting for each data point needs to be taken into account when results are compared. The restriction that customers are only included in the plot from feeders with 20 or more customers reduces this disparity.

It might be expected that customers towards the remote ends of feeders will unfairly be constrained more since they experience the greatest voltage rise. This is confirmed by the plots, but these customers are disadvantaged, relative to the full set of customers, by less than the advantage to the customers who are nearest to the substation. Rather than characterising the fairness such that the more distance customers are disproportionately disadvantaged and less able to earn income from V2G, it is more the case that customers near to the substation are disproportionately advantaged. For the cases modelled here, the impacts of V2G are substantial and it is necessary for many customers to reduce their active power so that voltages are maintained with limits. Those customers at the ends of feeders will almost certainly be included in the constraints, but they have no more exports to lose than the customers elsewhere on the feeder. Those who are nearer to the substation will have lower voltages and so reduce their active power less, and may also be fortunate that the more numerous customers elsewhere on the feeder take their control action first such that no further reductions are necessary.

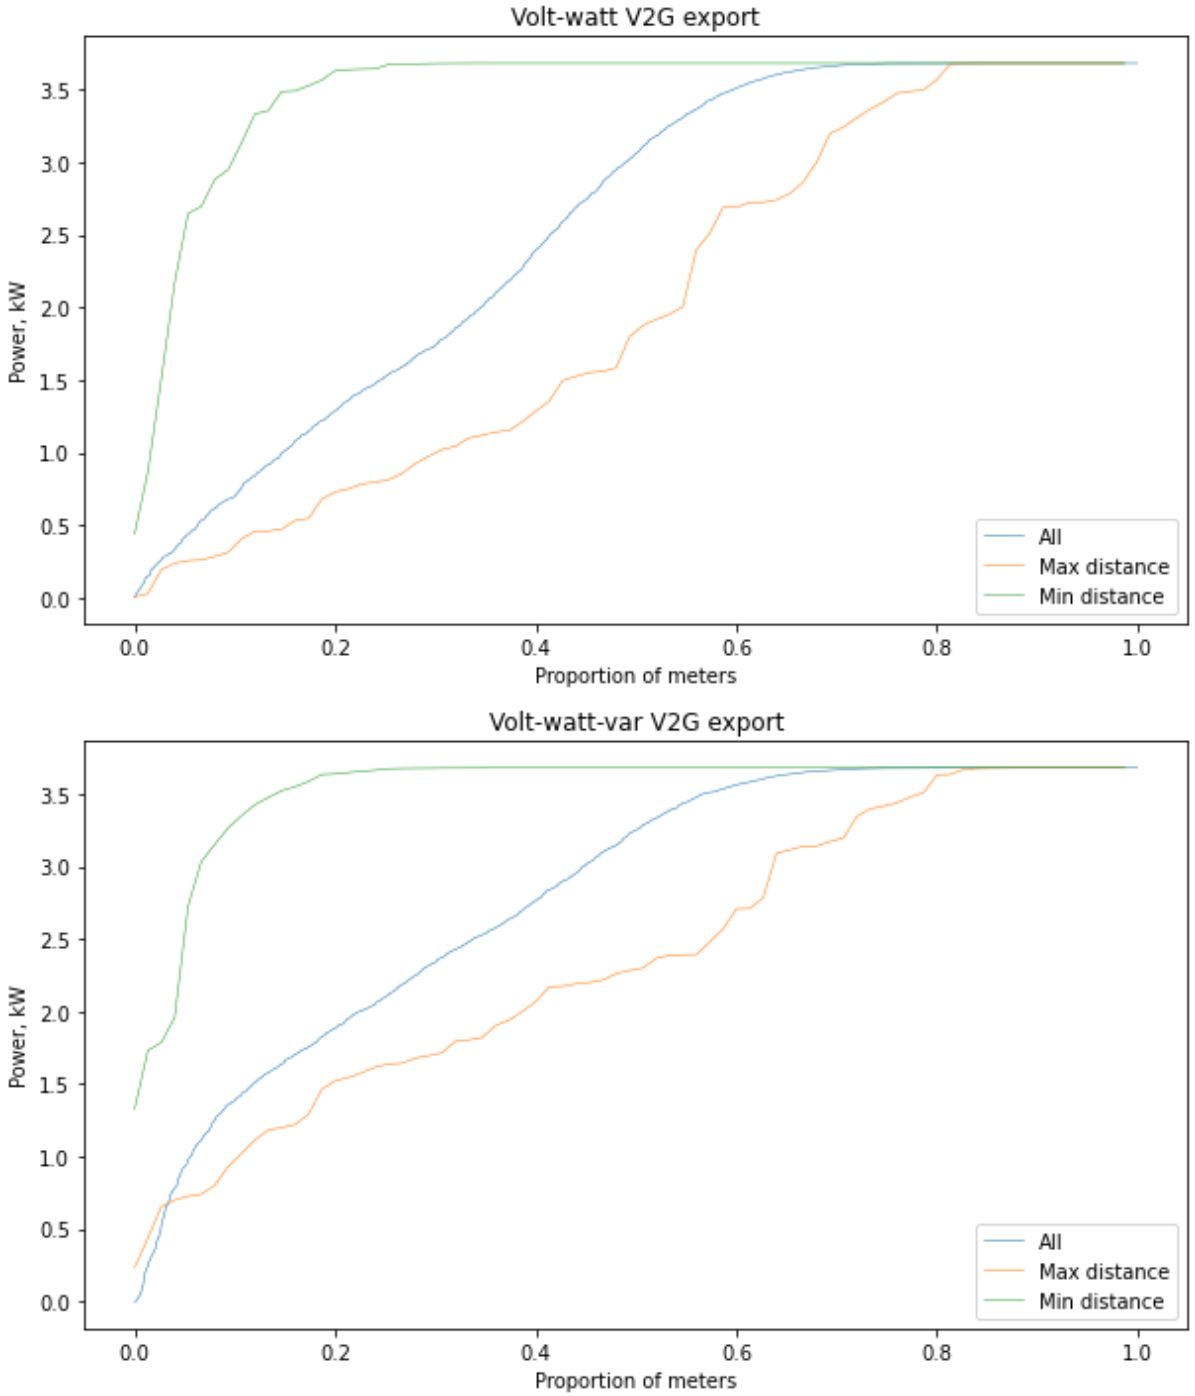

Figure 29 Mean V2G export power, minimum and maximum distances, volt-watt ramp 252 V to 253 V

With increased utilisation of the feeder cables for V2G exports, it would be expected that losses will increase. Exported power from V2G will offset existing demand imports but, for the simulation models considered here where all the V2G customers export at the same time, the exports will be greater in magnitude than the demand and so feeder cable losses are likely to increase instead of being reduced.

A slightly different approach is needed to calculate losses from the simulation results than has been used for the results shown above. These previous results only need to consider the impacts while the exports are active, and so the simulations have modelled a scenario where V2G exports occur all the time, so that the probability of over-voltage, or the mean export power, can be calculated for any time.

However, where losses are considered, it is more important to understand the mean losses over time, rather than only when the V2G customers are exporting. Since the proportion of time when exports are required has not been defined, results have been calculated for between zero and six half-hour periods in a day (ie. up to 3 hours per day, on average, of exports). These results select half-hours with a uniform distribution over the available smart meter data samples, such that there is no weighting to any specific time of day for exports.

The calculation of losses is presented here only for the first branch in the feeder, in effect the cable route from the substation busbar to the first tee-junction or service cable connection. This is a necessary restriction since the simulation model calculates a super-position of the V2G demands over the currents due to the existing customers, and this is only known from smart meter data where the demand from the full set of customers on the feeder has been aggregated. Although the disaggregated currents due to V2G demands are known throughout the feeder, the currents due to the existing demand are not. The losses in the first feeder branch are generally higher, per unit length, than losses elsewhere on the feeder, but there are typically many downstream branches beyond the first tee point and so the losses shown here are only a small proportion of the total. As a simple approximation, it might be assumed that total feeder losses increase in a similar proportion to the losses in the first feeder branch, on the basis that the majority of customers are assumed to have V2G exports and so the aggregation statistics may be similar to the existing demands.

Figure 30 shows the mean losses in the first feeder branch, calculated over all feeders, for V2G exports without any voltage-based control. There is a linear increase in losses from the case with the existing smart meter data, rising as a greater number of half-hours per day are assumed for V2G exports. In practice, there would be additional losses due to the recharging following the export periods, and this effect is not included here. For exports alone, losses increase from an average of 17 W per feeder, to around 70 W, assuming that a worst-case scenario (for losses) in which all V2G customers export at the same time for 3 hours per day.

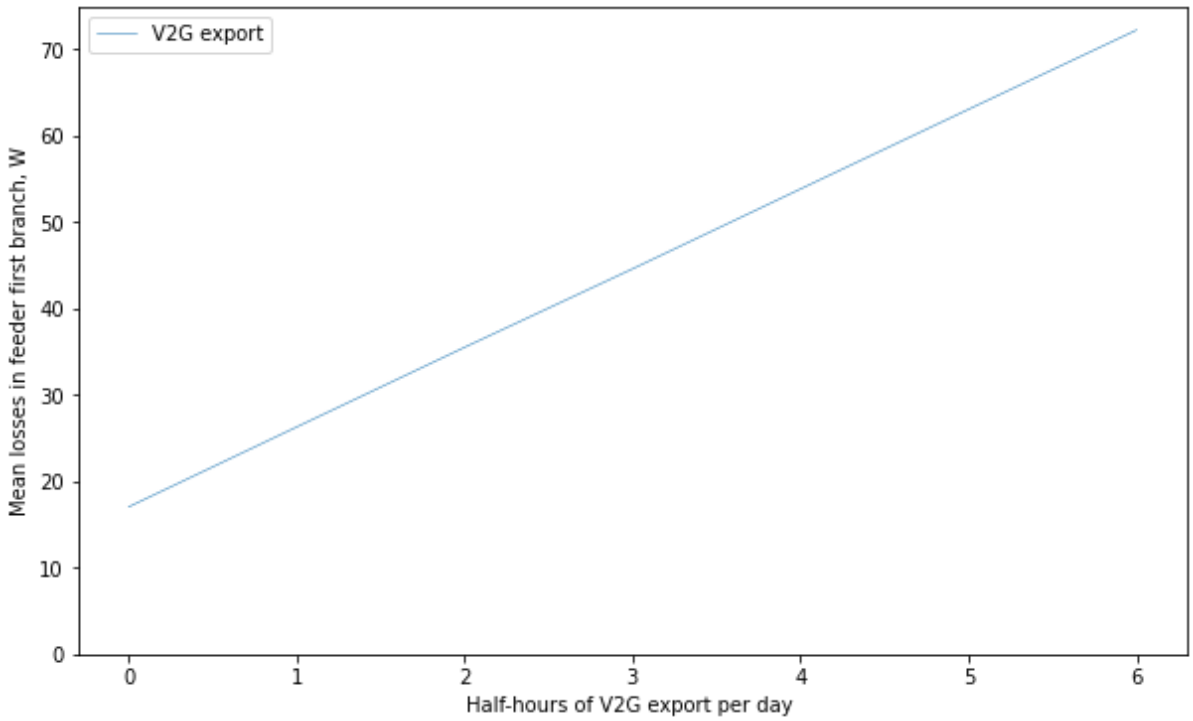

Figure 30 Mean losses in first feeder branch for V2G export

Figure 31 shows how these worst-case losses, again assuming exports for 3 hours per day, are impacted by the voltage-based control strategies. With volt-var control, losses nearly double relative to V2G without control, due to the increased current amplitudes associated with the reactive power

transfers. With volt-watt control, here assuming control thresholds ramping power to zero between 252 V and 253 V, losses are roughly halved. However, as noted above, this also significantly reduces the power exports achieved. Using the combined volt-watt-var control and these same thresholds, losses are similar to the case with no control, although exports are also reduced as shown in Figure 28.

While V2G exports clearly increase losses, and this may be considered a negative impact, this impact needs to be balanced against the value to be gained by providing power exports to balance the grid.

If a typical EV battery has a 95% efficiency when charging or discharging, and therefore a roundtrip efficiency of around 90%, then losses of 10% of the charging power are associated with V2G exports. This can be compared to the overall losses of the distribution network, typically around 7% of delivered power, although this figure is inevitably higher at times of peak demand (since cable losses increase with the square of power). While distribution losses are clearly a significant concern, they are therefore of a comparable order of magnitude to the losses inherent in supplying power from V2G exports. These losses may be reduced if the required exports were instead provided by grid-scale batteries connected at higher voltage levels, but with the additional capital cost (and embodied energy cost) associated with duplicating the storage capacity that could already by available from EVs. Losses may also be reduced if the V2G exports are required to meet domestic demand that is distributed throughout the network and therefore co-located with the EV chargers.

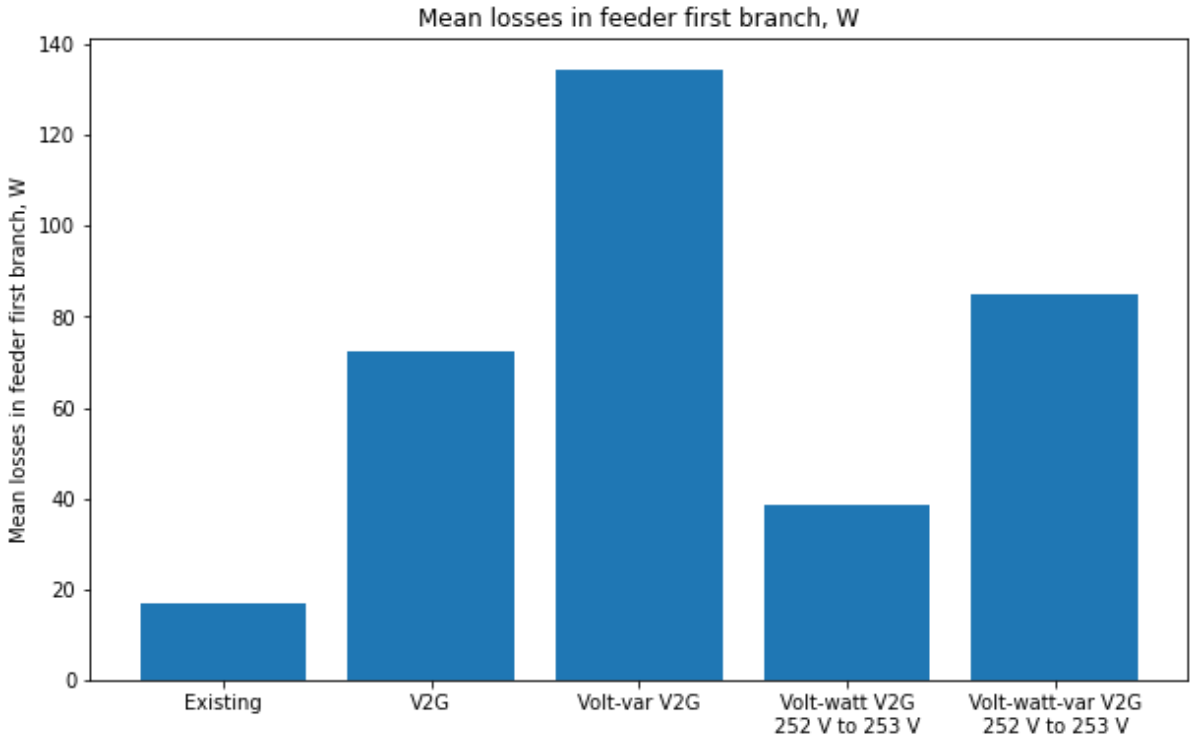

Figure 31 Mean losses in first feeder branch for V2G export 3 hours per day

5.2 Control algorithms for voltage drop

There is less scope for mitigation of voltage drop during recharging as the volt-watt control method defined so far does not apply in this case.

Figure 32 shows the impact of volt-var control on minimum voltages, reducing the worst-case voltage drop such that the minimum voltages are higher. For this selected test case, the proportion of feeders that remain above the minimum limit happens to be similar with volt-var control as without.

Figure 34 shows the corresponding impact on cable utilisation, where currents are slightly greater with the volt-var control than without. As seen above, the probability of over-load is around 50% and not diminished by the control techniques.

As noted above, further control methods would be needed to avoid problems with low voltages and to mitigate the associated high currents. An increase in the supply of reactive power would assist in managing the voltage ranges but would increase the thermal constraints, and so either a control of the active power, or a diversification in the timing of recharging, appear to be necessary.

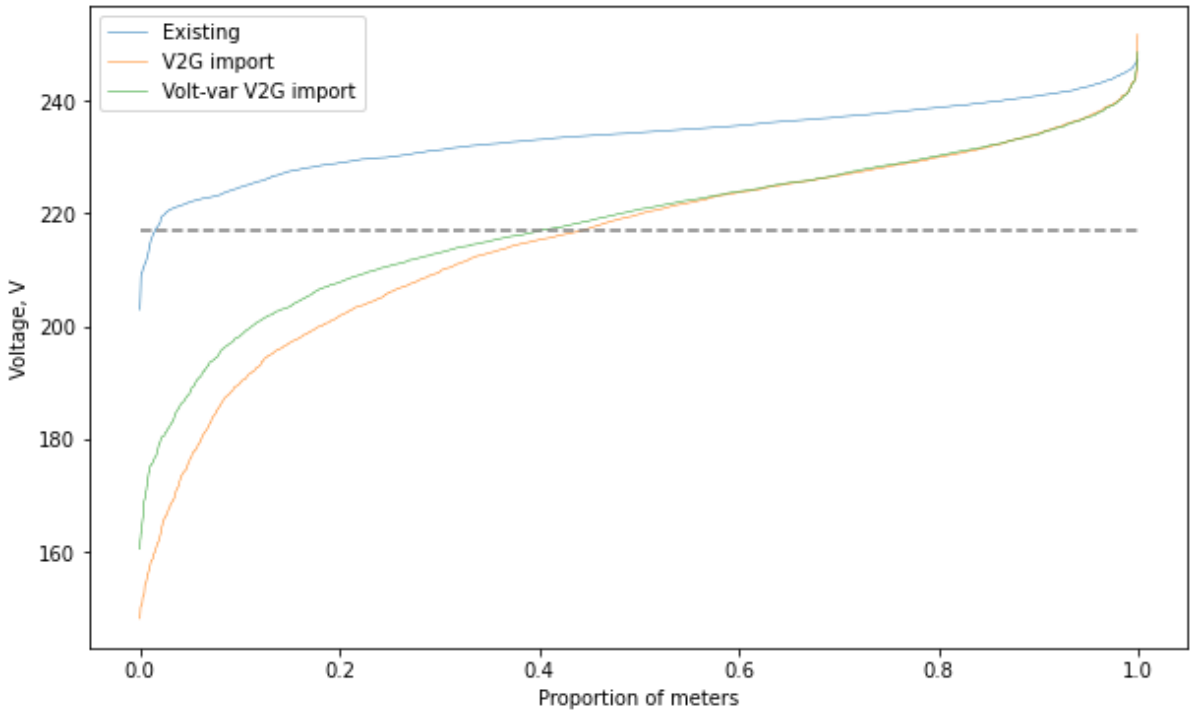

Figure 32: Minimum voltages with controlled V2G imports

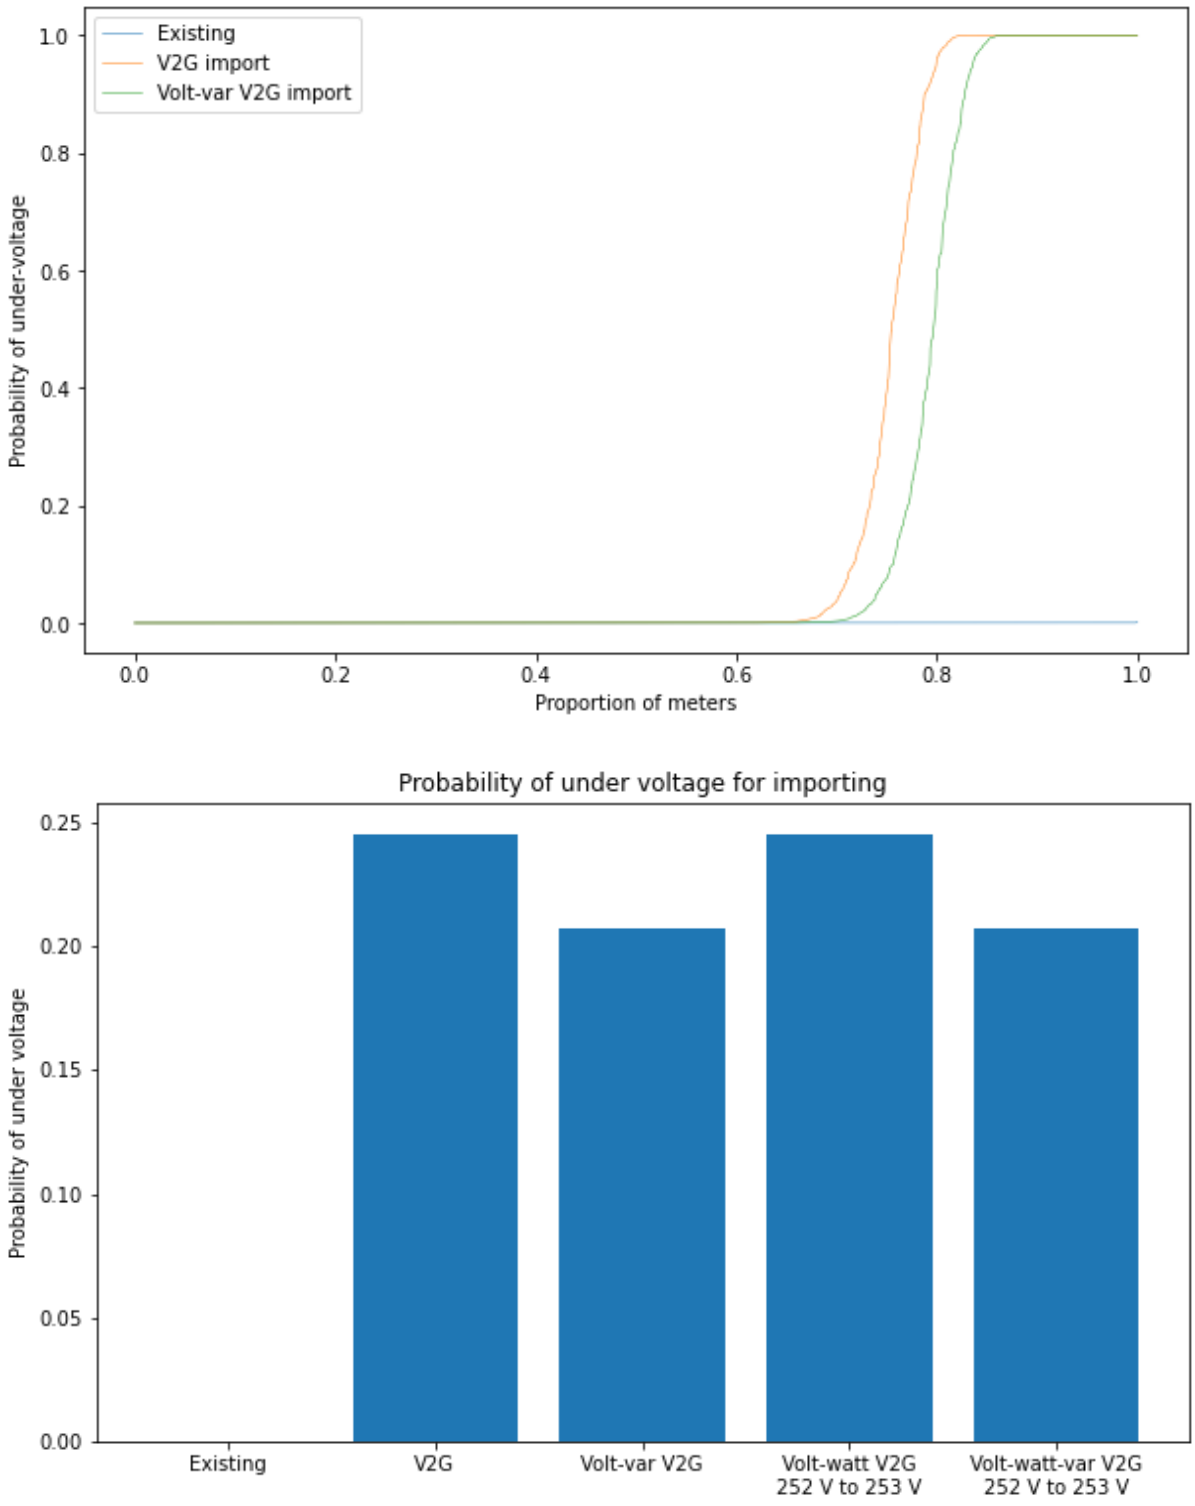

Figure 33: Probability of under-voltage with controlled V2G imports

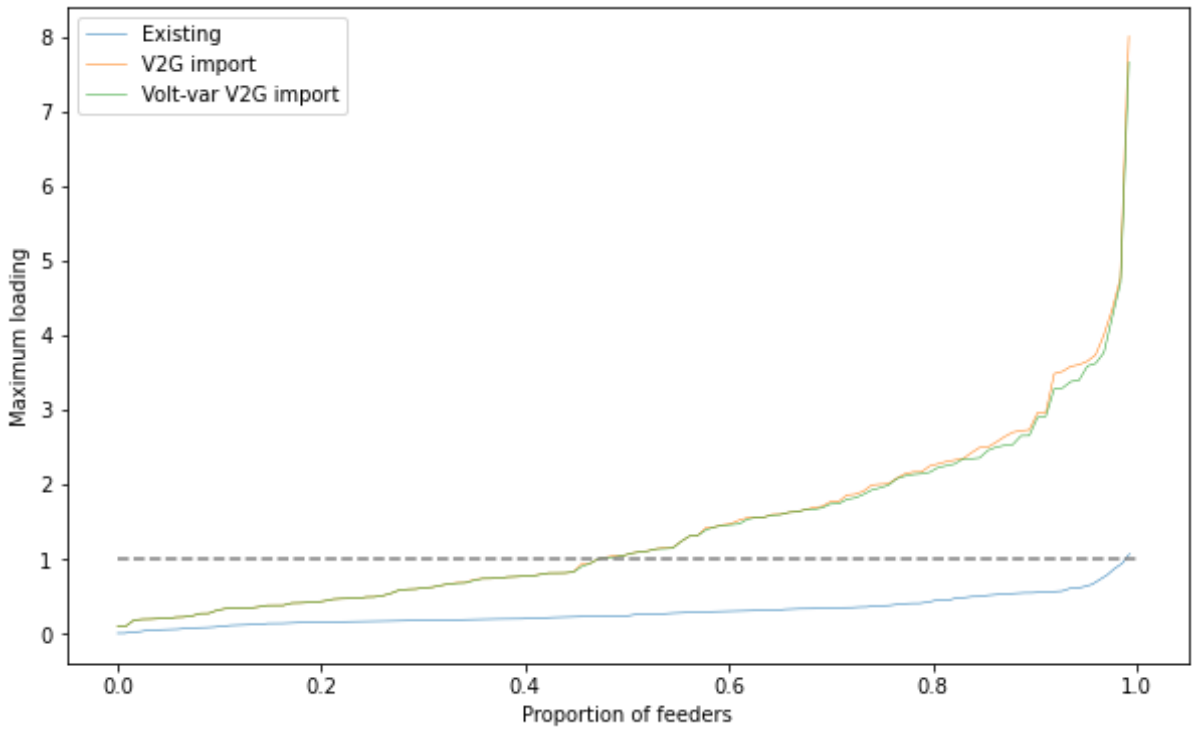

Figure 34: Cable utilisation with controlled V2G imports

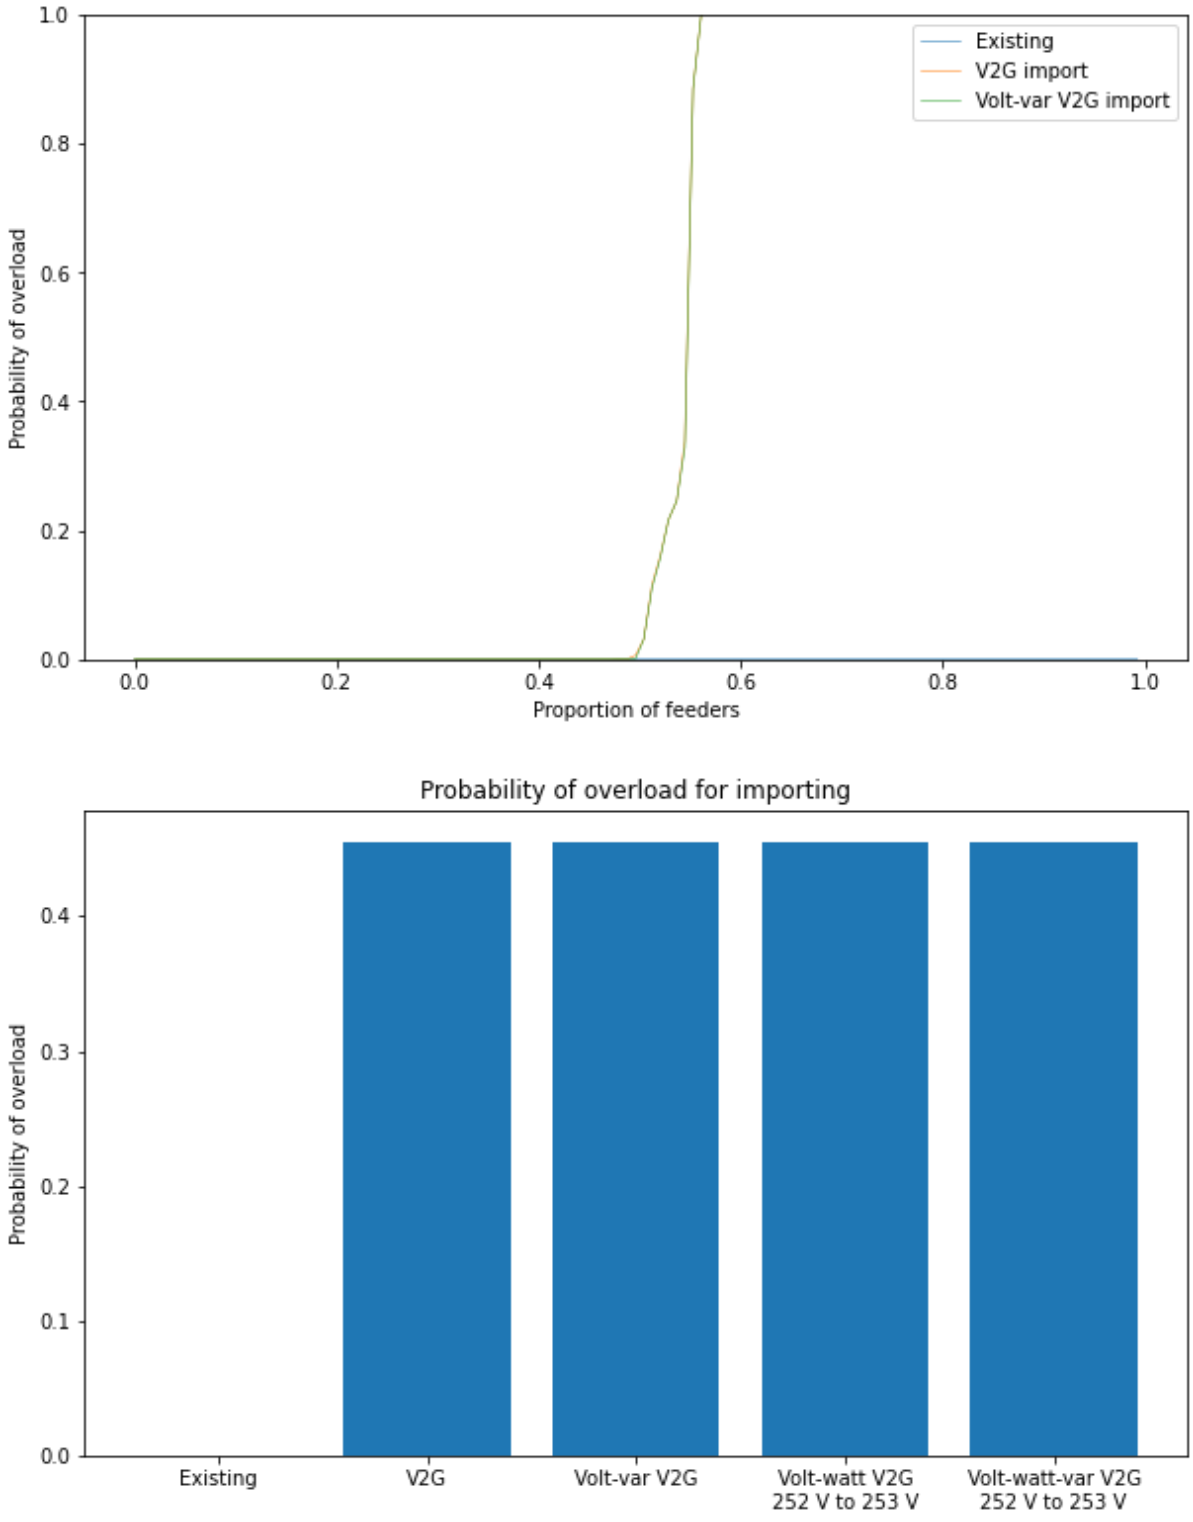

Figure 35: Probability of over-load with controlled V2G imports

5.3 Fixed threshold variations

This section presents results where the volt-watt thresholds are varied from a default range of 252 V to 253 V as in the previous section.

Figure 36 to Figure 40 show results for volt-watt and volt-watt-var control with threshold ranges of

- 240 V to 245 V,
- 245 V to 250 V,
- 250 V to 255 V,
- 255 V to 260 V,

As would be expected the lower thresholds reduce the maximum voltages more severely and the higher thresholds reduce voltages less. A lower threshold provides a higher certainty that the upper voltage limit will not be exceeded, whereas higher thresholds allow a greater number of feeders to have no constraints applied.

Lower volt-watt thresholds also significantly reduce the probability of over-voltages. Setting thresholds higher the upper limit of 253 V not only allows higher maximum voltages but also allows them to occur more frequently.

The use of a narrower ramp range does not significantly affect the results in terms of the impact on the network, although will mean that customers have a distinct transition between either being constrained or not. Wider ramp ranges are expected to provide greater stability in the control system, with a reduced likelihood that V2G inverters will repeatedly switch on and off in response to small voltage differences. This real-time effect is not captured in the simulations presented here so the following results show only case with the 5 V ranges.

The variations in volt-watt thresholds have a corresponding impact on the cable utilisation, as in Figure 38. Although there are a few feeders with very high cable utilisation (and also very high voltage, as in Figure 36), results for any of the threshold ranges are mostly acceptable, since there is greater headroom for cable utilisation than there is for maximum voltage.

Not surprisingly, the improved compliance with voltage and cable utilisation limits for lower volt-watt control thresholds has an impact on the export power from V2G. Figure 40 shows how each 5 V step reduction in threshold settings has a significant increase in the expected export power constraints, with V2G largely switched off for the tightest threshold setting where the ramp range is between 235 V and 240 V.

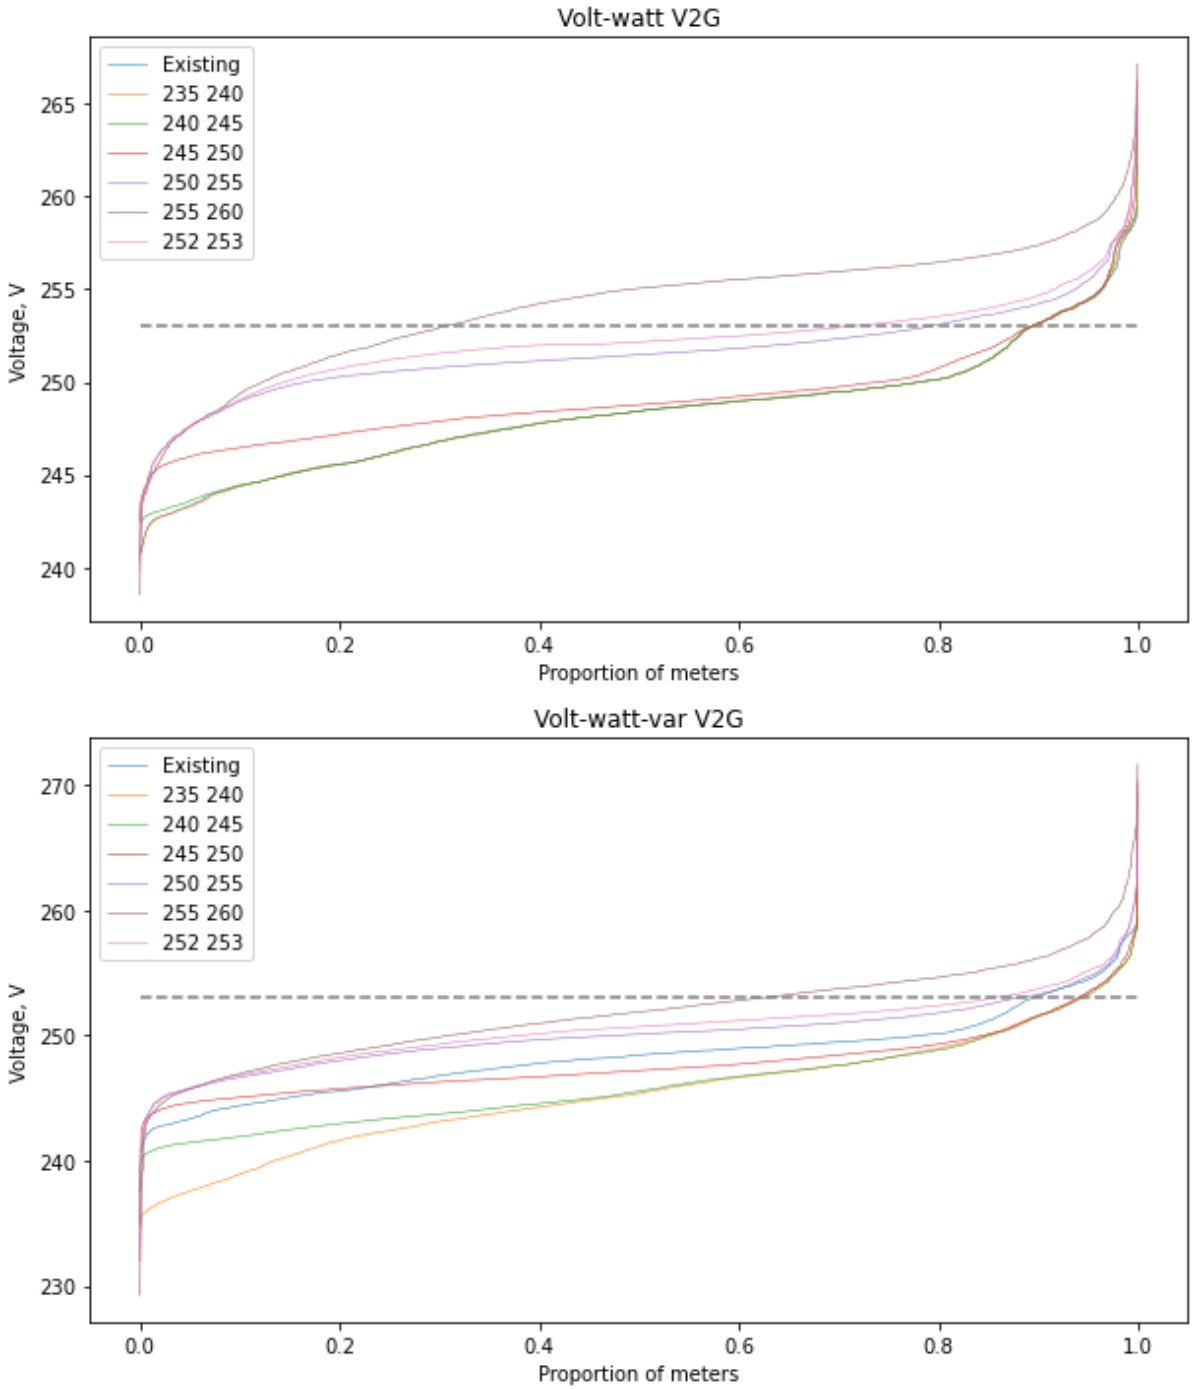

Figure 36: Maximum voltages with controlled V2G exports, varying ramp thresholds

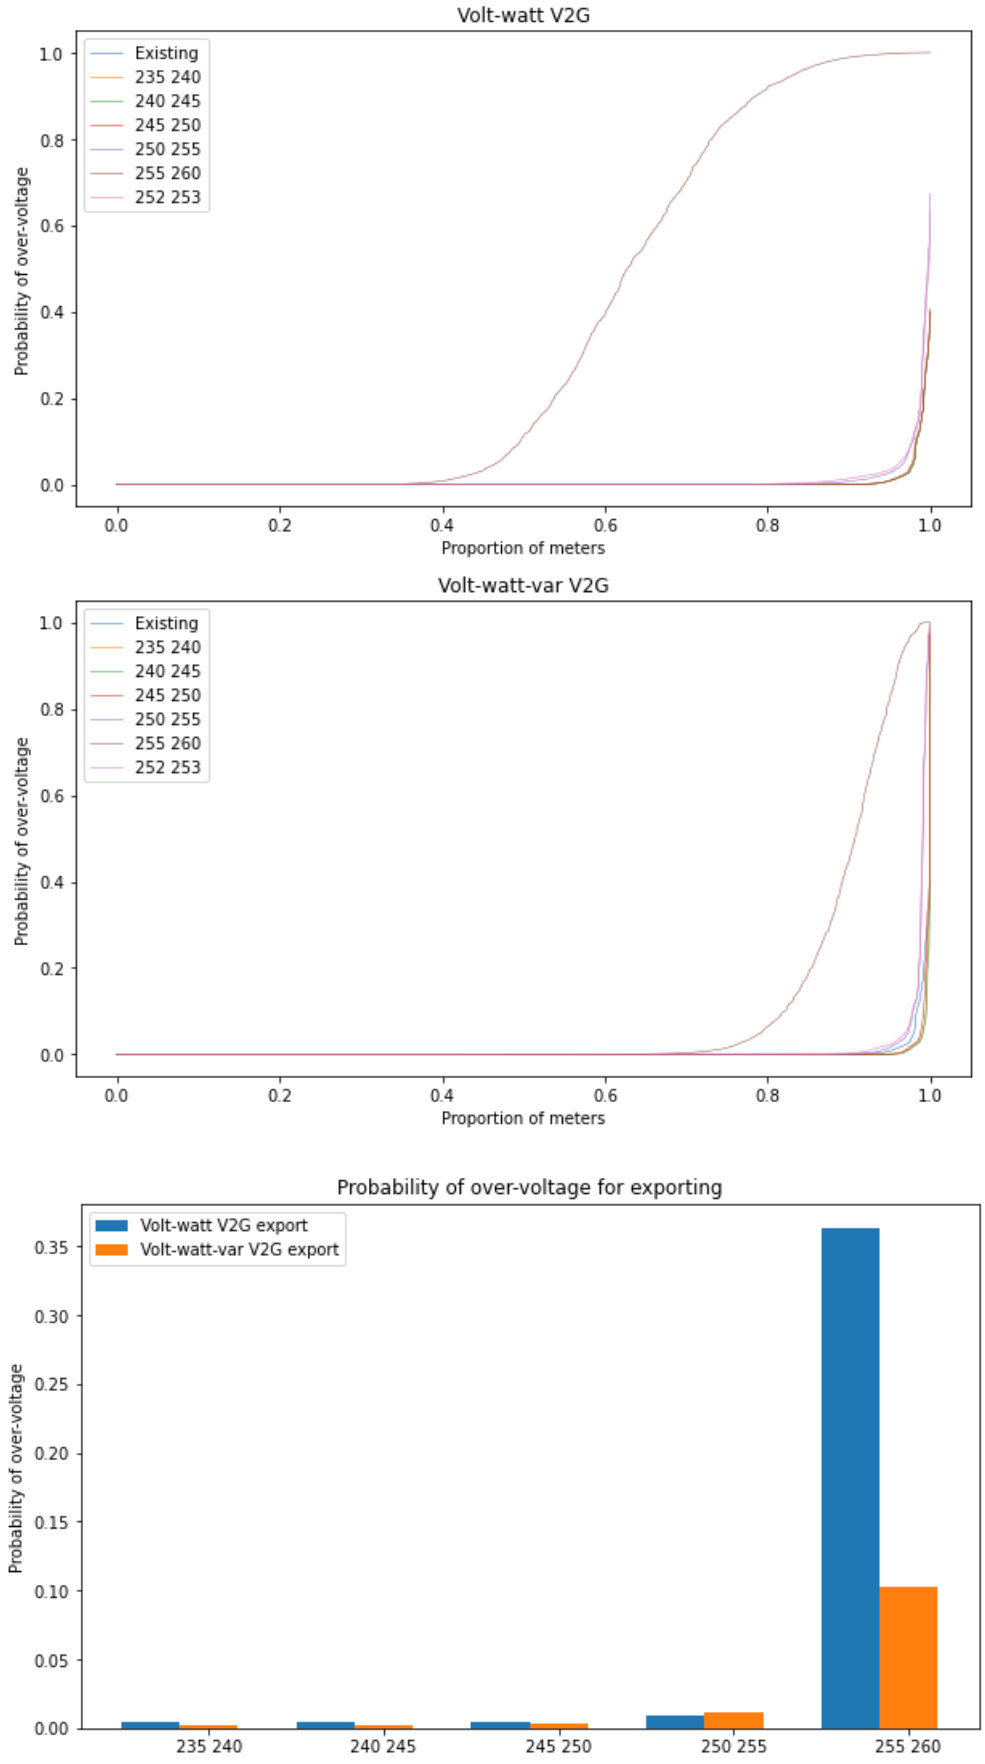

Figure 37: Probability of over-voltage with controlled V2G exports, varying ramp thresholds

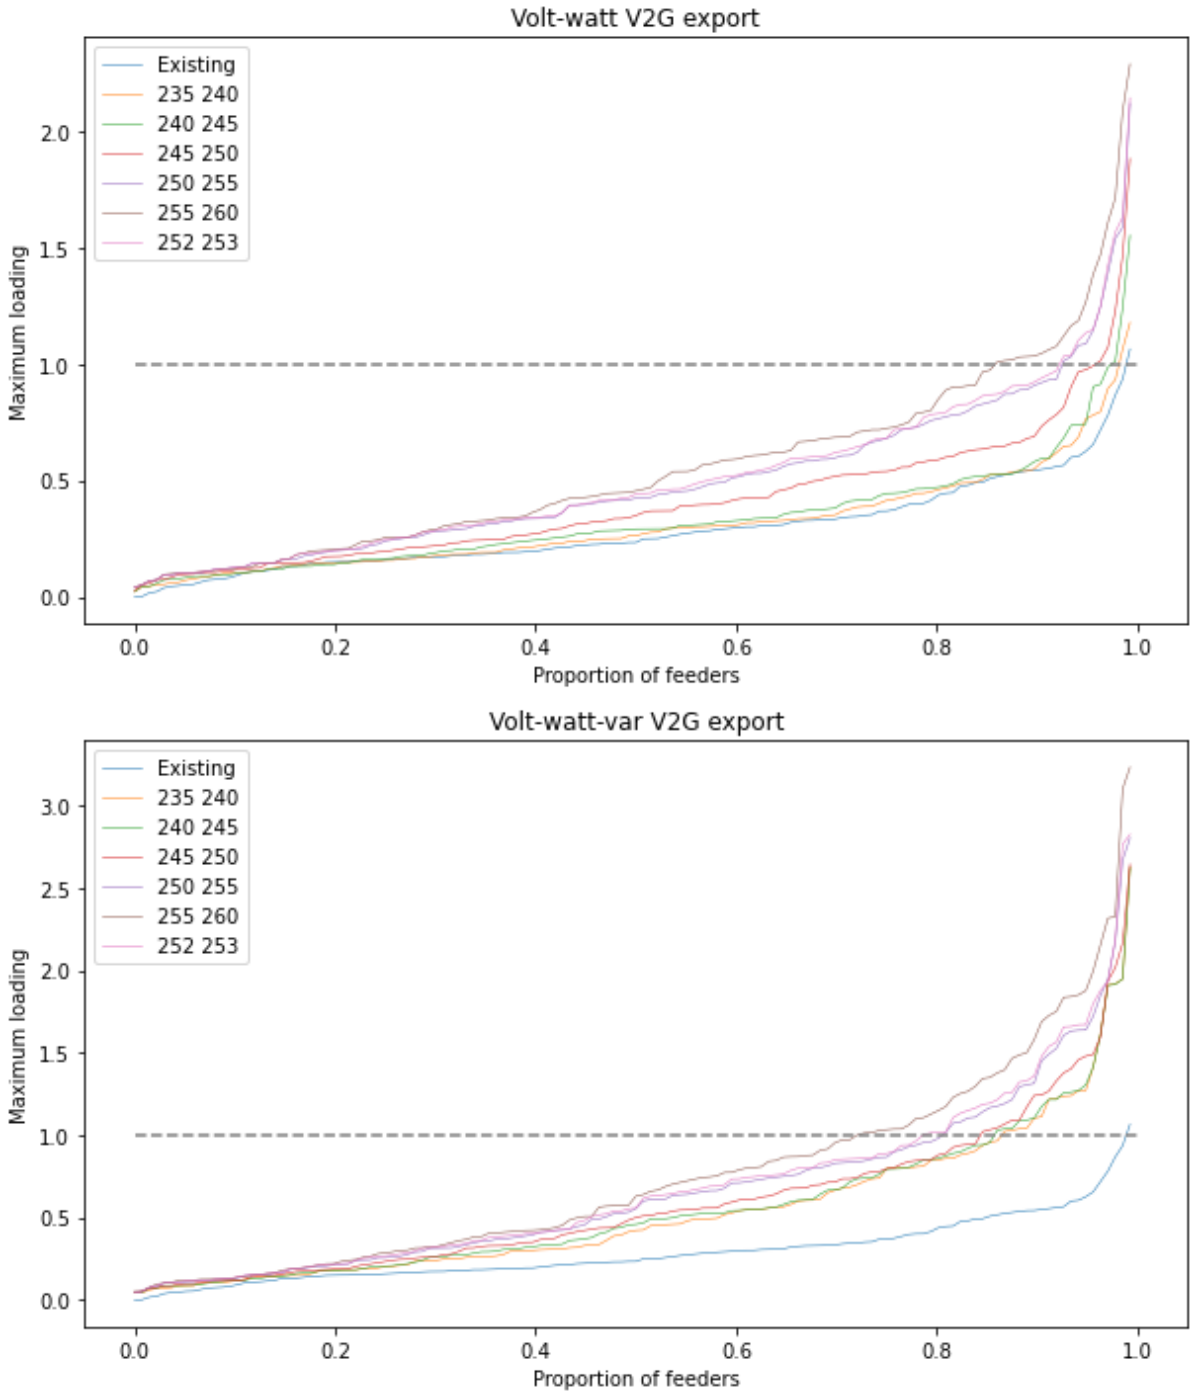

Figure 38: Maximum cable utilisation with controlled V2G exports, varying ramp thresholds

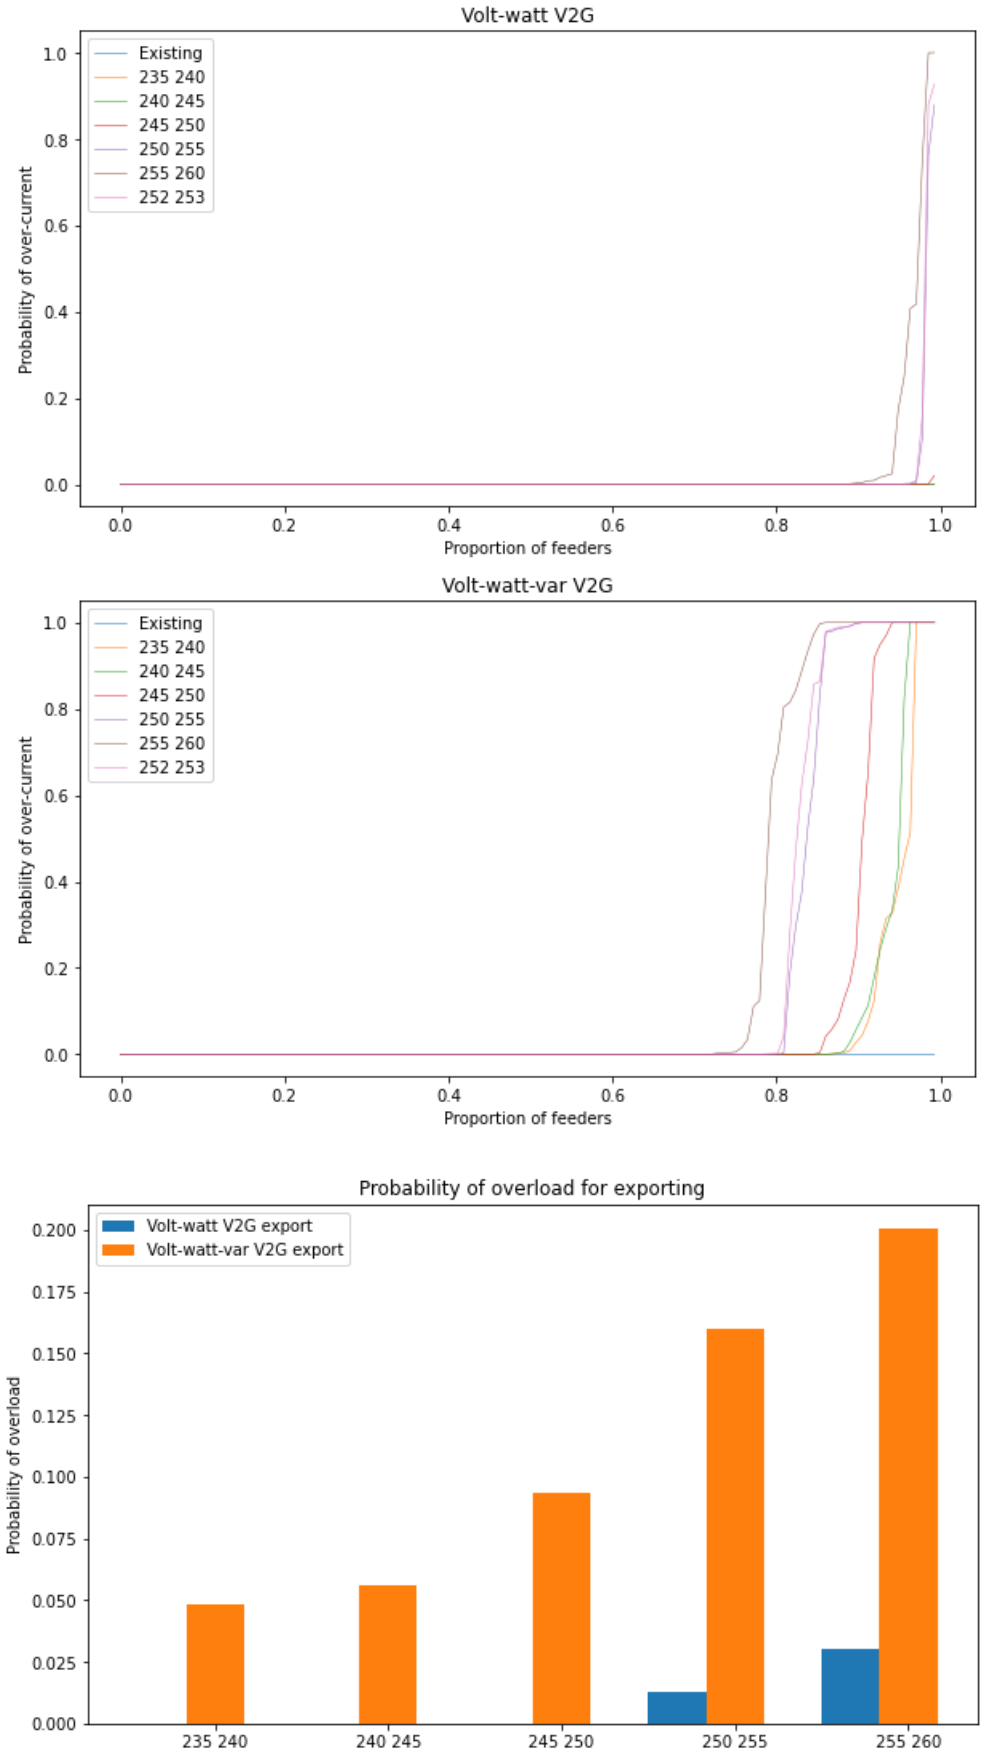

Figure 39: Probability of over-load with controlled V2G exports, varying ramp thresholds

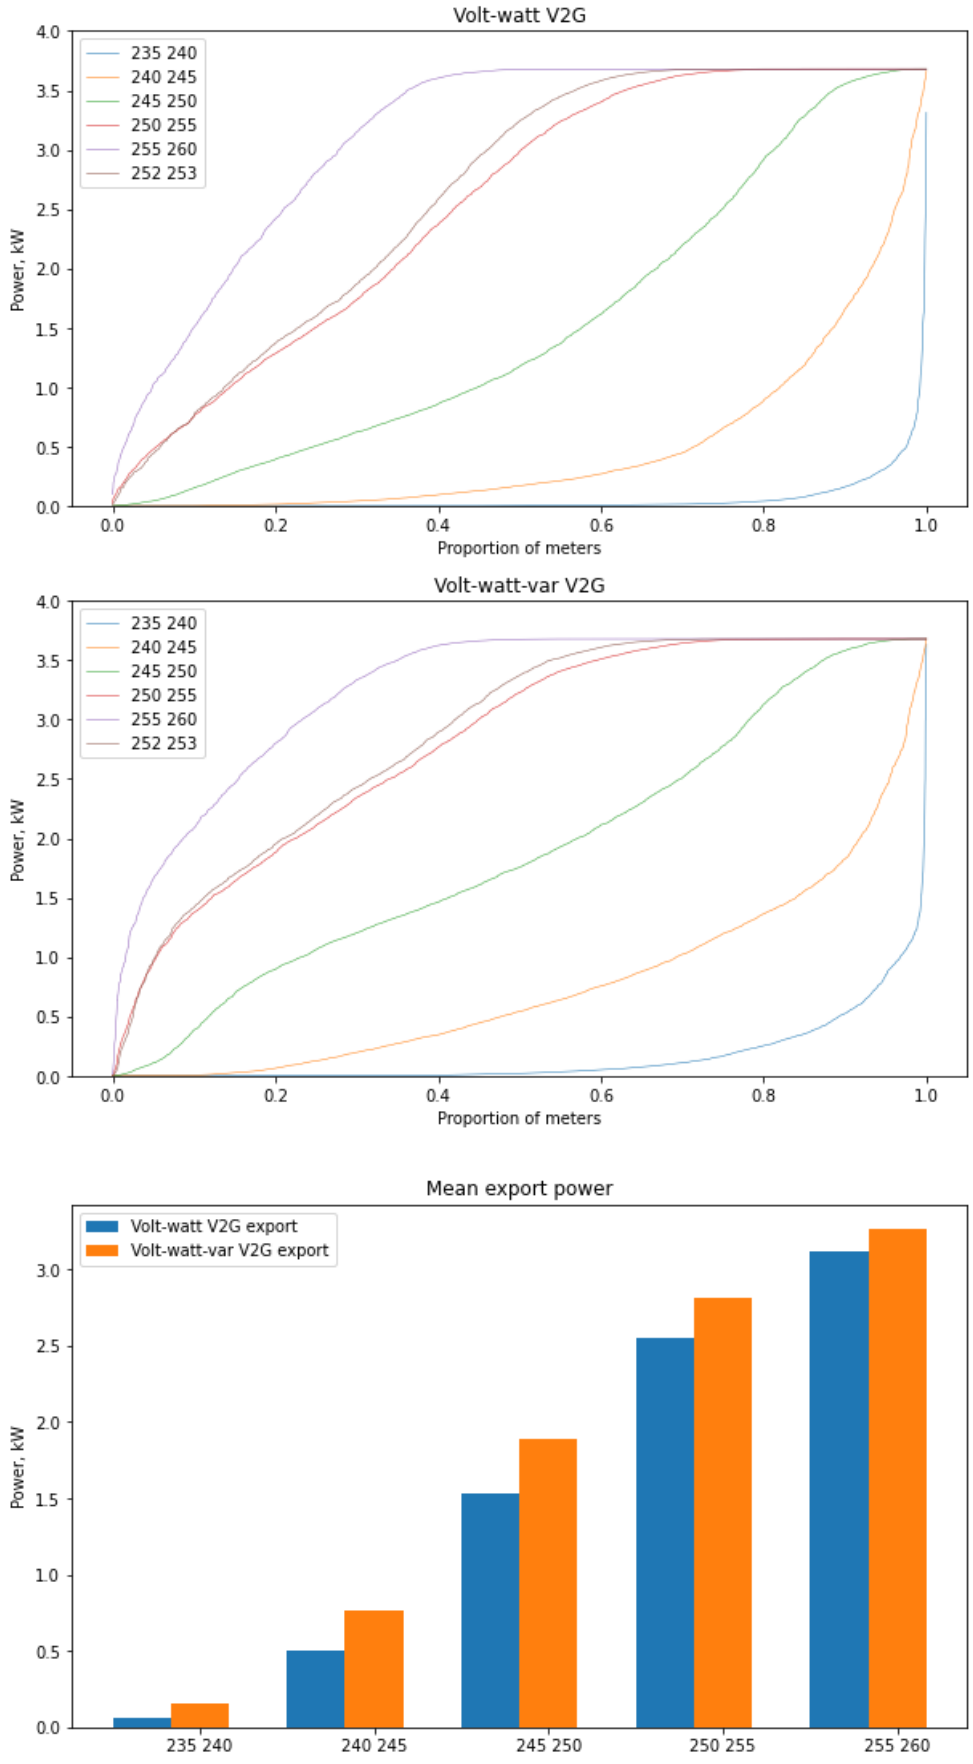

Figure 40: Mean V2G export power, varying ramp thresholds

5.4 Locally optimised thresholds

Results shown above use a common volt-watt threshold for all customers. This section considers an alternative approach where the optimal threshold range, selected from the ranges described in section 5.3, is determined separately for all the customers at each substation or on each LV feeder. This introduces a dynamic response to the headroom as it is assumed that this threshold would be selected on the basis of voltages and aggregated demand derived from previous smart meter data.

The optimisation method considers results with each of the 5 V threshold ranges shown in section 5.3 and excludes thresholds where either the maximum voltage after controls have been applied is above 253 V, or the cable utilisation is greater than unity. The selected threshold range is that highest range that will avoid the voltage and thermal constraints, while maximising the mean exported power from V2G. If no threshold is found to be valid, in effect where there is already no headroom, a default lowest threshold setting of 240 V to 245 V is used.

This process aims to find the minimum possible constraint to the total export power from each of the LV feeders at the substation, such that the maximum voltage and maximum cable utilisation on each feeder remain within voltage and thermal limits.

Where results are calculated per feeder, the threshold is selected only on the basis of the relevant feeder. Results for a threshold per substation select the lowest threshold of any feeder. A business-as-usual implementation would use a similar process for each, but with the group of customers selected accordingly. While the use of individual thresholds for each feeder would increase the complexity of the process in terms of the data handling, the requirement for smart meter data collection would be similar in both cases. An initial selection could determine which customers experience the greatest voltage rises or drops, such that only those meters would need to be included in the ongoing smart meter data collection.

Results for thresholds customised for each substation or each feeder are shown in Figure 41 to Figure 45. The maximum voltages for volt-watt control shown in Figure 41 are no higher than with the existing smart meter data where this is already above the limit, but can be higher where there is headroom available. With the default option of thresholds between 240 V and 245 V, there is virtually no increase in voltages above the existing baseline. With volt-watt-var control the maximum voltages can be lower than the existing baseline due to the reductions provided by the imported reactive power.

Figure 42 shows that the probability of over-voltages occurring is similar in all cases, although a very few customers would still not avoid having voltages over the upper limit. The probability of over-voltage shown here is very low because due to the threshold selection process which attempts to avoid any over-voltage occurring, other than those that cannot be avoided due to the voltage variations in the existing smart meter data.

The maximum cable utilisation with volt-watt control in Figure 43 is mostly within limits but there are a few feeders where overloads would occur with any of the volt-watt control methods, and this likelihood is increased relative to the baseline case. For these feeders, an even lower default threshold range would be needed in order to ensure that the current remains within limits, even though there may be no concerns with the voltage. With volt-watt-var control, the additional currents due to the consumption of reactive power increase the maximum cable utilisation.

This is also reflected in the probability of over-loads shown in Figure 44, where there is a very low risk using volt-watt control, but a certainty of over-loads with volt-watt-var control for around 5% of feeders, in the event that all V2G installations export at the same time. As above, the very low probabilities for volt-watt control are due to the threshold selection where over-loads are avoided other than where they already occur in the existing smart meter data.

The use of locally customised thresholds allows the mean export power from V2G to be increased, as shown in Figure 45. There is no trace here for the existing baseline case, for which there are no exports, but the more granular customisation to feeder level allows for a greater export than with customisation to substation level.

These plots all show curves relative to the arbitrarily selected default of a thresholds between 240 V and 245 V. Compared to this default, the local customisation allows for a significant increase in the mean export power, with the consequence that the maximum voltages and currents increase closer to their limits, making the best use of the available headroom.

A less conservative default could have been assumed, in which case the relative gains in export power would have been less. It may be difficult to impose a default standard, adopted by all V2G installations, where the export power is reduced to zero for a voltage of 245 V or above, especially as this is well below the permitted upper limit of 253 V. However, raising this default, for example to progressively reduce export power between 250 V and 255 V, would allow the risk of over-voltages or over-loads to remain.

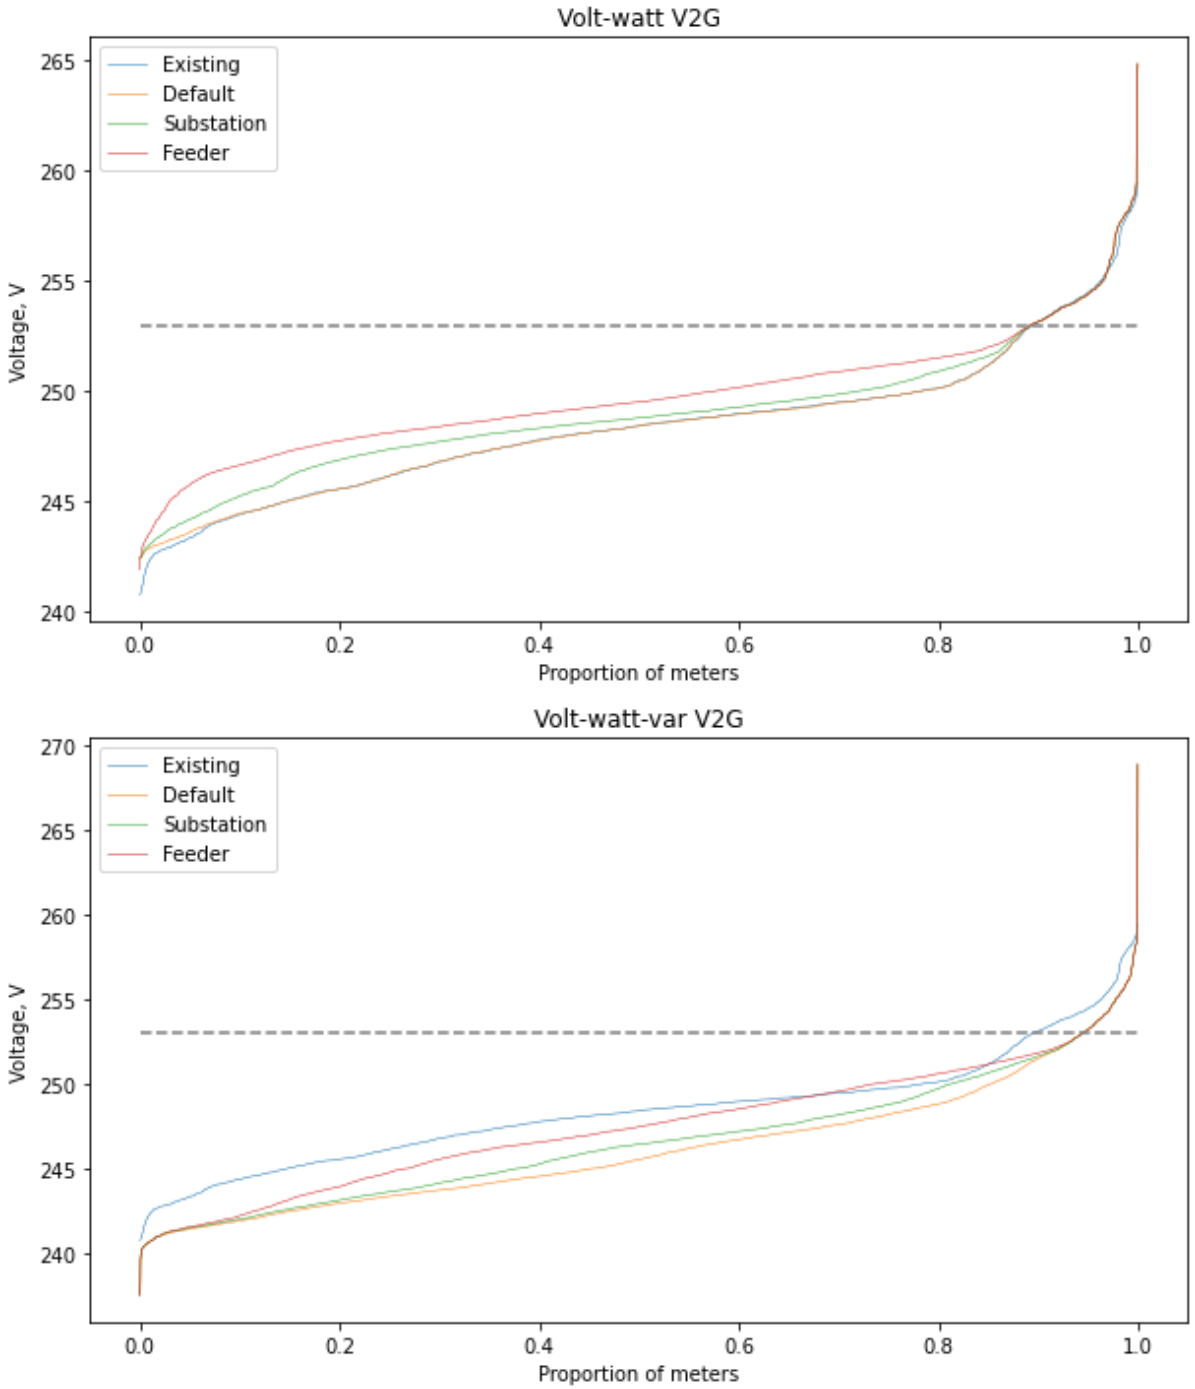

Figure 41: Maximum voltages with V2G exports using locally customised thresholds

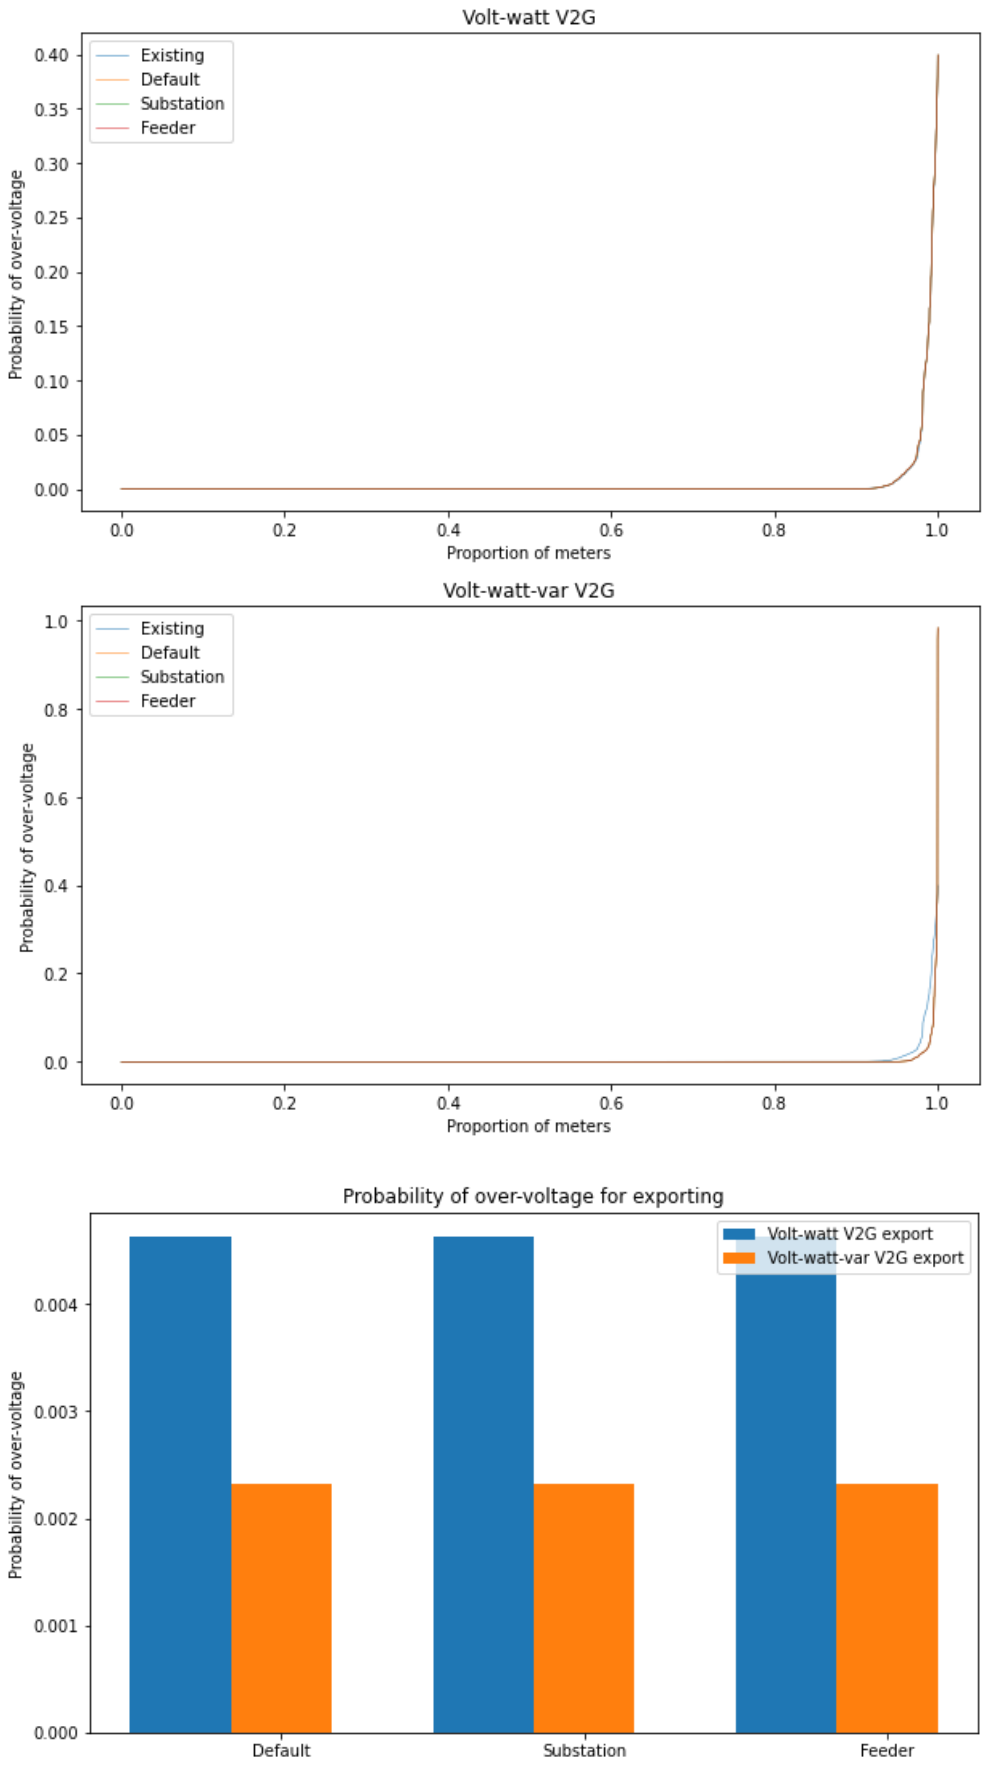

Figure 42: Probability of over-voltage with V2G exports using locally customised thresholds

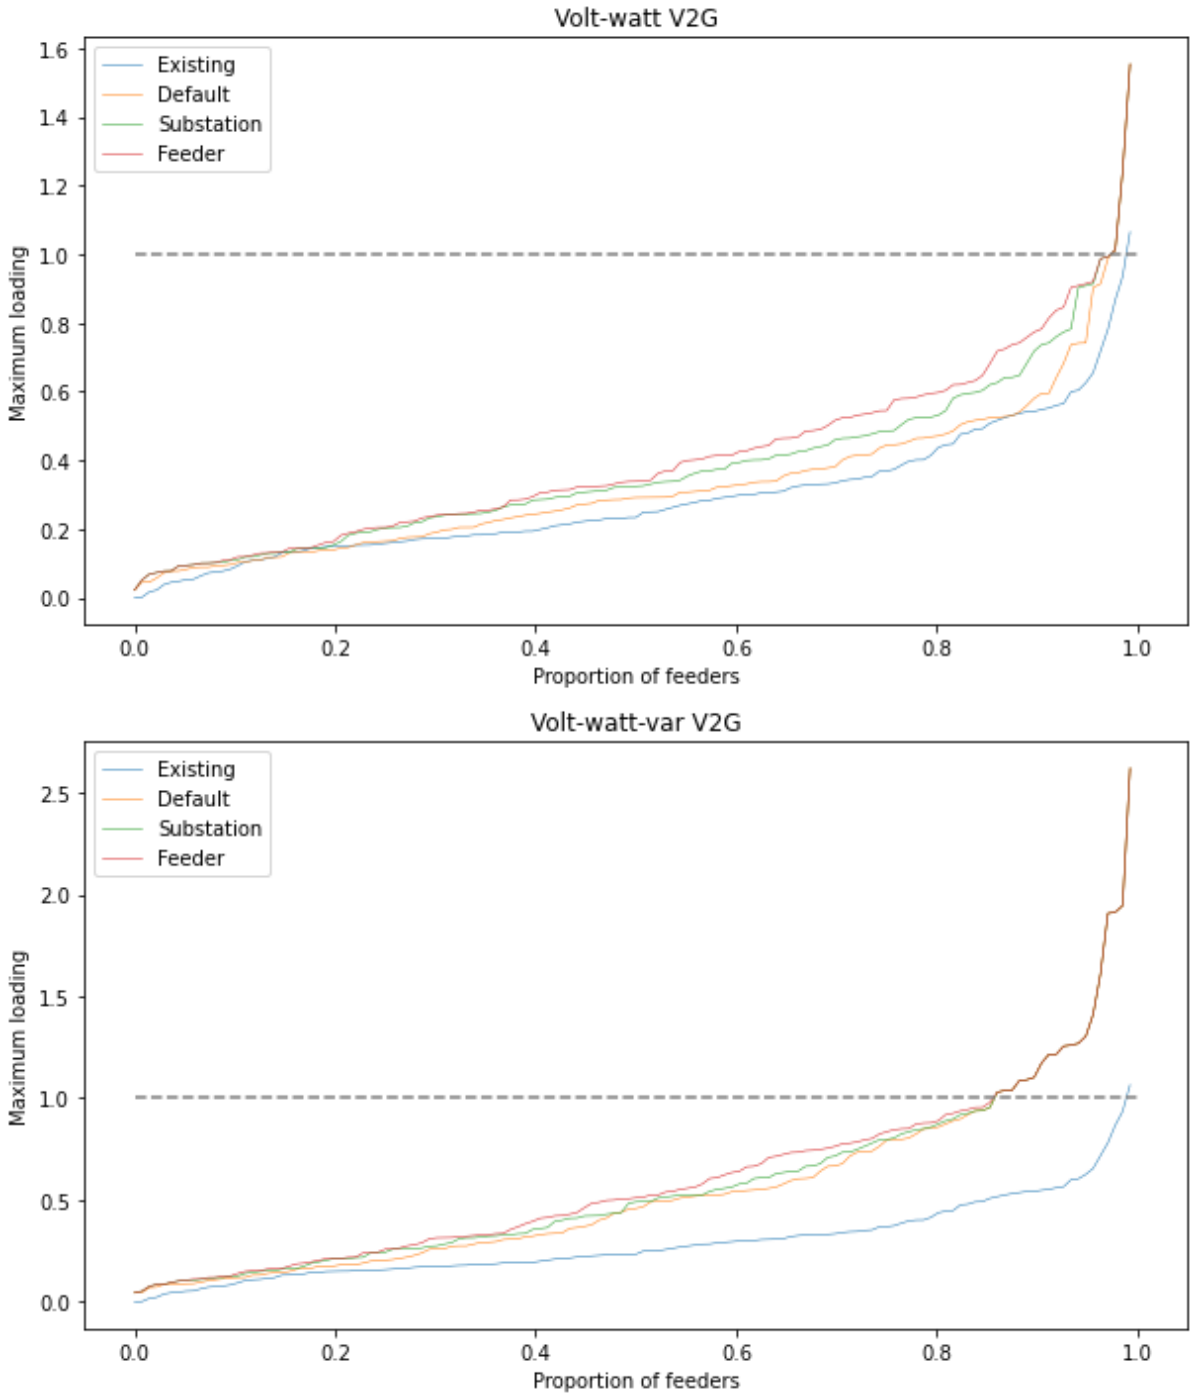

Figure 43: Maximum cable utilisation with V2G exports using locally customised thresholds

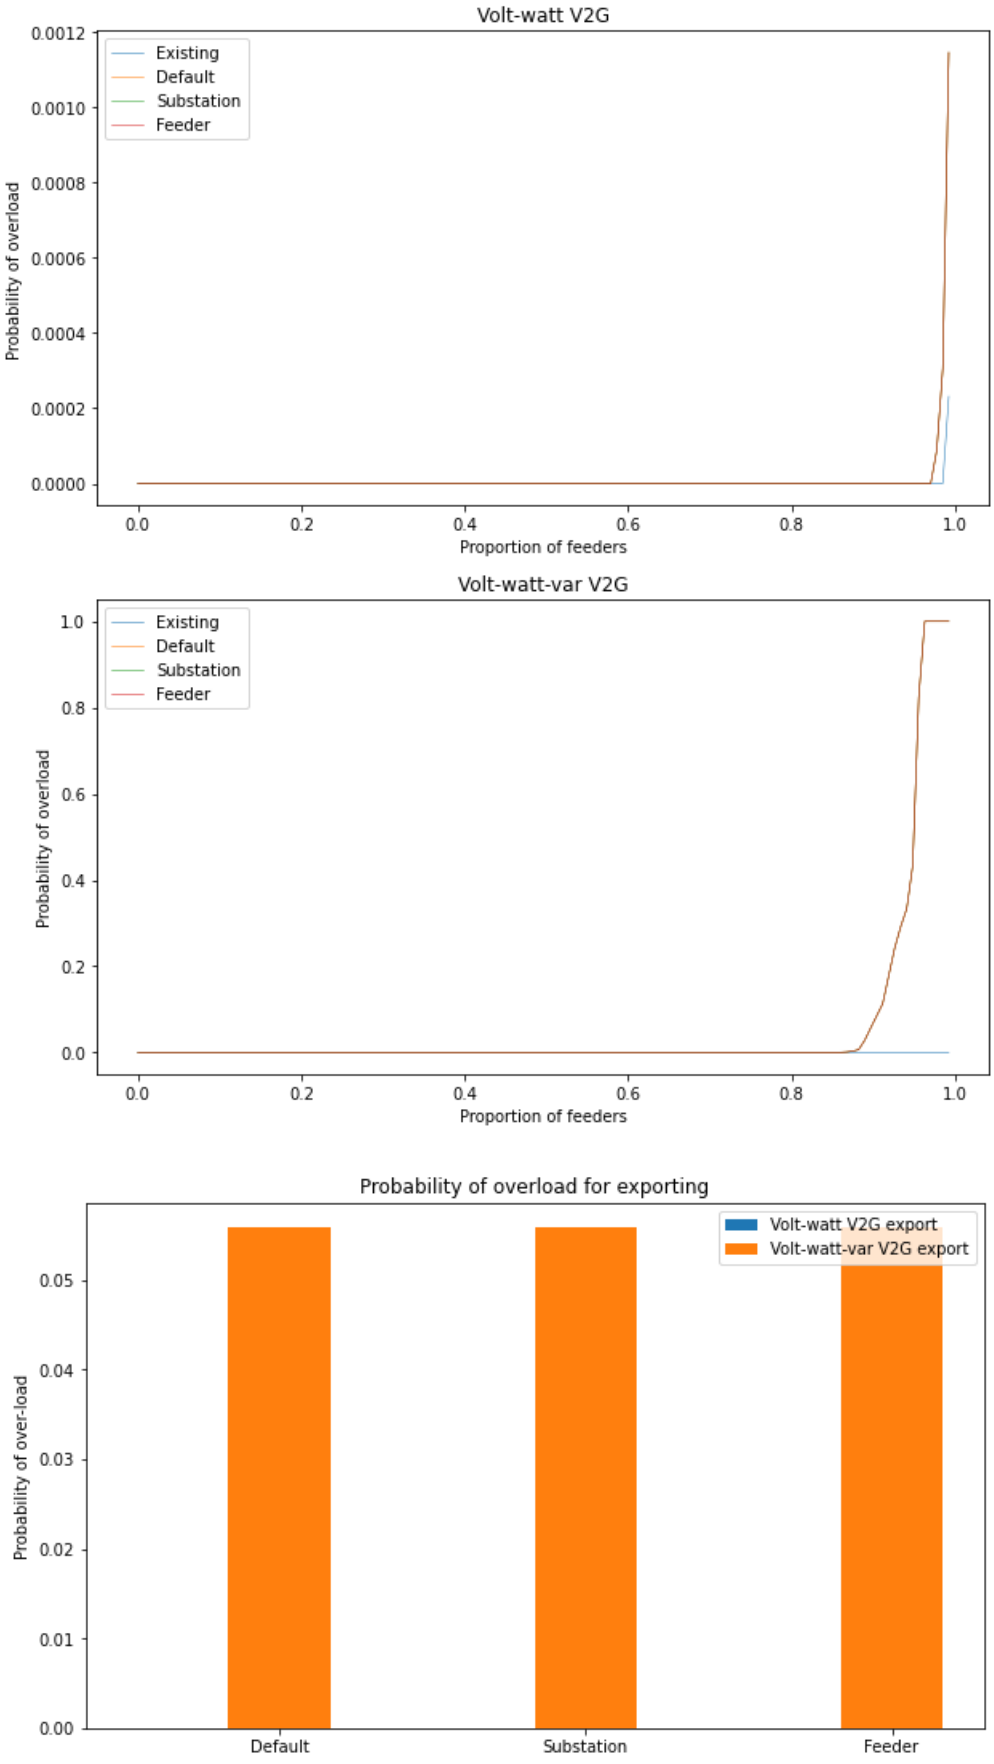

Figure 44: Probability of over-load with V2G exports using locally customised thresholds

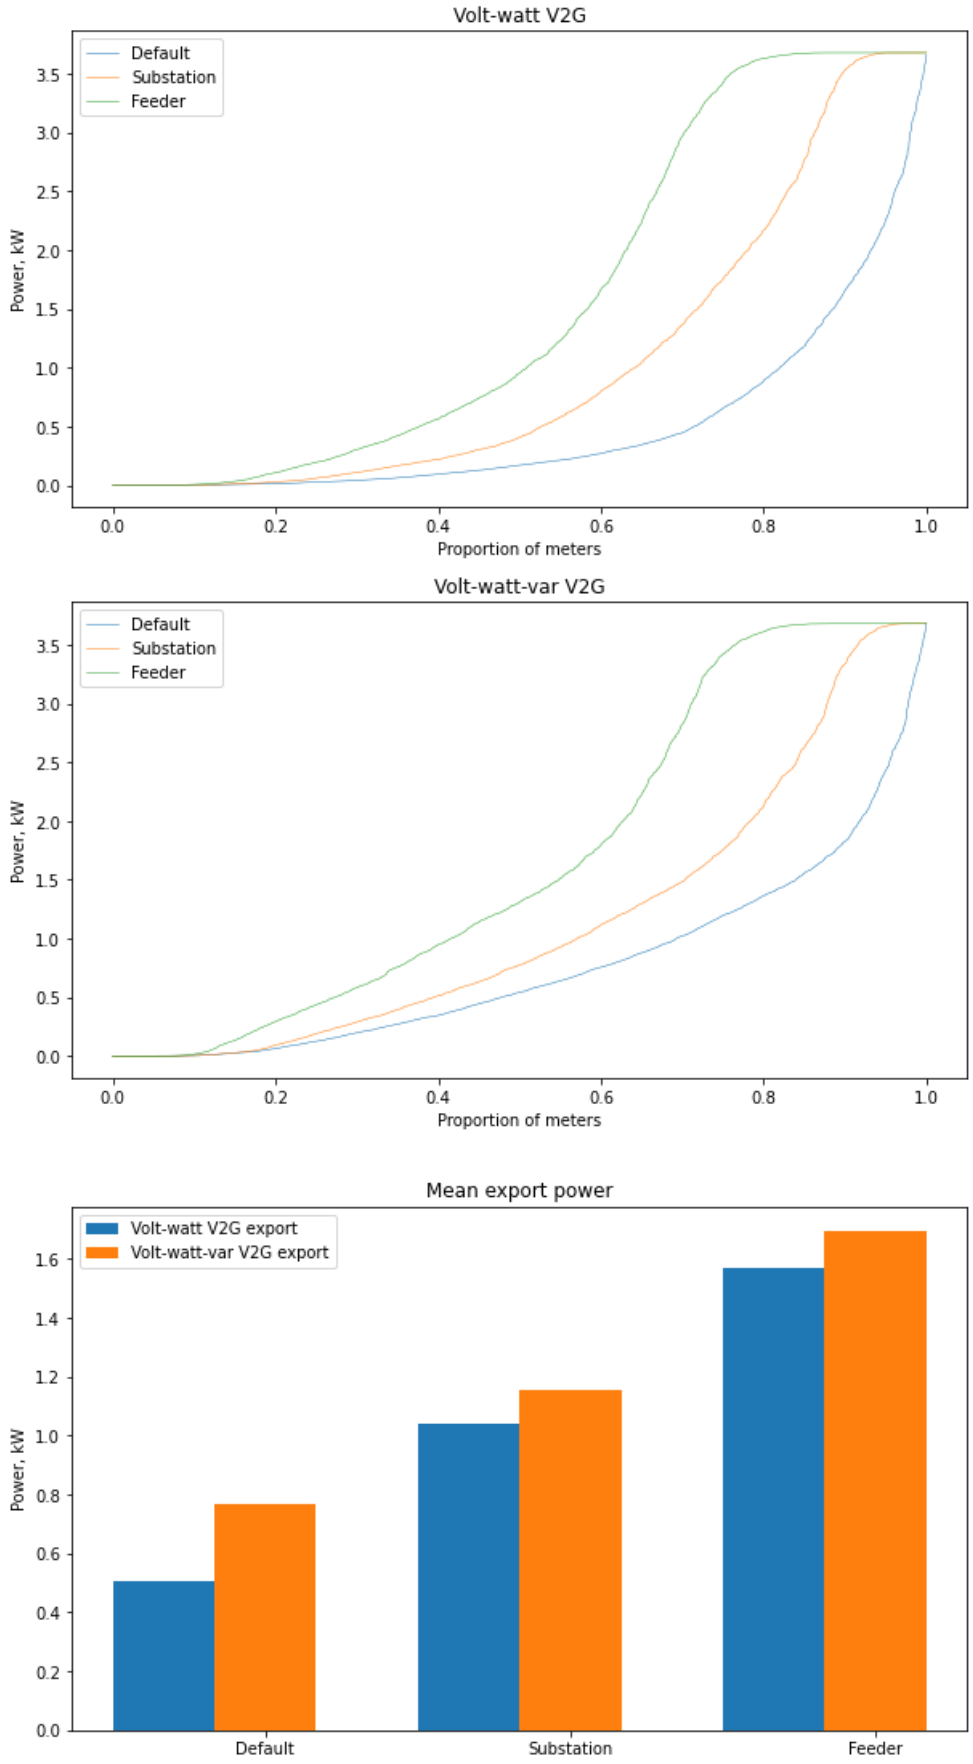

Figure 45: Mean V2G export power using locally customised thresholds

A practical implementation of this method would require a means of determining which thresholds to use for each feeder. This might operate based on experience of historical smart meter data.

It would be convenient if the maximum voltage in the absence of V2G export events could be used to indicate which thresholds should be set for each feeder. Possibly a lower threshold would be required where the voltages are otherwise higher.

Figure 46 shows the threshold settings for each feeder, plotted against maximum voltage for the baseline scenario. The y-axis value shows the mid-point of the ramp range, so for example control thresholds between 250 V and 255 V will have a mid-point of 252.5 V.

Unfortunately, there is no discernible pattern to the optimum threshold settings, although no feeders with a high voltage in the absence of V2G are assigned the higher threshold settings. Where the baseline voltage is lower, the optimum threshold settings could have any of three ranges, and the lower threshold range of 240 V to 245 V could be required for feeders with maximum voltages ranging from 245 V up to 260 V. However, the plot confirms that the higher threshold settings are always selected where the maximum baseline voltage is within the permitted limit.

Similarly, the maximum current loading for the existing baseline case is also unhelpful in determining the thresholds required, as shown in Figure 47 where the lowest threshold may be needed for feeders ranging from near zero loading to full utilisation of the rated capacity.

Figure 48 shows the equivalent case where the threshold settings are plotted against the V2G voltage with no controls applied. Intuitively this seems a more appropriate voltage to use for a threshold selection algorithm as it relates more directly to the extent to which controlled constraints are needed. However, this raises a practical difficulty that the feeder would need to be operated for some period without V2G constrained, potentially exceeding the voltage and thermal limits, in order for the optimisation algorithm to learn the appropriate voltage range and therefore to determine the settings. The mode of operation would likely be unacceptable.

Regardless of this practical issue, neither the baseline voltages with no V2G, or the voltages with V2G active and uncontrolled seem to suggest an obvious relationship with the optimum thresholds, and so it would be necessary for the control system to adopt thresholds and then adjust them periodically if the voltages or currents are found to be too high or exports are constrained unnecessarily.

Figure 46: Feeder threshold selections for volt-watt control relative to baseline maximum voltage

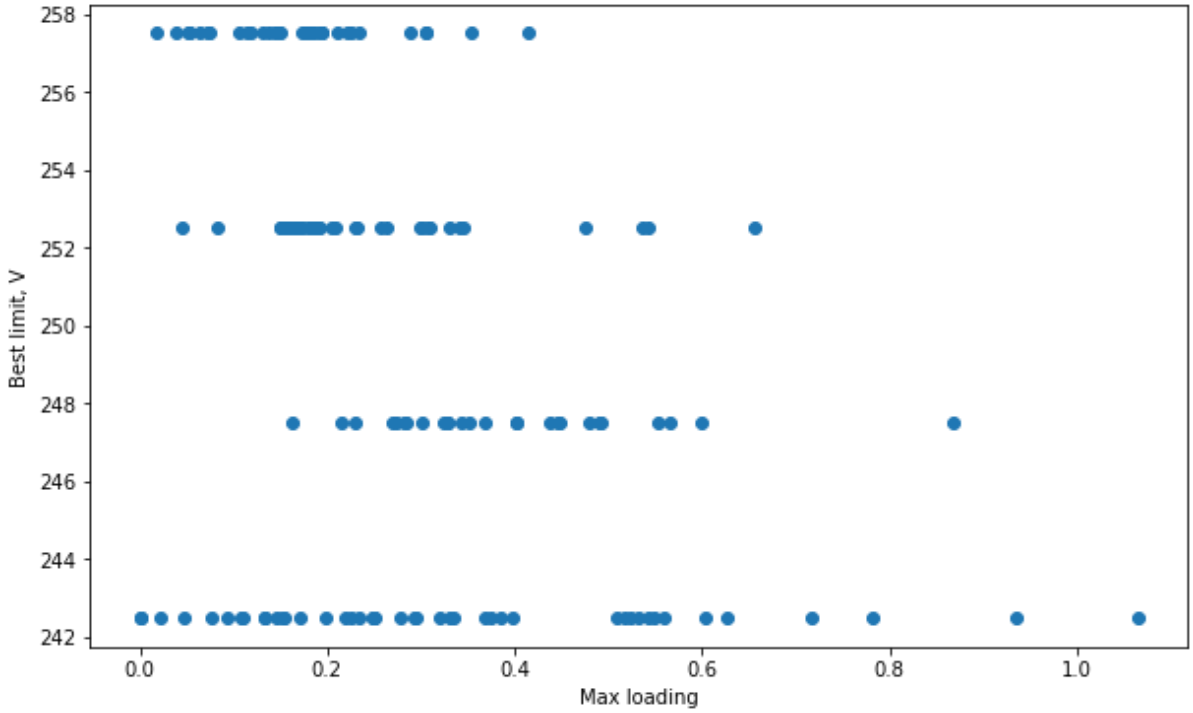

Figure 47: Feeder threshold selections for volt-watt control relative to baseline maximum loading

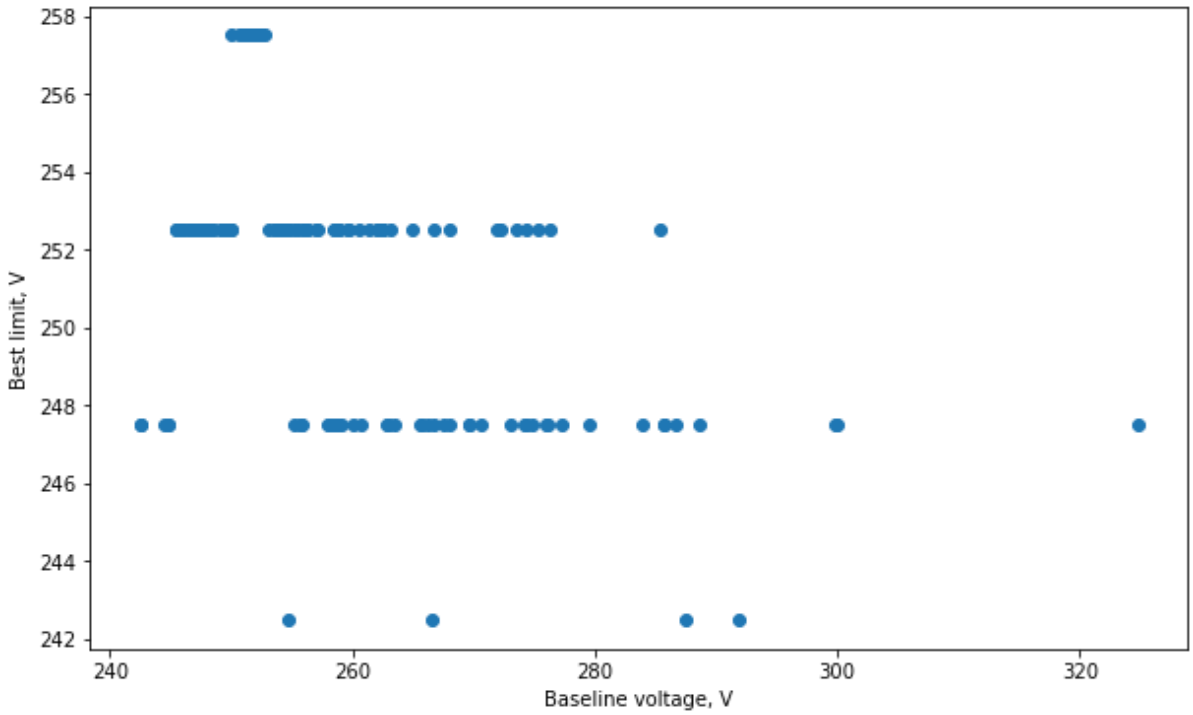

Figure 48: Feeder threshold selections for volt-watt control relative to V2G maximum voltage

5.5 Time-gated thresholds

The previous sections have demonstrated that thresholds for volt-watt control may need to be set significantly below the upper voltage limit to ensure that V2G does not cause either voltage or current limits to be exceeded. There is a risk that the use of these thresholds will unnecessarily constrain export powers if they are applied in time periods when the V2G would not have caused limits to be exceeded even without controls.

This section considers whether the mean export power could be increased if the volt-watt control is selectively enabled only in time periods when necessary.

An idealised implementation could be modelled by assuming that V2G could operate without controls for any half-hour in which the voltage and current limits are not exceeded. In half-hours where either of the limits are exceeded, a default threshold could be applied using settings that are common to all V2G installations.

This approach models an upper bound on the improvement in mean export capacity relative to the use of the default threshold for all time periods, where the thresholds are enabled individually for each half-hour period. A practical implementation would require a real-time communication system so that every V2G installation would be aware of the voltages and aggregated current elsewhere on the feeder. This model is therefore not achievable in practice using smart meter data that is collected retrospectively, but the results provide a useful upper bound on the increase in export power that could be obtained.

Results for this idealised approach are formed by combining results for V2G with no controls with results for V2G with a threshold applied. For each individual half-hour period, the results for V2G with no control are selected if this would not cause limits to be exceeded. Otherwise, the results for V2G with the volt-watt control are selected.

A more practical approach could use previous smart meter data to select half-hour periods in the day when voltage or current limits have not previously been exceeded. The selected half-hours would need to be reviewed periodically, and there is also a risk that half-hour periods could be excluded

even though the feeder would remain within limits on most days. This approach could be applied either for each LV feeder, or for all the LV feeders at a substation.

This approach is modelled here assuming perfect foresight of the half-hours when limits might be exceeded. Results for this approach are formed in the same way as for the idealised method described above, but with the selection being applied for each of the 48 half-hour periods in a day. V2G with volt-watt control is selected if limits are exceeded for the half-hour period on any day within the simulation time-series.

The mean export powers for results with time-gated thresholds are shown in Figure 49, where the default case uses volt-watt control with thresholds applied continuously. By definition, the probability of over-voltages or current over-loads is the same with time-gated thresholds as in the default case, since the thresholds are always applied if limits are exceeded.

The plot shows that there is a benefit of increased export power if volt-watt control is only enabled in time periods when it is required. The benefit is less for higher thresholds which apply less of a constraint, but would allow a greater risk of voltage or current limits being exceeded. For a default threshold range of 240 V to 245 V, there is an increase in the mean export power per customer of around 200 W if the time periods are selected on the basis of individual feeders, and an increase of around 100W if the time periods are selected over all the customers at a substation.

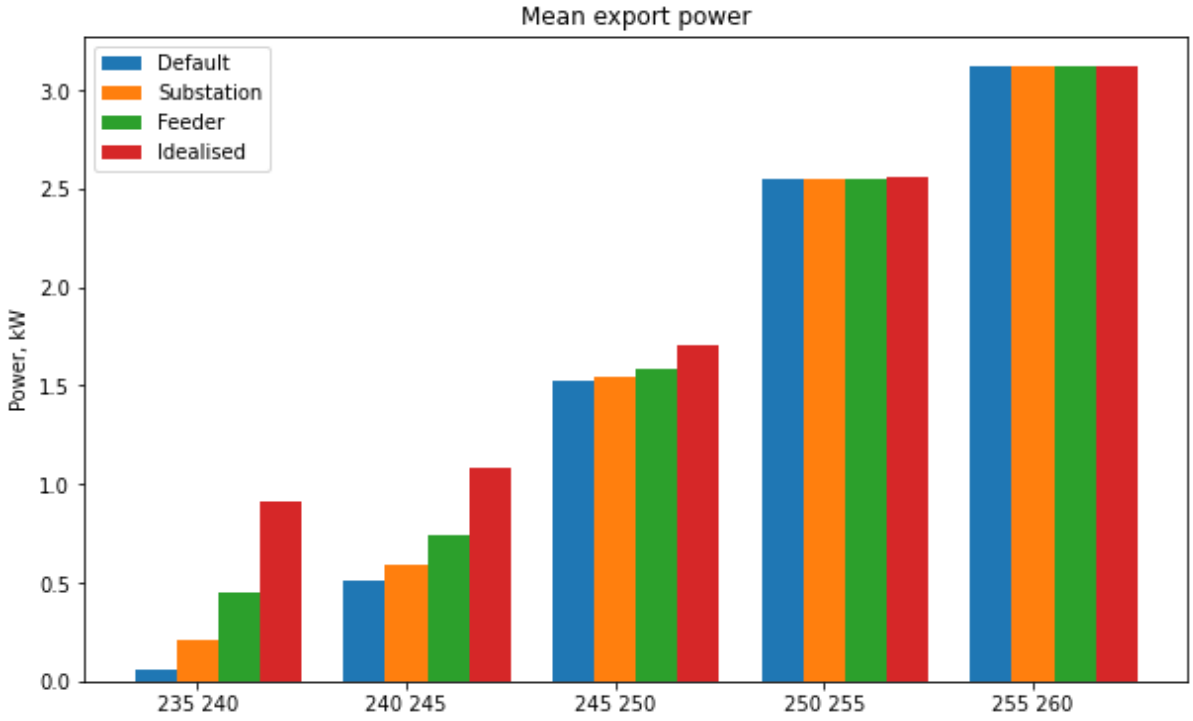

Figure 49: Mean V2G export power with time-gated threshold

6. Conclusions

Results from power-flow simulations are presented where smart meter data has been combined with the possible future demands from V2G exports to model the expected voltage rise, thermal loading and constraints to exported power. The smart meter data is used to provide a baseline representing the existing demand and the V2G demand is added to this by superposition.

We believe that these are novel techniques and that this may be the first time that smart meter data has been integrated into power-flow analysis within NGED.

This report extends previous work in Work Package 2 D2, using a longer periods of smart meter data so that results are more representative and include demand variations over the summer season when solar PV generation will be higher, and heating demand will be lower, so that there is a greater risk of voltage rise that may limit the headroom for V2G exports.

In preparation for this modelling the report has first examined the impact of the existing demand in terms of customer voltages and thermal loading. This work finds that there are many instances where measured voltages fall outside the permitted ranges, but that these are very infrequent events and possibly due to faults or maintenance work. However, there are a few feeders where voltages exceed the upper limit of 253 V more persistently. On 1% of feeders, voltages are above the upper limit for 50% of time samples, and around 3% of feeders are out of range for 10% of the time. Low voltages also occur but it is less common to find persistent under-voltages than over-voltages. There are around 1% of feeders that have a minimum voltage below the limit for 1% of time.

Most feeders have thermal capacity available, with the majority of feeders having an aggregated demand that represents no greater than 70% utilisation relative to the ratings of mains cables in the first branch out from the substation. This headroom is available for future LCT growth, and for V2G operation. The most likely concern for V2G operation would be any synchronised re-charging after an export event, where the imports would combine with the existing demand. Exports would likely offset the exiting demand, such that the headroom would be greater than the capacity implied by the cable rating. However, there is also the possibility that recharging will occur when there is already a net export, and therefore less capacity available.

V2G operation without any additional control can have a significant impact on voltage ranges and thermal loading. Assuming a worst-case where all V2G customers export at the same time, around 70% of feeders would exceed the upper voltage limit, with many being significantly above this. Over 20% of feeders would exceed 100% cable utilisation relative to the rated capacities of the branches directly connected to the substation.

The corresponding recharging is also a concern. If this were equally coordinated with all V2G installations charging at the same time, nearly 40% of feeders would have voltages below the lower limit. There would also be over 40% of feeders exceeding 100% utilisation. Thermal loading is a greater concern for the recharging than for the V2G exports as the imports can use the full 7 kW rated capacity of the inverter rather than the permitted export power of 3.7 kW.

The modelling then considers the impact of volt-watt and volt-var control algorithms, and volt-watt-var where both techniques are combined. Initially the models consider fixed threshold settings, as they would be with a standardised factory configuration deployed to all customers.

Volt-watt and volt-watt-var control are both effective in ensuring that voltage rise does not exceed limits. Volt-var also reduces voltages, but many feeders still have significant over-voltages.

Volt-watt control is also effective for managing thermal loading, but both the volt-var and volt-watt-var methods increase currents due to the additional reactive power. Volt-var when used alone increases thermal loading more than with uncontrolled V2G.

However, the improvements through volt-watt and volt-watt-var control are achieved at the expense of constraining the export power. Approximately 20% of feeders would have mean export power reduced by 1.5 kW or more. The additional reactive power in volt-watt-var control allows this reduction to be slightly reduced, although the risk of thermal overloads would increase.

The results also demonstrate an unfairness in the constraints between customers on the same feeder. Intuitively it would be expected that customers at the ends of feeders are constrained more. The results have shown this effect, but a slightly more notable inequality is that customers near to the substation are constrained much less. The exports from those at the ends of feeders are reduced only slightly more than for the average customer.

Variations in the fixed threshold settings for volt-watt and volt-watt-var control can be used to increase the confidence that feeders will remain within voltage and thermal loading limits. Clearly this improves the V2G operation from the network perspective but reduces the total exports that can be achieved from V2G, therefore offering less benefit for balancing power on the wider electricity system.

A more targeted approach where threshold settings are individually selected for each substation allows feeders to operate closer to their headroom limits and so increases the scope for power exports. The customised thresholds can be selected such that the cable utilisation remains below 100% and so that the maximum voltages do not exceed the available headroom.

Selection of customised thresholds for each substation allows significantly lower export power constraints than the use of a fixed threshold range that would have given similar confidence that the maximum voltage would not be exceeded.

This approach can be extended to a finer spatial resolution by defining a customised limit for each LV feeder. This gives a further increase in the total export power. Work so far has not identified any simple approach that might be used to determine the optimal threshold for each feeder, for example by considering how this optimal threshold varies with the maximum voltages in the baseline case or with uncontrolled V2G. More likely, the selection of thresholds would need to follow an iterative approach with ongoing monitoring of the smart meter voltage data and then thresholds being reduced as the uptake of V2G, or the frequency of exports, increase over time.

An alternative approach to customising the control thresholds could selectively enable the V2G control techniques according to the time of day. At times when voltage or thermal constraints are unlikely to be exceeded, V2G could operate without any controls, but with the controls being enabled when necessary. This might, for example, require V2G control during the middle of the day when solar PV exports are higher and demand is lower, but allow controls to be disabled in the evening when demand is higher and voltages are lower.

The benefits of this approach are relatively low, giving an increase in the mean export power per customer of around 200 W if the time periods are selected on the basis of individual feeders.

7. References

[1] Andrew Urquhart and Murray Thomson, “V2G Dynamic Headroom Control, Site Selection Report,” Loughborough University, 2025.

[2] Andrew Urquhart and Murray Thomson, “V2G Dynamic Headroom Control, Phase Identification Report,” Apr. 2025.

[3] Andrew Urquhart and Murray Thomson, “V2G Dynamic Headroom Control, Initial Simulations Report,” Jun. 2025.

[4] Matthew Pope, “Standard Technique SD5H/1,” Sep. 2021.

[5] E. & I. S. Department for Business, “Subnational Electricity and Gas Consumption Statistics.” Accessed: Jun. 30, 2023. [Online]. Available: https://assets.publishing.service.gov.uk/government/uploads/system/uploads/attachment_data/file/1126284/subnational_electricity_and_gas_consumption_summary_report_2021.pdf

[6] National Grid Electricity Distribution, “Distribution Future Energy Scenarios (DFES).” Accessed: Mar. 05, 2025. [Online]. Available: https://connecteddata.nationalgrid.co.uk/dataset/dfes

[7] R. Lowe and Department of Energy and Climate Change, “Renewable Heat Premium Payment Scheme: Heat Pump Monitoring: Cleaned Data, 2013-2015,” UK Data Service. SN: 8151. Accessed: Jul. 19, 2023. [Online]. Available: https://beta.ukdataservice.ac.uk/datacatalogue/doi/?id=8151#!#1

[8] Ofgem, “Average gas and electricity use explained.” Accessed: Jun. 30, 2023. [Online]. Available: https://www.ofgem.gov.uk/information-consumers/energy-advice-households/average-gas-and-electricity-use-explained

[9] Energy Savings Trust, “Measurement of Domestic Hot Water Consumption in Dwellings.” [Online]. Available: https://assets.publishing.service.gov.uk/government/uploads/system/uploads/attachment_data/file/48188/3147-measure-domestic-hot-water-consump.pdf

[10] Stark, “Degree Days For Free.” Accessed: Oct. 20, 2023. [Online]. Available: https://poweredby.stark.co.uk/SEO/SEO.aspx

[11] I. Richardson and M. Thomson, “Integrated domestic electricity demand and PV micro-generation model,” 2011. [Online]. Available: https://dspace.lboro.ac.uk/2134/7773

[12] I. Richardson and M. Thomson, “Integrated simulation of photovoltaic micro-generation and domestic electricity demand: a one-minute resolution open source model,” in *Microgen II: 2nd International Conference on Microgeneration and Related Technologies, Glasgow, 4 - 6 April.*, Glasgow: © Microgen’II Organizing Committee, 2011. Accessed: Mar. 16, 2013. [Online]. Available: http://hdl.handle.net/2134/8774

[13] Western Power Distribution, “Electric Nation Customer Trial Final Report,” p. 591, 2019, [Online]. Available: https://www.westernpower.co.uk/downloads/64378

[14] M. Bognar, “Weibull distribution.” Accessed: May 01, 2021. [Online]. Available: https://homepage.divms.uiowa.edu/~mbognar/applets/weibull.html

[15] Andrew Urquhart and Murray Thomson, “V2G Dynamic Headroom Control, Control Algorithms Report,” Feb. 2025.

[16] SA Power Networks, “TS129 Small EG Connections Technical Requirements-Capacity not exceeding 30kVA,” 2021, Accessed: Jan. 20, 2025. [Online]. Available: https://www.sapowernetworks.com.au/public/download.jsp?id=318532