WESTERN POWER DISTRIBUTION

INNOVATION

NEXT GENERATION NETWORKS

Solar Storage

Final Report

Report Title	:	Solar Storage: Final Report
Report Status	:	Issued
Project Ref	:	WPD NIA 004
Date	:	30.01.2019

Document Control

	Name	Date
Prepared by:	Christie Sims & Jenny Woodruff	01.10.2018
Reviewed by:	Jenny Woodruff & Christine Coonick & Luke Hosking	12.11.2018
Recommended by:	Jenny Woodruff	12.01.2019
Approved (WPD):	Roger Hey	14.01.2019

Revision History

Date	Issue	Status
30.01.2019	1.1	Revoked
01.02.2019	1.2	Issued

Contents

Glossary

Term	Definition
AAHEDC	Assistance for Areas with High Electricity distribution costs
ANM	Active Network Management
BMS	Battery Management System
BRE / NSC	Building Research Establishment / National Solar Centre
BSR	British Solar Renewables Limited
BSUoS	Balancing Use of System
BYD	Battery manufacturing company that supplied the solar storage battery.
CCGT	Combine Cycle Gas Turbines
CCL	Climate Change Levy
CfD	Contracts for Difference
CM	Capacity Market
CMZ	Constraint Management Zone
Co-Located	Within this report, co-located storage refers to energy storage that has been installed behind the meter of a solar park or other distributed generator.
DG	Distributed Generation
DNO	Distribution Network Operator
DUoS	Distributed Use of System
DSO	Distribution System Operator
DSR	Demand Side Response
EFR	Enhanced Frequency Response

EMS	Energy Management System
FiT	Feed in Tariff
FFR	Fast Frequency Response
GB	Great Britain
GSP	Grid Supply Point
HMI	Human Machine Interface
HV	High Voltage
LCNI	Low Carbon Networks and Innovation
LCF	Levy Control Framework
LV	Low Voltage
NIA	Network Innovation Allowance
PEA	Project Eligibility Assessment
PCS	Power Conversion System
PoC	Point of Connection
PQM	Power Quality Meter
PV	Photovoltaic
RCRC	Residual Cashflow Reallocation Cashflow
RES	Renewables company that managed the battery installation and developers of the RESolve control software.
ROC	Renewable Obligation Certificate
SOC	State of Charge
STOR	Short Term Operating Reserve
TNUoS	Transmission Network Use of System
WPD	Western Power Distribution

1. Executive Summary

Solar Storage was funded through Ofgem's Network Innovation Allowance (NIA), registered in April 2015 and completed operating in May 2018. The project consisted of installing a battery 'behind the meter' of an existing solar park, to investigate the potential use-cases and benefits of co-locating storage in this way.

The aims of the project were to:

1) Quantify the potential value to network operators and others of integrating storage with solar generation by demonstrating a set of use cases.

2) Use real-world operation of an integrated utility scale storage / generation system to provide data to regulators and potential investors.

3) Demonstrate safe, reliable operation of the system under operational conditions.

The battery chemistry was lithium iron phosphate, which is less energy dense than lithium-ion batteries but has the advantage of having greater thermal stability and is at lower risk of overheating. The battery began operation in October 2016 and operated via pre-programmed schedules. Despite a number of technical challenges that delayed the testing schedule, the schedule of tests was completed and during the project several improvements were made to the control software enabling more realistic testing of the use cases. By the end of the project the battery capacity fade was 7.4% however, round trip efficiency had not diminished due to better air-conditioning units being installed. At the end of the project, learning was extended by selling and relocating the battery, which was not known to have occurred previously in the UK.

The analysis showed that while the battery was technically capable of performing the use cases, the battery size limited the degree of impact in some cases, such as for managing network voltage either using real or reactive power. The arbitrage use case offered limited returns due to the site's power purchase agreement having a relatively flat profile which does not reflect the real differential between peak and off-peak prices. This is due to the value investors place on predictability when purchasing a renewable generation asset, over and above the potential upside from higher but variable prices. This is also true of banks that lend higher percentages or at better rates against predictable stable quarterly income levels. When co-locating projects in the future, there will be a balancing act of allowing the energy storage asset access to the volatile market while retaining the predictability of the solar park's revenues. Arbitrage is expected to play a more significant role in the future business case of storage as incomes from frequency response decline due to a saturated market.

The arbitrage use case could be combined with network peak lopping for sites which have peak load at the same time as peak prices. The solar output peak lopping use case could also potentially also be combined as the charging and discharging periods are complementary. Reactive power services could also boost incomes while sacrificing little of the battery's real power capacity, though impact would be limited for a small capacity battery. Services to third parties with constrained connection agreements could well be cost effective but are so location specific as to have little impact on the progress of battery development.

During the course of the project, significant changes have occurred to the price of frequency services but also the rationalisation of National Grid services has opened up the market such that flexibility providers can switch between services more easily, but conversely it is harder to make the business case for investing in storage. This is likely to limit the future geographic spread of storage to service Distribution Network Operator (DNO) requirements with existing storage located at sites where DNO services are not required. The market is continuing to change, with grid charging now also under review which reduces investor certainty still further.

New flexibility options are being offered and trailed by various DNOs as the transition to DSOs gathers pace, but often these are often too flexible in their contract terms and options to drive investment. These services offer existing asset owners and operators the option to bid in and out of the market easily while offering a good premium over the existing operating regime. However without more predictable revenues, such as those from more rigid contracts, the energy industry is unable to encourage investment and the pace of energy storage development could stall. The key to an investable energy storage project is having a justifiable and robust business case, which realistically involves having a dominant (and preferably proven) use case, with additionality provided by DSO services. Without the main business case, and with flexible markets unwilling to offer long term contracts that could become bankable and investable, then storage that isn’t reliant on National Grid contracts is likely to continue to be uncommon.

Connecting a battery to a solar park behind-the-meter, while potentially saving costs assuming they share export capacity, currently prevents the storage from accessing several lucrative and reliable revenue streams such as arbitrage and grid services. However, with increasing wholesale price variability coupled with future increased amount of solar parks connected to the grid, the midday price drop versus the evening price peak could make solar peak lopping economically viable. With this future potential in mind, as well as opportunities for more complex arrangements that could benefit both assets, any regulatory barriers preventing the future roll out of this should be solved now.

Connecting batteries in front of the meter (or behind the meter but using sub metering to separate them) on solar parks could be helpful, with unused land within solar fields able to be used and taking advantage of the infrastructure already installed. Batteries in these rural areas can provide grid support to the potentially weaker network.

New build solar and storage systems that are metered separately but share grid assets should be supported by the DNO, as this can reduce the grid costs significantly. This can bring forward projects that would otherwise not reach their investment targets and would fail to get built.

The new DNO commitment to flexibility shows an increasing awareness of the need to find alternatives to traditional expensive reinforcement, and the battery project has been able to prove the reliability of the battery in a variety of roles, and shown that the industry is in a good place to meet the demands of a flexible smart grid.

DNOs should continue to improve signposting of the required locations for future flexibility services, but it may be that battery development will be very limited if it relies on DNOs to provide contracts sufficient to justify investment.

2. Project Background

2.1 Overview

This project was initiated at a time when battery storage costs were steadily reducing at the same time as the use of flexibility services by DNOs was predicted to increase significantly as they transitioned to Distribution System Operators (DSOs). The scenario of a battery associated with a solar park was chosen because of the wide range of potential services, listed as the nine use cases below in Table 1, where a battery can provide benefits to different parties. These are described more fully in Section 9, Use Cases.

Usage Case	Beneficiary
1) **Arbitrage** - Sell electricity for a higher price per kWh.	Battery Owner
2) **Peak demand limiting** at the local primary.	DNO
3) **Local demand profile matching** e.g. as a service to a customer with a soft inter-trip connection who would otherwise be constrained.	DNO / load customer
4) **Low demand grid voltage support** - Raise minimum demand to limit voltage rise.	DNO
5) **Voltage control by reactive power.**	DNO
6) **PV Export limiting** - Peak lop generation to enable solar parks with an installed capacity over that of the connection agreement.	Solar Park Owner
7) **Variable PV export limiting -** Change peak lopping level (glass ceiling).	DNO
8) **PV power quality improvement** - Smoothing / Power Quality, Ramp Rate Control	DNO / Solar Park Owner
9) **Multiple storage system control** - (Not trialled, included for discussion only).	Multiple parties

Table 1: Use Cases

The project does not include the provision of services to National Grid, such as;

- Enhanced Frequency Response (EFR);
- Firm Frequency Response (FFR); and

- Short Term Operating Reserve (STOR).

While these services are currently major drivers of storage connections, it was considered that to trial these services would cause duplication of the Smarter Network Storage project[1] carried out by UK Power Networks.

As well as evaluating the efficiency and efficacy of the battery at delivering the use cases, the project also estimated the financial benefits and considered how these use cases reflected the potential for layering revenue streams.

To support the understanding of issues around battery sizing, investment and impact, the project also included some complementary elements which were;

- techno-economic modelling, provided by SRI Technologies;
- regulatory framework assessment, provided by Utilities Insight; and
- power quality monitoring, provided by Argand Solutions

2.2 Location

The solar park, where the battery was installed, is electrically connected to a clean 11kV feeder supplied by the Millfield primary substation. This has been altered to introduce an additional ring main unit to provide isolation between the battery and the solar park.

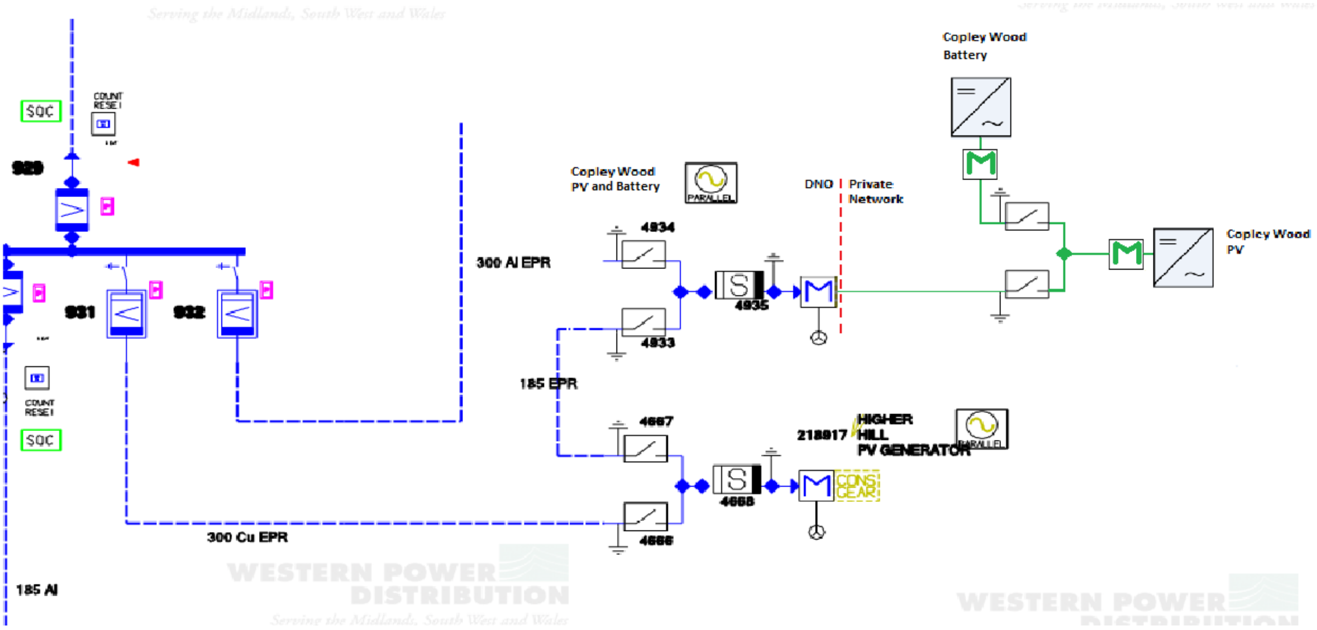

Figure 1: Electrical connection to Copley Wood battery

[1] https://www.ukpowernetworks.co.uk/internet/en/news-and-press/press-releases/Trailblazing-storage-project-leads-the-way-to-low-carbon-future.html

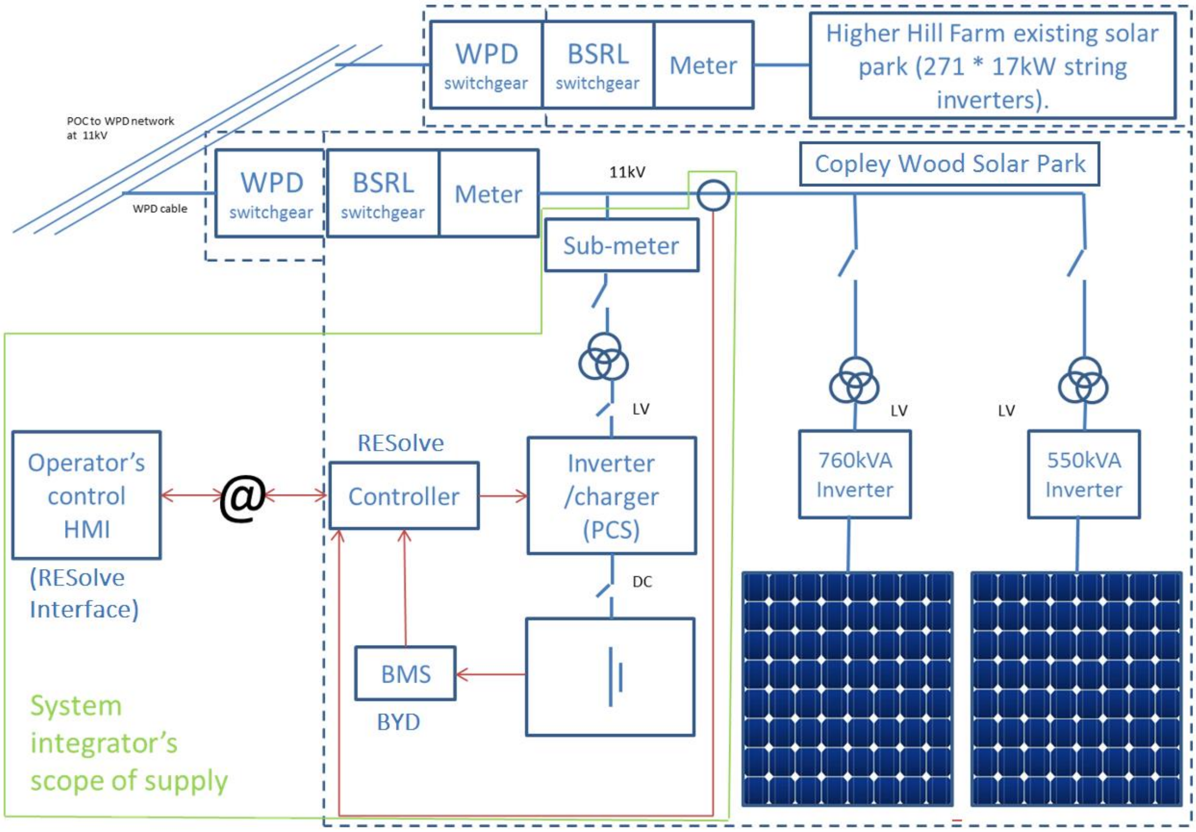

Figure 2: Battery site schematic

The battery was metered separately and connected via a low voltage (LV) isolating transformer.

The operator's control (HMI) was a system called RESolve, provided by RES. This umbrella term covers a diverse array of control modes and a SCADA data storage system that was also managed by RESolve. The key components discussed within this report are the state of charge (SoC) manager, the use cases/modes and the underlying SCADA data storage system.

2.3 Timeline

The project was registered in April 2015 and originally expected to run until April 2018. The battery was commissioned on site in October 2016. While the testing was completed in April 2018, the project was extended further to allow for a process to sell and relocate the battery as required.

2.4 Learning objectives

The project objectives were to;

- quantify the potential value to network operators and others of integrating storage with distributed generation (DG);
- use real-world operation of an integrated utility scale storage/ generation system to provide data to regulators and potential investors; and
- demonstrate safe, reliable operation of the system under operational conditions.

The learning objectives that followed on from these were:

- How well can the battery perform each use case and what is the impact on the network?
- Are there any seasonal variations in the battery performance and costs?
- What is the financial benefit of each use case?
- Which use cases can be combined effectively and what are the combined financial benefits?
- What are the practical issues that investors and battery operators need to consider?
- Are there barriers inherent within the regulatory arrangements that would prevent the investment in co-locating batteries with storage?
- How do the economic factors affect the battery sizing?

2.5 Battery specification

The factors affecting battery sizing were potential impact and cost. Anything smaller than 300kW would not make a measurable difference to voltage levels on an 11kV network. A prototype on this scale was considered desirable to give industry stakeholders sufficient confidence for a larger roll out to be possible. The battery capacity (731kWh at 0-100%) was such that it could run at full power for two hours as this is more likely to represent the type of usage by DNOs for peak lopping applications.

No particular battery chemistry was specified as a requirement of the procurement process. The battery chemistry of the successful bid was lithium iron phosphate.

3. Project Phases

The project can be divided into the following phases:

- Design and procurement.
- Construction.
- Testing.
- Data analysis.
- Decommissioning / transfer.

The following sections describe the activities and learning from each of these phases, with the exception of the learning from analysing the use cases themselves which is provided separately in sections 10 to 17 of this report.

4. Design & Procurement

4.1 Design & procurement phase activities

British Solar Renewables Limited (BSR), who were at the time owner operators of the solar park at Higher Hill Farm and had an interest in investigating options for storage, were already confirmed as project partners at the time of project initiation. The battery was procured via a competitive tender process to ensure value for money for customers. There were four submissions for the tender, and after evaluation RES was the successful bidder.

The process of obtaining planning was relatively onerous and non-material amendments to the planning permission were required when the fenced area was altered due to site conditions. The total area of the enclosure was reduced to allow for improved access across the BSR site without compromising vehicle access to the battery itself.

The design sign-off was a two-stage process that covered the battery itself followed by the balance of plant.

Examples of the drawings from the design process are given below in Figure 3: Example Design Drawings. It can be seen that the container was divided into two compartments for safety reasons, such that the battery operator was separated from the battery itself and the fire suppression system. The drawings also show that only part of the usable space within the battery compartment was used and that it would have be possible to approximately double the battery capacity if desired. The capacity of the project only required a standard 20ft container, but a 40ft container was available with significantly less lead time at a similar price which is why there was additional space inside.

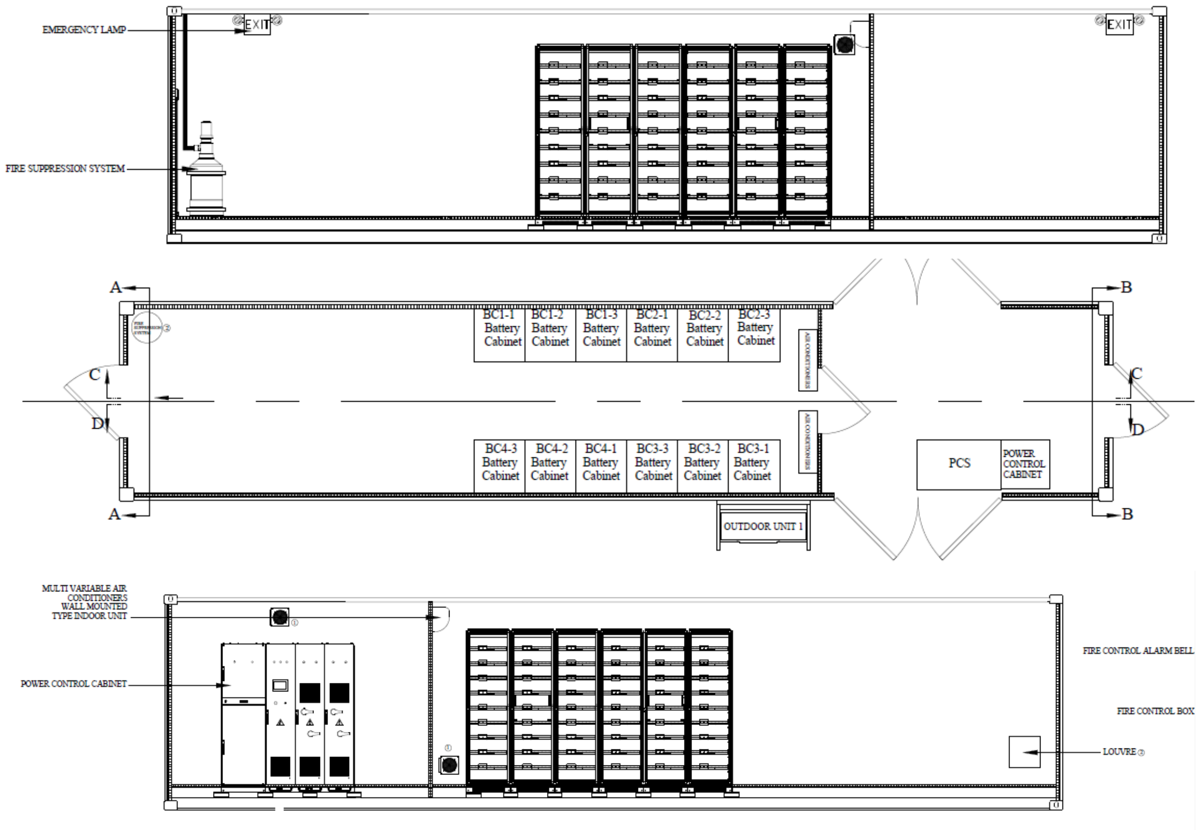

Figure 3: Example Design Drawings

The design phase also included negotiating a lease for the battery site. This proved to be a far more lengthy and complex process than had been anticipated for this research project. In order to avoid lengthy delays preventing installation, a 'Licence to Occupy' was used as a temporary measure until the lease negotiations were complete. Timescale pressure of this nature was due to the nature of the research project, and it is not believed that this would occur for a standard commercial project.

4.2 Design & procurement phase learning

The learning points from this phase are summarised below.

- The use of a partner to assist with the procurement of the battery was essential as DNO staff were not yet sufficiently familiar enough with battery technology to carry out procurement unaided;
- Including more flexibility in the Statements of Works would have avoided the significant work of updating the documents and getting the updates signed off;
- Having as much access to technical detail as possible during the procurement stage is beneficial;
- The contractual conditions covering the battery operation should have included a clause concerning the imbalance between strings. It appears this is a standard clause in other battery contracts; and

- Identify any issues with the contractual limits for items such as power factor early on. This required modelling by Western Power Distribution (WPD) staff to ensure that if the algorithm were to fail to operate correctly, that the reactive power element would not cause network issues. The selection of a clean feeder for the trial has limited the potential impact on other customers from voltage fluctuations during testing.
- While the process to negotiate a lease started very early in the project, this aspect took far longer than anticipated. Future projects might benefit from the use of template legal documents to flush out potential issues at the feasibility stage.

5. Construction

5.1 Construction phase activities

Construction was completed in October 2016, with the exception of a couple of minor snagging items which were resolved within three months. Initial values from the commissioning tests are included in section 8.3 Pre-sale battery performance testing. A further description of the commissioning tests is given in Appendix D Commissioning Test Learning Summary.

Issues encountered during the construction phase included;

1) location of cables differing from plans;
2) damage to communication cables during the erection of fencing; and
3) the requirement for a specialist driver to transport the battery due to its hazard rating.

The following photographs show some key stages of the construction and the battery internals.

Figure 4: Battery Arrival

Figure 5: Battery Offloading

Figure 6: Battery on plinth before and after fencing

Figure 7: Secondary Access Door (Battery Compartment)

Figure 8: Local Control Panel

Figure 9: Circuit Breakers and Emergency Stop

Figure 11: Fire Suppression Equipment

Figure 10: Battery Strings

The additional external inter-trip signal was originally going to be provided by a standard WPD inter-trip device. The lead time and cost of this was relatively high, and it was agreed to replace the signal generated from the modified soft inter-trip panel with a surrogate signal. The surrogate signal was generated by an enhanced Raspberry Pi device built by a BSR member of staff which was equally valid for the testing

Argand Solutions was commissioned to install some advanced power quality monitoring equipment at Millfield Primary. This served two purposes, firstly it provided high resolution data to support detailed analysis of the impact of the battery operation as measured at the primary substation, but it also allowed for WPD to assess the potential benefit of a different type of power quality monitoring solution. To enable high quality monitoring, a specialist current transformer was purchased due to space limitations around the existing Millfield switchgear.

Figure 12: Power quality monitoring and communications equipment installed at Millfield Primary

An overview of the Argand solution is given in Figure 13. While traditional power quality monitoring relies on analysis software and hardware provided by the same supplier being used together, the solution implemented allows greater flexibility by making the data available to other software and other purposes via a cloud server.

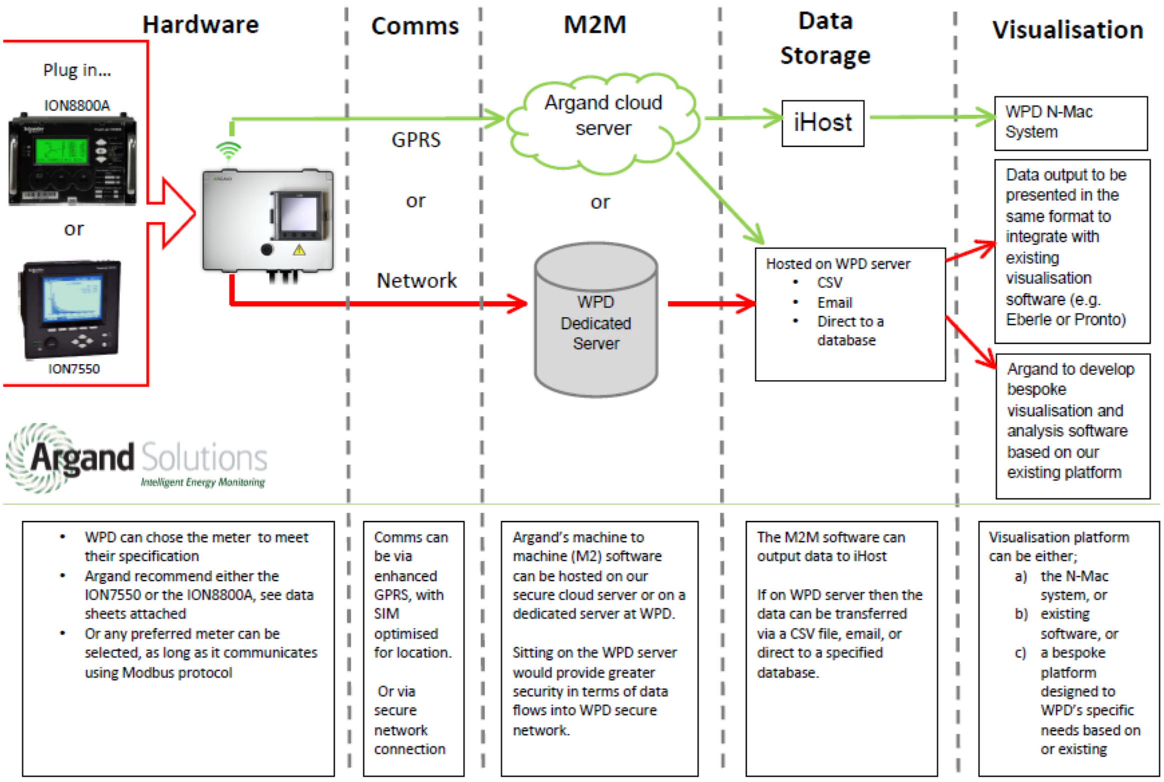

Figure 13: Argand Power Quality Monitoring system architecture.

5.2 Construction phase learning

The key learning points from the construction phase are;

- ensure that legal issues are resolved early in the construction schedule;
- expect a degree of inaccuracy in plans; and
- power quality monitoring was something of an afterthought to the project, and procurement and installation took longer than expected.

6. Testing

6.1 Testing phase activities

A test schedule was produced that would ensure that each use case would be tested for sufficient iterations in different seasons both individually and combined with other appropriate use cases. The test schedule was originally drawn up with a significant amount of days where the battery was not scheduled to operate, allowing time to resolve minor issues and giving an opportunity for early analysis of the data to inform the rest of the testing schedule. Due to the high number of technical issues that were experienced, these non-operational days provided a useful contingency allowing the test schedule to be re-planned several times during the project.

During the testing phase there were also a number of site visits from parties interested in the battery technology. These included Innogy and the local MP. Innogy are a renewables subsidiary of the German energy company RWE. The collaborative visit shows the ability of innovation projects to share learning between companies that are competitors.

MP is energised by the potential of firm's battery project

JAMES Heappey MP paid 'ambitious' British Solar Renewables (BSR) a visit and learn about its innovative battery storage project.

The company, which is now based in Shepton Mallet, has developed one of the first industrial-scale battery storage facilities in the UK.

The project is in collaboration with Western Power Distribution. Mr Heappey said: "It is great to see such an ambitious and quickly growing company relocated to Shepton Mallet. Clean technologies are the heart of the Government's industrial strategy and companies like BSR are leading the way in delivering clean growth for our economy.

"British Solar Renewables' partnership with WPD is a great example of how the energy industry is working together to understand this technology and what it can do for our energy system and how they can be used to ensure security and supply to drive down prices for consumers."

Managing Director at BSR Graham Harding, said: "It was fantastic to hear James Heappey's enthusiasm for our services to the low-carbon industry.

"James was particularly interested in the energy storage research we're undertaking with WPD. And rightly so – this study is providing fascinating insight into how batteries can benefit the UK's power supply as the grid decarbonises and becomes increasingly decentralised."

Figure 14: Press coverage for MP visit

Figure 15: James Heappey MP visiting

6.1.1 Technical issues encountered

The commercial battery industry was much less developed or advanced in 2016, compared to the end of the project. While the hardware itself was relatively similar (although newer cells have greater energy density), the most recent major advances have been in software. Monitoring of State of Charge (SoC), communications between different manufacturers equipment, advanced and multi-use control systems are significantly more common now.

The battery received numerous software and occasional hardware upgrades throughout the test process, as flaws and edge-conditions were discovered as the tests became more advanced. One of the earliest challenges was that the battery strings appeared to be becoming unbalanced, a single charge-discharge cycle could report imbalance by up to 2%. After several tests the resulting imbalance could reach up to 15%, although this didn't appear to directly affect any operations of the battery. It did, however, limit how far the user was prepared to charge the battery as at a 90% overall SoC, at 15% imbalance one string would appear to be at 97.5% SoC, which is considered high within the industry. A similar problem would occur at low SoC. In retrospect, this issue was more of a concern for the operators of the battery rather than being a direct problem or fault, but it caused delays and significantly more cautious operation of the battery than anticipated.

BYD, the battery manufacturers, stated that the cell strings themselves were not out of balance, but rather that it was merely an artefact of the State of Charge calculation algorithm. They were able to install a software update which rectified the issue and, from that point onward, all strings remained completely balanced, without even a percentage difference between them.

The RESolve State of Charge manager was not properly functioning at the beginning of the project due to an unset parameter. The issue manifested itself during the start-up of the battery when various subsystems falsely reported erroneous state of charge readings at different points, which RESolve would then try to act upon. To avoid this issue, the manager start-up has a delay parameter (usually set to 30-60 seconds). In this case the value had been set to zero, effectively infinite, meaning it never started. Even though there was a simple fix, the problem was not spotted immediately due to the imbalanced string readings. This meant that each test had to be carried out carefully, trying to ensure that no schedule would excessively charge or discharge the battery as the programming issue meant that RESolve would stop this.

During the early testing there was an issue with the battery discharging at midnight, despite no schedules being set. This was reported as an unusual occurrence, with an ‘invisible’ export schedule set for every day at midnight. This ghost schedule was deleted and caused no more problems. This issue has not been encountered with RESolve on any other site and it did not manifest itself again during the rest of the project. This delayed testing by approximately a week.

A very short-term problem appeared to be that RESolve was ignoring an import instruction. There was a missing parameter in a particular import schedule, even though this should have been impossible. The parameter was set, and the schedule functioned normally after this. This caused test failures only if tests were carried out on the day that that particular schedule was used.

RESolve used the BYD Master State of Charge output for most of its calculations. A fault developed where RESolve read the true SoC, then 50% and 25% of the SoC, every few seconds. To avoid any further testing delays, RES altered their software so that it read the SoC calculated by the Battery Management System (BMS) instead. This fix remained in place until the completion of the testing programme and it remains unknown whether this issue with the BYD Master was properly fixed. This issue caused a three-day delay.

Each of the control methods in RESolve were custom-written for the project. The standard use case for RESolve was previously operating one service, so the use of so many advanced schedules, some of which took inputs from local sensors and set a response set point, was a significant undertaking. Some of the schedules had small software issues which would have been discovered and fixed quickly in a fully commercial battery but were left unnoticed for a significant time due to not being required in the test schedule until later.

The solar peak lopping algorithm initially produced exactly 50% of the required response from the battery. This was easily fixed with an extra term within the internal algorithm. Later in the process, the response of the algorithm was also improved with some optimisations aimed at speed of response. This was found before testing started in earnest, so had minimal impact

When using reactive power control in a combination method, the system erroneously ignores the level of reactive power limits set, responding only to the inverter limit. This issue was not noticed until the very last stage of testing as the reactive limits had not been used before. This impacted the effectiveness of the combination methods, as RESolve prioritises reactive power modes over active power. However, it is expected that this issue is easily fixed, and the maths behind the reactive-active power output of inverters is well understood, so it is not considered a major impediment to the usefulness of the project.

There was an issue with the connection agreement for the Copley Wood battery. The connection agreement prohibited operating outside the standard power factors, however the WPD test schedule required reactive power testing. This was resolved after discussions with WPD's network management team, which ran additional simulations on the local network to investigate the impact of the reactive power flows. They granted specific limited permission to operate outside of the confines of the existing connection agreement, but only for the purposes of the reactive power tests. This delayed reactive power testing by three weeks, although this time was used for other active power tests.

A major problem that prevented successful testing was an air conditioning unit failure. This allowed the system to heat up to its official limit, 40°C, after which the inverters de-rated to prevent further detrimental effects due to excessive heat. There were two air conditioning units, but only one external condenser. This single point of failure, the condenser, is where the issue occurred. Visibility of this issue was hampered by a lack of dedicated SCADA communications between the air conditioning units and RESolve, in addition high and low temperature alarms were not set. The battery at this point was operating at 1.5-2 cycles per day, a heavy work load, so the units were replaced with larger systems that had separate external condensers. RES installed temporary air conditioning to reduce down time to 6 days.

Another single point of failure became apparent with the design of the system, in that only one remote control unit was provided for the two air conditioning devices. The loss of the remote control caused further delays whereas providing an additional remote control would not have been expensive.

The combination method had a relatively naive way of calculating the combined setpoint. Regardless of which active power use cases were selected, RESolve would sum the setpoints to tell the battery what to do. While this initially sounds logical, it causes problems when trying to stay within network limits, as each control mode was unaware what the other control mode was doing. The most obvious occurrence was this was when combining ramp rate control and solar peak lopping. When the solar generation was below the peak limit (e.g. 800kW) the ramp rate control worked as expected and the peak lopping algorithm was idle. However, if the generation was above the limit, at 950kW, the peak lopping mode would maintain the export at 800kW. If the generation dropped from 950kW to 850kW quickly, the output should still be stable at 800kW with an associated reduction in import from the battery. However, the ramp rate control algorithm responded to the drop in generation, increasing export to try and reduce the ramp, **even though** the actual site export was already being held stable. The site export then increased above the peak limit, which in the real world would be a breach of export limits. This suggests that a more complex combination method is required, with additional safeguards and potentially allowing the control modes to become 'aware' of each other's actions.

6.2 Testing phase learning

The number of technical issues experienced was unexpected and initially the process to report and resolve these issues was ad-hoc. When initially reporting the peak lopping issue, for example, a screenshot of RESolve and an email would be sent to RES from the BSR operator. There would then be some follow up emails and phone discussions if the issue wasn't obvious, or didn't have a simple solution. In addition, initially all issues were sent directly to the project manager, who then had to pass them on to the software engineers. While issue reporting had been anticipated, the process to share the issues encountered was improved during the testing phase and this led to improved issue resolution. This included emails sent to both the project manager and the software engineers at the same time, so the project manager was aware there was a problem and the software engineers could begin investigating immediately

Realistically, if the battery was installed today it is likely most of the issues would not be encountered. Advances in control systems and wider adoption of energy storage has driven improvements in reliability. Running several brand-new algorithms on a control system for a research project meant that some minor issues were inevitable, and it reflects more on the nascent nature of the industry rather than a specific failing of any component. However, one of the key learning points was linking the air conditioning system to RESolve so that an alarm could be raised if there was an issue. Under the original configuration, the only way to detect an over-temperature de-rating was to observe the battery and ensure it was following its setpoint appropriately. It wasn't possible within the budget and timeframe of the project to install SCADA communication between the air conditioning units and RESolve, but high and low temperature alarms were set up for the cells after discussions between BSR and RES, providing a fast, low-cost solution. Alarms such as this should be standard on all new systems, as should an alert function when the system has de-rated itself. This would make it significantly easier to remotely identify problems and minimise downtime, which, if the battery was to be used for DSO services, would be essential.

The problem with over-heating (and the resulting issues with the air conditioning) was exacerbated by the choice of paint colour for the battery container. While the dark green colour was intended to minimise the visual impact of the battery container a lighter colour would have reduced the solar heat gain, thus reducing the load on the air conditioning. One option that could have been a compromise between visual impact and heat gain could have been to only paint public facing faces of the battery containment green and paint the faces that are not visible to the public (i.e. top, rear and furthest end) in white.

The single point of failure with the air conditioning remote control could have been prevented by providing multiple units and tethering the handset to the battery container to prevent its removal.

7. Analysis

7.1 Analysis phase activities

The original project plan was to carry out much more of the analysis alongside the battery testing, making use of the planned battery down-time. However due to the technical issues the analysis of the data did not start until the battery operation was largely completed.

The analysis activity involved downloading the operational data from the battery and completing analysis on the technical performance and economic value of the services.

The analysis phase also included the calculation of the net impact of the solar storage project on the revenues of the solar park, i.e. the additional costs of charging the battery from either the grid or the solar park output, less any additional value gained from the energy released from the battery as it discharged, either to supply the solar park or to export to the grid.

7.1.1 Data sources for analysis

The main resource for the analysis was the data recorded in the RESolve SCADA system. This system recorded all of the schedules programmed for the battery, and all parameters required for the performance analysis (power in and out, battery setpoint etc). The SCADA also receives information from two Power Quality Meters (PQMs), one for the battery and one for the solar park. In this way it is possible to separate the effects each asset has on the network. These can include harmonic disturbances, investigation of sources of power flow, voltage fluctuations and impacts on the power factor.

A secondary source of data was the Argand power quality monitoring device installed at the Millfield substation. This monitors the 11kV feeder that was connected to the solar park and battery, as well as one other solar park (Higher Hill). This allows us to independently verify any effect (or lack thereof) that the various use cases have on the network, rather than solely relying on the instruments installed on the battery site.

To help with the financial analysis, the details of the import agreement for the solar park have been obtained. Due to commercial interests, exact figures could not be provided, though close approximations were given by the energy supplier. The export power purchase agreement (PPA) has not been altered, and is based on the N2EX next day auctions, which have an hourly price.

For comparison, and with the knowledge that were this battery a commercial concern the import/export agreements would be altered, the half hourly system price of power (also known as the wholesale price) has also been used. This is the price that energy is traded at on the energy market, although renewable energy doesn't actively take part in this trading. The wholesale price cannot be directly accessed by an asset: it is the pure price without supplier margins or network costs added. In this way energy can be traded in a location-agnostic way. PPAs which follow this price 'live' are available, although the standard is to fix at a more consistent and predictable pricing structure. It is expected that the system price plus network costs and supplier percentage would be a more accurate import charge for an energy storage system, while solar parks usually have simpler (but costlier) arrangements as they import so little. For example, Copley Wood solar park had a flat set-level day-night tariff, with adjustments for summer and winter.

7.1.2 Data cleansing

The RESolve SCADA system used in the project had an unexpected method of recording data which was used to conserve space. The system checked each sensor and compared the value to the last recorded value. If nothing had changed, no value or timestamp was recorded. This meant that when downloading the data, there were large gaps of time with no values recorded. In addition, when downloading months of data at a time, the 1 second data quickly grew beyond the limits of Microsoft Excel and became unmanageable. A Python script was written to extract the data and place it in a form that could be easily analysed. This software translated standard calendar dates into Epoch-milliseconds, which is what was used by the server[2]. It then only extracted data that fell between the user set dates, and only for the data that matched a pre-written list, rather than downloading the 200+ system values that are recorded. The software also extracted the data for the previous day, so that values which had not changed on the current day could be forward filled from the previous day. This, in nearly all cases, allowed a complete dataset to be extracted. The software also averaged each half hour as a separate comma de-limited file (CSV) so that the data could be visualised by Excel.

7.1.3 Calculation methodology for financial gain / loss

Each use case has been compared against the 'base-case' for energy storage in the UK, which was the Fast Frequency Response (FFR) and capacity market contracts. The values for these in 2016 (when the battery was constructed) and 2018 (when the testing ended) were calculated and given as a simple £/hour for a battery of this size and capacity.

This gives a single value against which all the other cases can be measured, as well as an opportunity cost which can be reduced as required (e.g. if a use case required half of the

[2] Epoch-milliseconds are the number of milliseconds that have elapsed since midnight on January 1st, 1970. This is how several operating systems keep time, including Unix/Linux variants and macOS, and many programming languages. For example, the Epoch Millisecond date stamp 1514764800 is equivalent to 01/01/2018 at midnight, in Greenwich Mean Time. The software conversion made the data extraction process significantly easier.

power output, the battery can still bid half of its capacity into the FFR auction and receive half the standard income). It has been assumed throughout the analysis that the battery can receive FFR contracts whenever it isn't operating in other modes, although changing market conditions mean this is no longer necessarily true.

The amount of income that the battery will miss out on can be used for investigating if DSO network services (such as network demand peak lopping) are commercially viable. The cost/benefit of providing this service, plus any administration costs and required profit levels, must be above the threshold that the energy system can generate by staying within the FFR base case.

7.1.4 Calculation methodology for round trip efficiency / battery efficiency

The commissioning efficiency tests involved charging or discharging the battery to 25% SoC, then charging for a set amount of time at a target rate. This can be seen in more detail in Table 2: Efficiency test charging durations. All these tests except the one at 310kW were carried out during commissioning. The battery was then discharged to 25% again, and the kWh used to charge the battery are compared to the kWh exported while discharging the battery. Charge power divided by discharge power gives the effective round-trip efficiency. These values were read from the PQMs, which were installed on the HV side of the transformer. This gives a real-world round-trip figure, including transformer losses, which is the figure that investors and developers will be most interested in.

7.1.5 Testing for seasonal impacts

The original test plan included testing of all the different use cases with an equal spread throughout the year. The impact of the technical issues meant that the majority of testing was carried out in the Autumn/Winter, although there were some successful tests in the summer. This was not considered to be a problem, as variables such as temperature, solar production and network load have been monitored throughout the year, and tests have been completed under all weather conditions. Thus, tests can be theoretically run based on the existing data, using real world battery behaviour to inform the results. In addition, battery operation (other than auxiliary loads) appears to be similar regardless of the season.

7.1.6 Base case for comparison

One of the initial challenges of financial analysis of the different use cases, was creating a base income for them to be compared against. Without this anchoring the analysis, it is difficult to draw any meaningful conclusions about the viability of the operating modes. Peak lopping could have twice the value of arbitrage, but that wouldn't necessarily make it profitable.

To provide the base income, the current income streams of energy storage were investigated. The standard business case for most commercial batteries currently relies on the National Grid FFR contracts and the Capacity Market Auction. These combined created the investment case for batteries - however there has been a significant shift in the market since this project was envisaged and installed. FFR rates have plunged from over 20 £/MW/h (pounds per MW per hour of service) to less than 10 £/MW/h between 2016 and

2018, while the capacity market has seen a similar reduction. Further to this, the capacity market payments are now de-rated based on the number of hours the installation can discharge at maximum capacity for, further reducing payments.

To highlight the difference between when the battery was built and current market conditions, two base incomes have been calculated. One uses 2016 FFR and Capacity Market rates, free from de-rating, while the other reflects the latest FFR and Capacity Market auction incomes from 2018. When scaled to the size of the Copley Wood battery, the 2016 figures equate to £7.31/hour, while current rates provide an income of £2.79/hour for being available in these markets. The calculations for this can be found in the virtual PPA appendix B, and are based on what was considered the usual commercial business case for energy storage at the time that the battery was operating (i.e. 24/7 FFR contracts and Capacity Market revenue). These are considered the hurdle rates which must be met by DSO/ third party contracts to incentivise battery operators to make their business models more complex.

These two figures give a baseline number with which to compare the cost of operation and expected payment price for each use case. By providing an hourly figure, use cases such as peak lopping, that take several hours during the day, can have a 'lost opportunity' cost associated with them, demonstrating whether they are economically viable or if there are more profitable options available to a battery asset.

7.2 Analysis phase learning

This section only includes the learning that was gained in the process of analysing the data. For the results of the analysis of the individual use cases please see sections 10 to 16.

The format of the data storage within RESolve, while reducing the volume of data recorded made data transfer and analysis more complex.

It is difficult to provide business case analysis when prices are changing rapidly which is both the case for batteries but also for the value of services that can be provided. It is also very difficult to estimate the future values of services to DNOs or third parties that are very specific to a location. The services may differ in value or not be required at all at different locations or at the same location at different times.

Another learning point related to the various metering points that were used to calculate the net impact on the solar site revenues. It had been assumed that the various meters, though installed at different times by different parties, would be synchronous. It became clear from the data that the values for the total import and total export of energy were not summing correctly. Ignoring the impact of losses, it was expected that the total power through the site meter would equal the sum of the import by the battery and the solar park. This value would often have a large error in one time period followed by a large error in the opposite direction in the next time period suggesting that the meter synchronisation was causing problems.

8. Decommissioning

8.1 Decommissioning phase activities

The decommissioning phase of the project was delayed due to the testing phase taking longer than anticipated. The lease for the battery site included stipulations about the site being restored to its prior condition and therefore it was necessary to understand the work required to remove the battery and restore the site. There was provision for the battery to be sold to the owner of the solar park, negating the need to restore the site. However, due to the planned sale of the solar park it was preferable to be able to transfer the entire site unencumbered from other leasehold arrangements.

Much of the work in the decommissioning phase related to managing the tender process to sell and relocate the battery which are covered in a separate report[3]. This was a 'first of its kind' activity which allowed for greater learning about the commercial and technical issues surrounding the sale and removal of a containerised battery.

The project team supported the sale process by providing interested parties with information about the batteries performance and having an "open day" for bidders to see the battery and assess the site for removal and reinstatement works.

The following section describes the process completed to test the battery condition at the end of the project.

8.2 Capacity fade / age-related battery degradation.

The project lifetime was significantly shorter than that of expected commercial installations (~2 years versus 10 years), giving less time for the energy capacity to be diminished. Commercial installations are governed by their warranties, which tend to be structured such that they guarantee a residual capacity at the end of the warranted life, which is commonly before the expected end of the project life, so developers will plan for replacement of all cells. It is unusual that a battery would reach the end of its project life before the end of its usable life, as this would suggest an over engineered (and therefore overpriced) project, leading to lower project returns than could have been achieved.

However, in this case the battery was specifically being used for a research project and exposed to an unusual operating regime that had never been tested on a battery before. In addition, the testing requirements changed over the life of the project, making the usage impossible to simulate or predict at the beginning of the project.

This, coupled with the desire to sell the battery after the research was complete, meant that a complete health check-up was required. This would be valuable in reassuring potential purchasers that the systems were still functioning correctly and provide a valuable insight into the ways the new use cases affected an energy storage asset.

[3] https://www.westernpower.co.uk/innovation/documents Solar Storage - Battery Disposal Report - January 2019

8.3 Pre-sale battery performance testing

When the battery was installed it had undergone a series of commissioning tests, ensuring the system was working as intended and meeting all its contractual obligations. These included a full capacity test and a suite of roundtrip efficiency tests, demonstrating the performance across a range of charging rates.

It was decided that these tests would be suitable for the end of project tests, as direct comparisons could be drawn with the asset as originally installed, and any differences or degradation should then be attributable to the testing regime.

The efficiency tests involved charging and discharging at the same rate and calculating the difference between the imported and exported energies. Each test was started at 25% SoC to ensure that the testing wouldn't be interrupted by the State of Charge Manager. The battery then imported for a set time, then exported until the SoC read 25% again. The times for charging are listed in Table 2: Efficiency test charging durations.

Efficiency Testing Plan	
Rate of charge/discharge (kW)	**Time for charge (mins)**
310	60
300	60
200	90
150	120
75	240
30	600

Table 2: Efficiency test charging durations

The commissioning tests were completed on-site, while the end of project tests were conducted remotely, using the kWh import and export readings of the battery power quality meter to record the energy.

8.4 Decommissioning phase learning

8.4.1 Pre-sale battery performance test results

The initial results are taken from the commissioning paperwork and are shown in Table 3.The efficiency rate of the battery at 30kW was not recorded, but it was agreed that it was higher than the contractual hurdle rate of 32%. All other percentages were recorded. The battery was not tested at 310kW in the commissioning tests as officially it was a 300kW battery that had been specified. The 310kW inverters were oversized to ensure sufficient capacity. However, throughout the project the battery has frequently been used at 310 kW and so this was included in the final tests.

The battery was well in excess of its contractual efficiency obligations and had a surprisingly high efficiency at 200kW, compared with the other efficiencies recorded. While lithium batteries can technically reach this level of round trip efficiency, it is usually not achieved on grid-connected batteries due to auxiliary loads. It is possible that the storage system had no need for engaging the cooling systems during this test, giving an inflated efficiency score. The capacity recorded at commissioning was 702kWh on a discharge from 98% to 2%, giving a calculated capacity of 731kWh for a 100% to 0% discharge.

The end of project tests are listed in Table 4, demonstrating a very reliable efficiency of approximately 88% over a wide range of charging rates. The efficiency only drops off over the 30kW charging rate, at which point, with a charge-discharge time of 20 hours, the parasitic loads have a greater impact.

End of Project Tests										
				Import meter readings		Export meter readings		Calculated values		
Rate of charge (kW)	Start SoC (%)	Max SoC (%)	End SoC (%)	total kWh start	total kWh end	total kWh start	total kWh end	kWh imported	kWh exported	Roundtrip efficiency (%)
310	25	64	25	251078	251388	208371	208645	310	274	88.39%
300	24	62	25	249069	249369	206943	207206	300	263	87.67%
200	24	63	25	249392	249692	207233	207498	300	265	88.33%
150	24	63	25	250774	251074	208101	208367	300	266	88.67%
75	25	64	25	250451	250751	207809	208075	300	266	88.67%
30	24	61	25	250124	250424	207539	207782	300	243	81.00%
Capacity Test	99	N/A	1	N/A	N/A	208969	209632	N/A	663	N/A
Calculated 0-100% Capacity (kWh):									676.53	

Table 3: End of Project Round Trip Efficiencies

All tests were conducted at unity power factor for import and export.

The more consistent test results are potentially due to a change in the batteries hardware partway through the project an air conditioning unit failure resulted in the replacement of both units with those of a different manufacturer, with additional systems put in place to allow one to operate in the event of total failure of the second. This greater redundancy may have resulted in a more consistent power draw for cooling, rather than the longer, more intense on-off cycle of the previous units. In addition to this, the settings for target temperatures have been changed on the project twice to optimise the internal temperatures.

The capacity of the battery was still well above its nameplate capacity of 658kWh, although the heavy-duty cycle of the testing has reduced it from 731kWh to 677kWh. From this the capacity fade[4] can be calculated to be 7.4%, over 320-350 cycles (depending if calculated using starting or final capacity). This can seem relatively high compared to current commercial models and predictions, but the battery has been charging to 90% and discharging to 10%, sometimes as often as two cycles per day. Whereas commercial batteries which engage in the FFR market usually operate at approximately 50% state of

[4] (1-(new capacity / original capacity)) x 100

charge, with small deviations as they are called on, which is significantly less harmful to the battery.

8.4.2 Conclusions

It is not easy to directly compare the results of the re-run of the commissioning tests[5] with the initial commissioning test results due to the interim change of hardware. However, it was clear that the battery was performing consistently and efficiently, with no significant problems in absorbing or releasing energy. The calculated capacity fade is not unexpected from a two-year-old asset, especially given the tougher services it has been providing vs current commercial energy storage systems, and it was still above its initial nameplate capacity. The inverters were able to produce/absorb power at their peak ratings and showed no signs of faults or thermal derating. The replacement air conditioning systems were keeping the cells within their thermal parameters, aiding efficiency.

It did not appear that the use cases tested had caused excessive capacity fade from the lithium iron phosphate cells, albeit degradation was considered to be at a higher rate than that expected from the delivery of services such as frequency response.

9. Use Cases

Having described the learning from the various phases of the project, the next section of this report addresses the results and learning from the various use cases that were tested during the project. These use cases demonstrate the different ways in which the battery can be used to provide value whether the beneficiary is the battery owner, solar park operator, DNO or another third party. They are not mutually exclusive and multiple use cases can be carried out at the same time (as long as they all require either charging or discharging). The use cases are summarised in this section before each use case is examined in more detail, analysing the results from the battery testing over the various days of the week and seasons.

9.1.1 Use case 1 - Arbitrage

Originally viewed as the ‘default’ use case, arbitrage is the simplest scenario. Buying power when its cheap and selling it when prices are high is an income stream available to any storage asset anywhere on the network, independent of location. However, co-located assets may not be able to access the more volatile power markets. The battery owner would be the direct beneficiary of this battery service. It is included to test the profitability of this operating mode and have a baseline to compare the other use cases with.

[5] Tests completed on the 30/7/2018 as part of the end of project/pre-sale performance testing

9.1.2 Use case 2 – Peak demand limiting

This use case is designed to reduce the maximum demand 'seen' by the DNO's network. The theory is that if the peak load is reduced, it can delay required reinforcement for network assets, saving customers money. This will only delay replacement in areas where thermal ratings, rather than voltage or fault level limits, are the reasons for upgrades. This can only apply if the battery is connected in a way that is electrically relevant to the DNO asset that is overloaded. i.e. the battery connected at Higher Hill could have helped reduce the load at the primary transformer but would have had little impact if another 11kV feeder were overloaded because the flow of energy from the battery would be the same as the usual flow of energy from the primary transformer.

9.1.3 Use case 3 - Local demand profile matching

Similar to use case 2, this relies on the battery being connected downstream of a localised network restriction. For example, if a factory downstream of a substation wanted to expand operations but in doing so would require a complete transformer upgrade, than installing a battery on the same feeder and linking via telecommunications to the factory would allow the battery to export whenever the demand increased, keeping the load on the primary below the transformer limits.

9.1.4 Use case 4 – Low demand grid voltage support

Rural DNO networks have a large number of feeders that are lightly loaded overnight. This can cause a voltage rise, which the battery can help reduce by loading the line on demand. The effectiveness of this depends on the size of the battery and the nominal voltage of the line.

9.1.5 Use case 5 - Voltage control by reactive power

The network experiences high voltages on some lines during the day, due to the penetration of large grid connected solar sites and other distributed renewables. In many more rural areas, this voltage rise has reached the point where any additional generation potentially can make the feeders operate outside of their statutory limits. By using the battery inverters as a STATCOM (static synchronous compensator, a regulating device used to generate or absorb reactive power), they are capable of importing or exporting reactive power as a way to raise or lower the network voltage. This could allow more generation to be connected and/or generators with an Active Network Management (ANM) connection to be constrained less often.

9.1.6 Use case 6 – PV export limiting

A significant percentage of the DNO's network, especially in the south of the UK, is now reaching its generation capacity limits. In addition, distributed generation needs to be above a certain minimum size to allow economies of scale to take effect, especially for new parity projects. By peak lopping generation, a larger capacity solar park can be installed on a smaller capacity grid connection, with export controlled by the battery to remain within connection limits. Any excess generation can then be exported later in the evening after the solar park has finished generating, benefiting from the higher export prices (times of peak demand). The battery could be dedicated to peak lopping during the day, unless there was another service that was more lucrative than the loss of generation.

9.1.7 Use case 7 – Variable PV export limiting

This is similar in function to the use case 6 but is based on a flexible ceiling caused by temporary network conditions or outages. These reduce the standard capacity of the network and thus the output of the generator needs to be reduced. However, it has been treated as functionally the same as use case 6, because the solar park is unconstrained, so any peak lopping is automatically a changeable glass ceiling set by the user for the test.

9.1.8 Use Case 8 – PV power quality improvement (ramp rate control)

Solar generation can 'ramp' up and down almost instantaneously, dependent on cloud cover. This can cause disturbances on the network, with power flows changing direction quickly and tap changers working more often to maintain statutory voltage conditions. The battery should be able to smooth out these peaks and troughs, to maintain a steadier output and reduce the swing effects on the network.

9.1.9 Use Case 9 - Multiple storage system control

This use case suggests the potential benefits from having multiple energy storage systems in the same area of the network that co-ordinate and work together. This could bring greater benefits by optimising over a greater aggregated battery capacity. There would be the potential for problems where multiple storage systems operate independently e.g. if two energy systems are monitoring the same variable (e.g. voltage) independently, they would both respond slightly differently and could end up opposing each other while trying to stabilise the network. However, grid services (and DSO services looking forward) are likely to have a set-point signalled from a centrally monitored location, rather than batteries independently making decisions based on internal logic. In this way, the Network/ DSO can avoid any issues with different software implementations and maintain reliability. With only one battery installation, it was not possible to test this use case during this project, but has been included as another potential means of providing additional revenue in the future.

10. Results - Use Case 1 - Arbitrage

10.1 Introduction

Arbitrage was viewed at the outset of the project as the 'default' use case of the battery. The logic was that whenever no other service was required, or the battery was out of contract, the inter-day price variation would be a standby source of income. This use case was considered 'safe' as it wasn't reliant on contracts with any third parties, and the theory of buying power at a low price and selling high was simple and instantly understandable, compared to the more complex use cases. Coupled with the increasingly variable system prices seen every winter, it seemed the logical fall-back operational mode.

However, the early assumption that there would be arbitrage opportunities every day was proven to be unrealistic. While the system (wholesale) price varies considerably within each 24-hour period, this makes up less than half of the total import cost of power delivered to the battery. Use of System charges, climate change levies, energy supplier margins etc. all increase the price of power, impacting the potential revenue for the owner/operator. The cost of using the battery for arbitrage also has to be quantified, as other use cases could be a more profitable use of time.

10.2 Method

As the project had no price-prediction algorithm available to it, or an energy trading floor, the ability to capture fluctuating system prices was limited. In addition, as the battery was set up to mainly to operate on pre-programmed schedules rather than responding to external signals, it was decided that specific 'windows' of import and export should be chosen. RESolve was able to respond to external signals in a specific schedule, and RES also offered an option to include an 'arbitrage mode' which takes price feed information provided by the owner and imports and exports whenever the price hits predetermined limits. However, this functionality was not included in the trial project.

There is a huge variation each day in the system price of energy, with a wide range of variables combining in random ways to influence the price. Certain days may see a baseload generator go down for unscheduled maintenance, or an excess of wind during a time of low demand. Coupled with this, the time of use charges - Distribution Use of System (DUoS), means that while the wholesale price could be low, it is still expensive to use energy. The inverse of this is that if it is possible to export during the network peak times, the battery can benefit from additional revenue.

However, in general, prices overnight are cheaper than during the day, and the UK has a peak demand between 5-7pm. In addition, between 7-9am demand rapidly increases as the population prepares for work, requiring power stations to ramp up quickly and inefficiently, thus pushing prices up. Therefore, the peak prices each day are typically between 7-9am and 5-7pm. This is borne out by the 'average day' of 2017's system price, calculated from Elexon Portal data, shown in Figure 16.

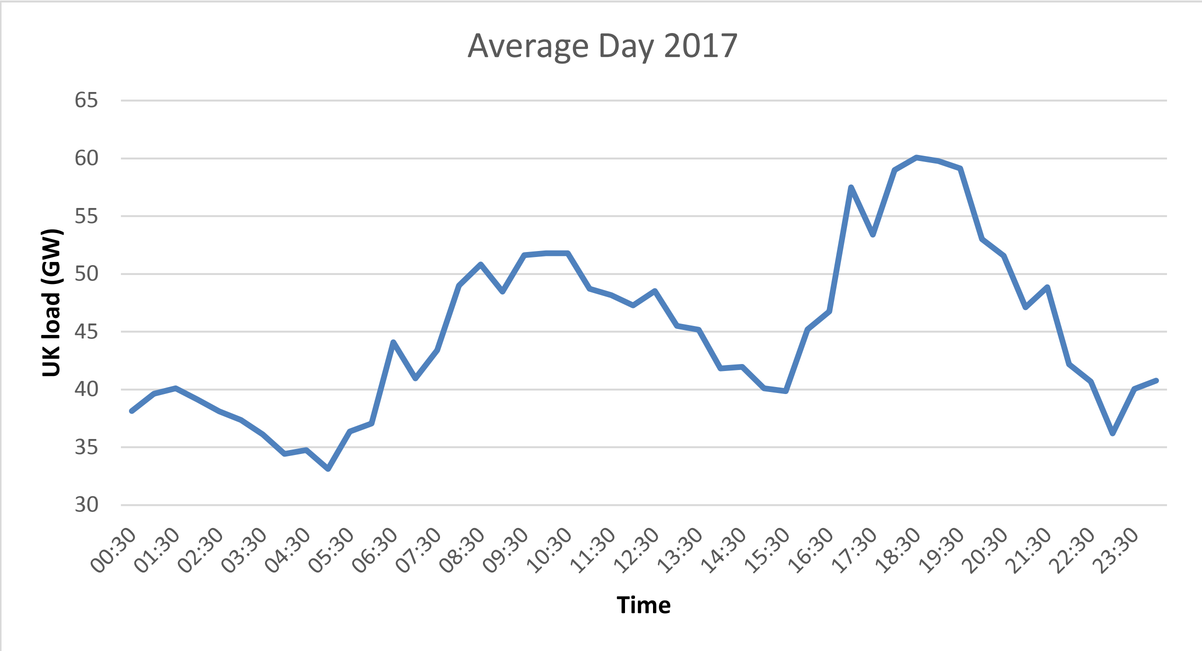

Figure 16: Average Wholesale Electricity Prices

In the summer months, the morning peak coincides with the start of solar production. In use cases 6 and 7, a solar park with a constrained grid connection was simulated, which requires the level of export to be kept below a specific value at all times. If the battery was still exporting for arbitrage at 9am, then not only was the battery unavailable for peak lopping; it was actively increasing the export level of the park. Indeed, on certain days it could increase the export beyond the agreed grid export limit of the park, not just the virtual constraint. For this reason, the time of export was set an hour earlier in the morning. This allowed the system to capture some of the expected system price increases while keeping in line with other commitments and restrictions.

Arbitrage was the longest running test of the project, having occurred at the beginning of each test day (other than the test days for ramp rate control) and thus has the widest seasonal coverage. To enhance the learning available from this test, the entire year of 2017 has been assessed as if the battery had been scheduled every day. This gives a broader picture, including seasonal variations of the potential arbitrage opportunities.

The key to successful and profitable commercial arbitrage is access to the most volatile energy prices; the system price. This is the opposite to the typical operating regimes of solar parks, which value certainty and stability of income over potential, but less certain, increases in revenue. Many solar parks enter in to long term PPAs or sleeved PPAs (allowing the sale of energy to a specific remote customer rather than an energy supplier) at a set rate (£/MWh). If market prices increase beyond this set rate, the park is generating less revenue than it potentially could, but the certainty is more valuable to an investor than the possibility of upside.

The Copley Wood Solar Park had a short-term export PPA, to allow flexibility for future opportunities. The short-term agreement was designed to reduce uncertainty and exposure to price swings, as demonstrated in Figure 17. The import prices are from the Copley Wood PPA, the system price is from the Elexon Portal, and the N2EX next day auction price is from Nord Pool[6].

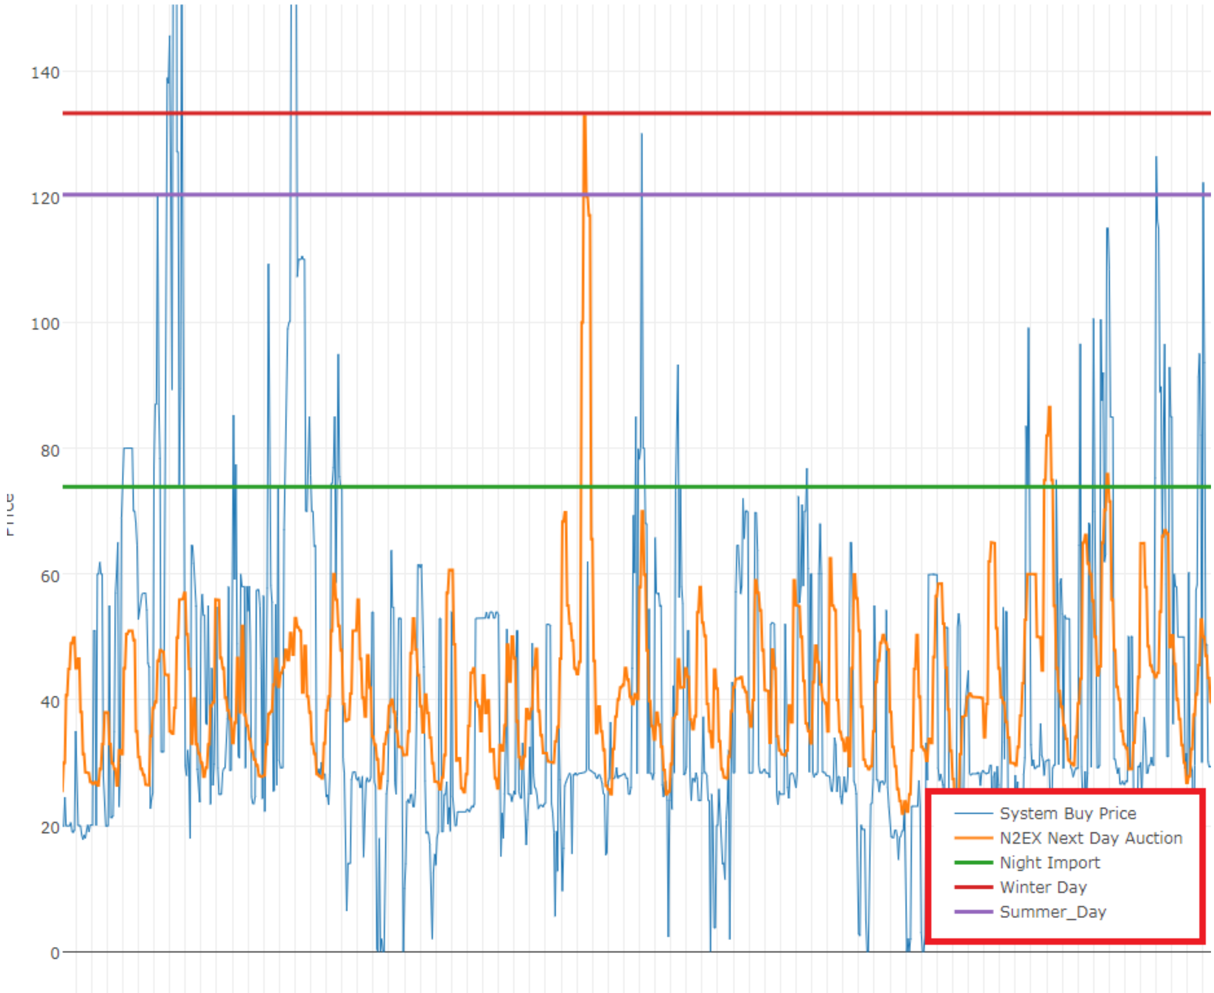

Figure 17: 3 months of Wholesale and N2EX Energy Price Variations

The next day auction smooths out the excesses of the system prices, reducing the solar parks exposure to sudden dips during the middle of the day. Indeed, while the system price goes negative several times, the N2EX price does not.

This shows the opposite aims of the battery and solar park. While the battery is connected to the solar park it cannot access the extremes of the system price, which makes arbitrage much less profitable. The import price is even less variable, with a flat summer and winter day rate and night rate (Figure 17: 3 months of Wholesale and N2EX Energy Price Variations). There are only 556 half hour periods in the whole of 2017 (out of 17,520) where

[6] https://www.nordpoolgroup.com/Market-data1/GB/Auction-prices/UK/Hourly/?view=table

the N2EX export price exceeds that of the overnight import price. The battery was capable of exporting half its rated capacity (i.e. 155kWh) every half hour. If the battery was 100% efficient and dedicated only to arbitrage, with no conflicts of use or grid restrictions due to solar export, this would give a theoretical arbitrage income for the year of £981. The decommissioning tests (as detailed in section 8) showed that the batteries efficiencies at charge levels over 150kW was approximately 88%, so by multiplying the import price by 1.12 gives a more realistic maximum of £482.

This potential income doesn't consider any loss from being unavailable to other revenue streams, or any cost savings from being attached to the solar park. It does demonstrate, however, that batteries retro-fitted to existing generation assets can suffer from an inability to access arbitrage markets, as they are essentially at the mercy of existing grid agreements. It is also possible that some arbitrage opportunities will be inaccessible due to the solar park taking priority over the shared grid agreement.

A standalone battery would negotiate different terms (potentially a co-located battery could also do this, provided the battery guaranteed to repay any shortfall in the income of the solar park), so to investigate the arbitrage opportunity of a new-build battery a theoretical simulated PPA has been created. This provides calculated figures of import and export prices, based on what is available on the market. Import costs include system price, supplier profit margin, AAHEDC (Assistance for Areas with High Electricity Distribution Costs), Red, Amber and Green DUoS charges[7] and 7% system losses. The export price includes a 7% loss credit (awarded to distributed generation for generating locally), a Red DUoS credit, BSUoS and RCRC credit, minus a 5% supplier profit margin. Climate Change Levy (CCL), Feed in Tariff (FiT), Contracts for Different (CfD), Capacity Market and Renewable Obligation Certificate (ROC) levies are not included in this PPA as, at the time of writing, whether energy storage should be paying these final consumption levies is in consultation, with the expectation that they will be exempt going forward. This should allow the analysis to still be relevant after the changes have been brought in.

The simulated PPA is available to view in Appendix A.

This simulated PPA follows the system price much more closely than the N2EX price, giving access to the market volatility that makes arbitrage profitable. It is based on the full year of 2017's Elexon data, which is when most of the battery testing was carried out.

It is difficult to calculate the arbitrage potential from this simulated PPA as there is not a flat overnight rate. Each export period was compared to the previous 48 half-hourly import periods, looking for profitable trades. The comparison assumes that the battery would charge and discharge at 310kW for the full half hour, for a total of 155kWh, and that the round-trip efficiency remains stable at 88%. It ensures that no two trades are using the

[7] The Red/Amber/Green are additional charges from the DNO on each kWh consumed. These are time-of-use charges and are set at three levels, with red being the highest and green being the lowest. Each DNO sets their own red amber and green times according to the peak load on their network. In this way, they act as a signal for companies to reduce their consumption during the red hours, which also reduces the requirement for network reinforcement.

same import or export periods, then produces a file listing all profitable trades. Over the course of the year, again assuming no loss from inability to enter other contracts, the total estimated income is £8,266, significantly less than the circa £20,000 that can be earned with 2018 FFR and capacity market incomes, or circa £64,000 earned from the 2016 rates (Table 4: Arbitrage Revenue and Cost Summary).

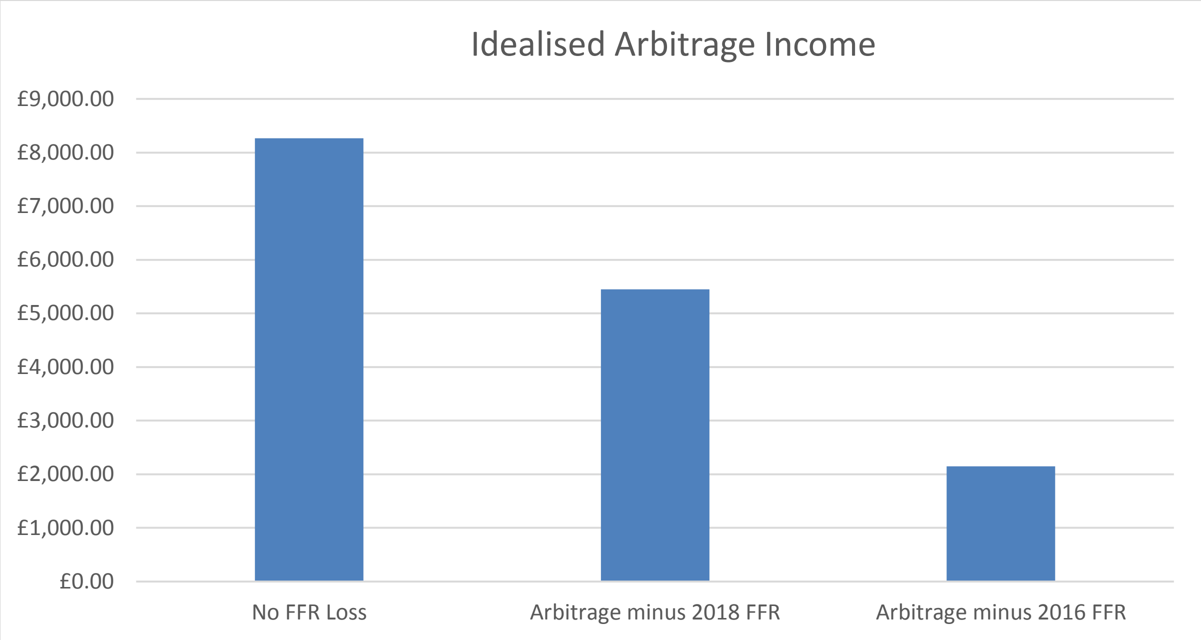

Figure 18: Idealised Arbitrage Income

The time taken for each trade to occur comes with a cost; in this case, the lost income from the FFR and Capacity Market. The comparison parameters were adapted so that the trade time was recorded and the FFR 2018 base case rate deducted from the income, providing the opportunity cost. This reduced the number of profitable trades, only leaving those that brought in more money than was lost from leaving the FFR market. This was repeated for the 2016 FFR rates (Figure 18: Idealised Arbitrage Income).

The analysis showed that, should the battery be commissioned in 2018, there is an additional potential £5,000 to be gained from the arbitrage market by leaving the FFR market to trade power. This could deliver more than a 20% increase in income, were it possible to capture all of the profitable trades. The FFR market currently operates monthly - each asset must decide for a whole month which market it wants to operate in, rather than changing each half hour depending on which is more profitable. Therefore, the profit for each calendar month needs to be greater than an entire month of FFR.

Table 4: Arbitrage Revenue and Cost Summary demonstrates that at no point does the arbitrage income exceed or equal the potential income from even the low FFR prices of 2018, when extended over a whole month. This demonstrates that, at the time of this project, arbitrage cannot meet the expected return rates for batteries on its own, although it does show a reliable fall-back income should an asset fail to gain an FFR contract. As the

project has progressed, the likelihood of this has increased significantly as there is now a large number of assets competing for FFR, with prices expected to be depressed even further. The data in Table 5 assumes that 24/7 dynamic primary and secondary FFR contracts are available, however, National Grid has begun only procuring overnight FFR, leaving arbitrage as a viable, location-agnostic revenue stream. There is no predicted requirement for daytime frequency response before April 2019. National Grid is on track with their procurement curve and so are not accepting bids for this time yet. With the FFR market already saturated, arbitrage seems to be the most obvious base-case revenue to convert to.

As the industry moves towards more flexible contracts, it is possible that the market will reach a point where it is possible to bid in and out of the FFR market in a much more flexible way, potentially to the point where bids are accepted an hour ahead, but this is considered unlikely for the next few years. In that scenario, a spike or drop in system price would result in a large percentage of the energy storage market bidding out of FFR to chase the arbitrage opportunity, causing instability on the grid. Therefore, the extra income shown in Figure 18 is unlikely to be realised.

Date	Arbitrage Income	Low_FFR (2018)	Arbitrage Loss	High_FFR (2016)	Arbitrage Loss
31/01/2017	£828	£2,076	-£1,247	£5,439	-£4,610
28/02/2017	£589	£1,875	-£1,285	£4,912	-£4,323
31/03/2017	£624	£2,076	-£1,452	£5,439	-£4,815
30/04/2017	£569	£2,009	-£1,440	£5,263	-£4,694
31/05/2017	£1,377	£2,076	-£699	£5,439	-£4,062
30/06/2017	£708	£2,009	-£1,301	£5,263	-£4,556
31/07/2017	£697	£2,076	-£1,379	£5,439	-£4,741
31/08/2017	£578	£2,076	-£1,497	£5,439	-£4,860
30/09/2017	£642	£2,009	-£1,367	£5,263	-£4,621
31/10/2017	£660	£2,076	-£1,415	£5,439	-£4,778
30/11/2017	£408	£2,009	-£1,601	£5,263	-£4,855
31/12/2017	£585	£2,076	-£1,491	£5,439	-£4,854
Total	£8,266	£24,440	-£16,174	£64,036	-£55,770

Table 4: Arbitrage Revenue and Cost Summary

So far, the analysis has focused on the theoretical 'perfect capture' of every arbitrage opportunity, but this is unrealistic. Battery maintenance and downtime would have a small impact, but the larger issue is that of predictability. It is not possible to predict with 100% accuracy the arbitrage opportunities and take advantage of them, so real-world numbers are likely to be lower.

The Copley Wood battery only had scheduled times for importing and exporting and continued to operate regardless of whether the prices were favourable or not, as the system was not linked to any price index. With the import and export agreements of the Copley Wood solar park in operation, the battery project lost money every time it operated (with only 556 half hour periods over the whole year where it could have made money). Referring back to the simulated PPA, however, and expanding the testing to every day over the whole year, gives an insight into how well schedule-based arbitrage works.

By 'blindly' importing and exporting between 4.00-6.00am and 6.00-8.00am for the full year of 2017, based on the theoretical PPA, the battery would have generated an income of slightly more than £400. For comparison, had the battery only done this on days when there would have been a net positive result, it would have generated nearly £1900. This demonstrates that nearly a quarter of the total arbitrage opportunity is available at this time, but this is reliant upon accurately deciding in advance which days will be profitable and which will not. Once the 2018 base case costs for 4 hours of missed FFR are added, the £400 income drops to more than £3,500 in cost, while the 2016 FFR case causes a loss of more than £10,000.

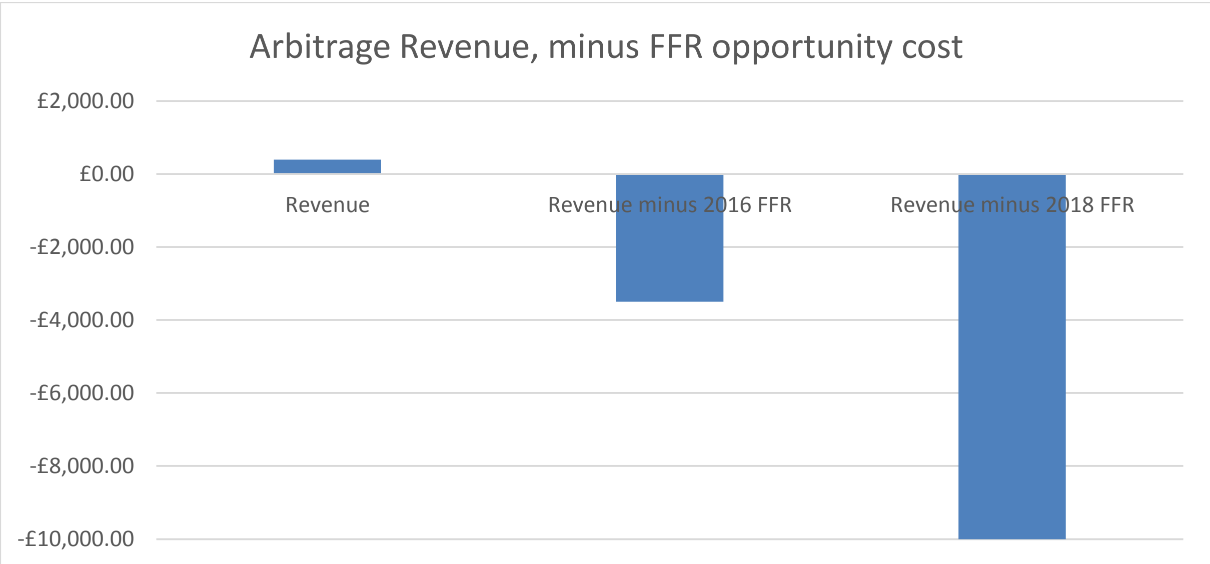

Figure 19: Arbitrage minus Opportunity Cost

Importing and exporting between 15.00-17.00pm and 17.00-19.00pm, had a smaller profit of £270 when operated every day during 2017, with 2018 FFR losses and 2016 FFR losses of £3,600 and £10,000 respectively. The most interesting development is that when looking at only the positive trades, this regime would have generated nearly £2800, backing up the prediction that volatile prices at this time of the day would be useful for arbitrage.

Nevertheless, the cost of bidding out of the FFR market for four hours each day over a year is estimated to be £4,073[8], so at 2018 prices this is not cost effective.

10.3 Conclusions

Arbitrage is the simplest of the use cases tested, requiring no complex software or data feed to operate. However, it is apparent that to capture the fluctuations in the marketplace, the battery needs to be under the control of energy trading companies who can predict with accuracy the direction of the short-term future prices. Even with this dedicated professional team, it is unlikely that all opportunities will be captured.

The high round trip efficiency of the battery helps smaller price gaps remain profitable, as only 12% of the power is wasted. This only holds true at high charge speeds, but if a half-hour trade is profitable then the battery will be operating at full rated power to gain the most volume possible at those prices.

At 2016 FFR and Capacity Market prices, when the battery was installed, arbitrage represented a very small opportunity relative to the base case income. £8,300/£64,000=13%. (from) As a fall-back revenue stream or a potential market to bid in and out of, it didn't offer a viable alternative. At 2018 prices, however, arbitrage income is more than 33% of FFR income. While still not a preferred choice for a storage asset, should FFR be unavailable it would ensure that the battery could still generate money, and this option is available to any battery installed directly on the grid.

Finally, arbitrage is significantly affected by being connected behind the meter of the solar park. The existing power agreements with the solar asset effectively block the batteries access to the variable price market. With new build solar parks this may be mitigated by drafting an agreement between the park and the battery, making the battery responsible for any loss the solar park experiences by selling and buying at system price versus being locked into a more stable market, but this is unlikely to work with retrofitting batteries to solar parks. These assets are now often owned by institutional investors whose attitude to risk is extremely cautious, making them unlikely to give up the long-term certainty provided by their current agreements.

[8] Based on 4 hours at £2.79 for 365 days a year (£2.79x4x365)

11. Results - Use Case 2 – Peak demand limiting

11.1 Introduction

The rise in the volume of distributed generation has reduced the spare capacity for additional generation within the DNO networks as developers have sought out locations where minimal reinforcement was required.

It might be expected that the reverse would be true of load, with the output from distributed generation offsetting load and increasing the capacity for further load connections without reinforcement. However, the intermittency of renewable generation means it cannot be relied on in network modelling, and solar power, for instance, stops generating before the peak evening demand is reached. The result of this is that transformers and lines must be upgraded for load capacity at the same point as they traditionally would be before distributed generation was installed. Indeed, the exported power from distributed generation has the same ability to cause thermal overloads on cable as load current and so in some cases distributed generation not only fails to offset issues caused by load current but can cause additional separate issues.

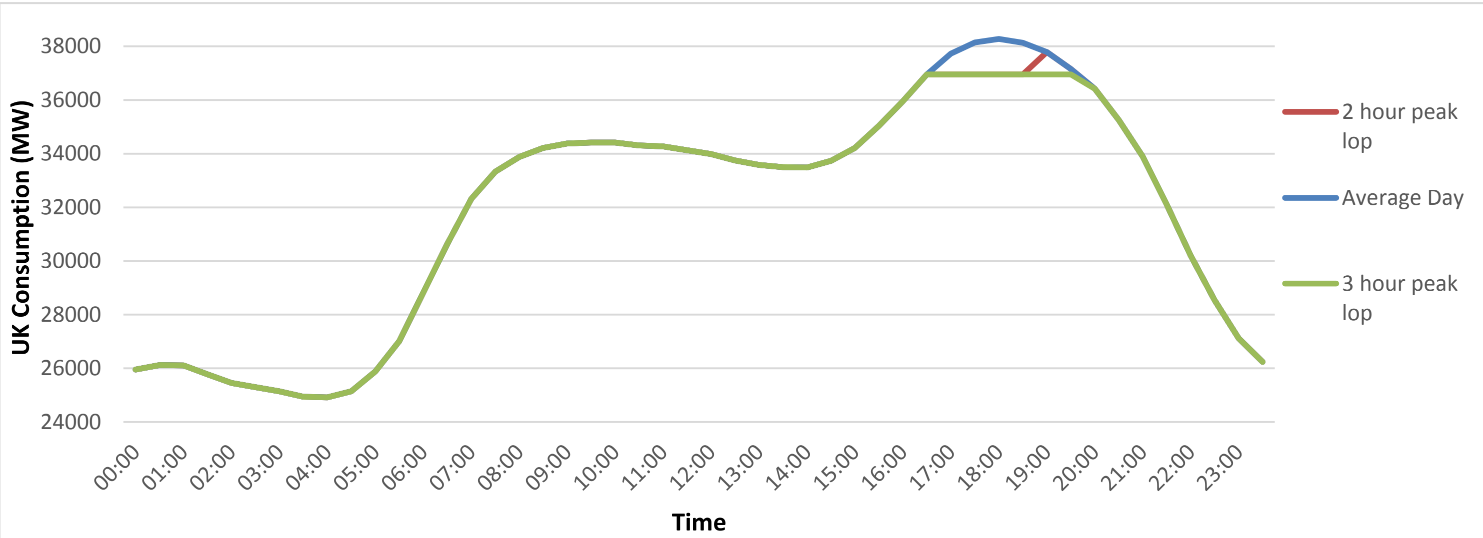

Figure 20: Average National UK Demand

There has been wide discussion for many years about the 'evening peak', the highest demand section of the day. Figure 20[9] shows the average UK demand over 2017, taken from

[9] https://gridwatch.templar.co.uk/

http://www.demand.ac.uk/wp-content/uploads/2015/11/DEMAND-insight-1-peak-electricity-final.pdf the BM (Balancing Mechanism) Reports data (National Grid data). The evening peak at the national level is between 5-7 pm.

This evening time period is therefore what drives infrastructure reinforcement at a national level. It also increases prices, as the less efficient and more expensive generators need to come online to meet the demand. This high load represents an opportunity for peak lopping: a two-hour or three-hour peak lop reduces the total amount of generation needed on the system.

On a more local level, the load profile often doesn't reflect the national trend. Factories with 24-hour operation typically have flat consumption profiles, while schools and business are likely to have a fairly predictable increase in load between 8am and 6pm. In fact, it is domestic customers that cause the evening increase.

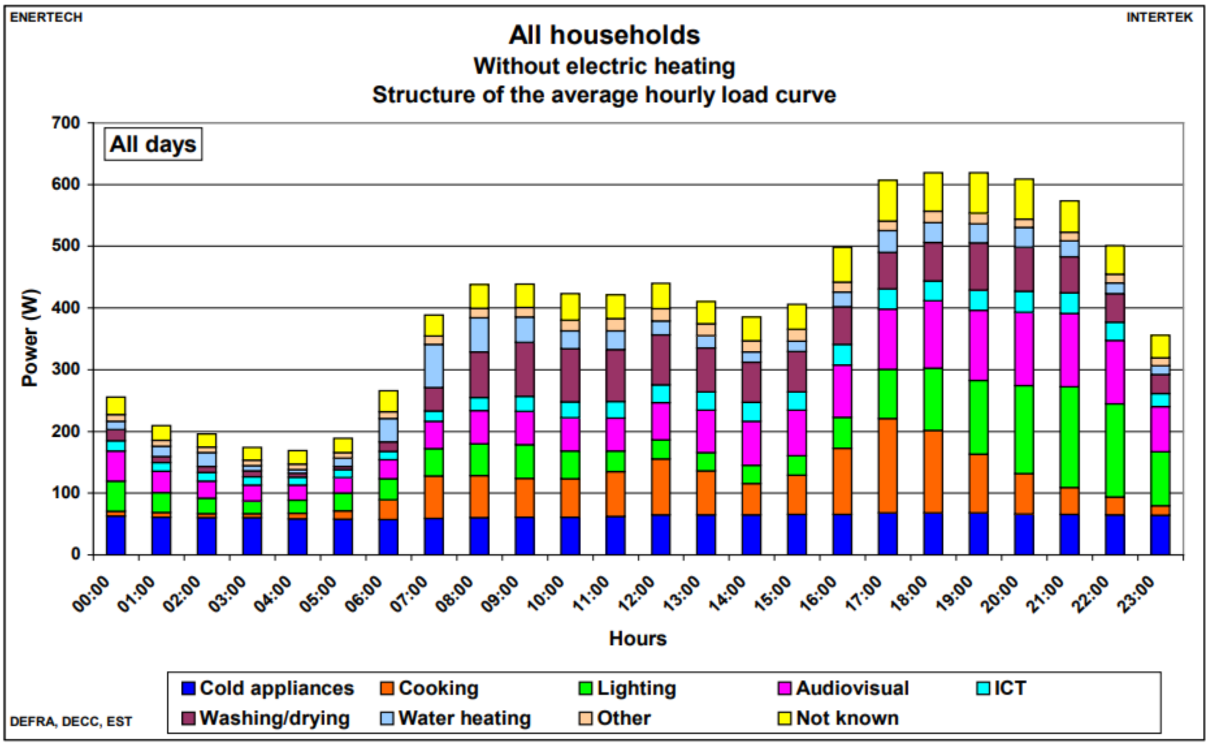

Figure 21: Domestic Electricity Consumption

It can be clearly seen from Figure 21[10] that the evening increase is much more pronounced for households. This demonstrates that DNO-level peak lopping will be significantly more viable and useful on substations that predominantly feed domestic customers rather than industrial or commercial properties.

The potential for avoiding reinforcement using batteries with timed/triggered export has a set of specific requirements. The substation needs to be predominantly feeding domestic customers and the load must be close to the limits of the substation assets or the higher voltage feeding circuits. The predictions for future evening load increases in the area should be static, or increasing very slowly year by year. If the load profile is not domestic, then it often isn't suitable for peak lopping and without the load being close to network capacity then there is no advantage to procuring peak lopping. Finally, if a large increase of load is predicted then the upgrades will be needed regardless, meaning there is little to save from operating peak lopping.

11.2 Method

Time-based peak lopping is one of the simplest use cases to implement. The battery was programmed to export between 5-7pm, during the red period of the DUoS charges, as it was

10 https://assets.publishing.service.gov.uk/government/uploads/system/uploads/attachment_data/file/208097/10043_R66141HouseholdElectricitySurveyFinalReportissue4.pdf

expected these will coincide with the highest use of the system. This decision was made early in the project, as the analogues dataset for the Millfield substation were only available at the end of the project. The Argand monitoring system only monitored the 11kV feeder to the solar parks and battery, and so could not see the load on the substation. WPD's analogue datasets are at a much lower resolution than the Argand system, only recording every hour, which is still sufficient to show a profile.

Use case 1 (arbitrage) complimented the network peak lopping, based on the assumption that the red DUoS signified the peak load times on the network. (Red DUoS times are set for when the load is highest across the DNO region, although further testing at Millfield Primary during the project showed that the peak load in this localised area occurred at a different time) This meant during the early, more simplistic test schedules, the battery exported all of its energy between 5-7pm. As the tests evolved, and more use cases had to be satisfied, the network support supplied by the battery had to be weaker. The local peak lopping use case required export to be available at the same time in case of an inter-trip signal , which reduces the power available for the network peak lopping. In addition, the later use cases were using solar peak lopping throughout the day which did not always charge the battery sufficiently.

To combat this, a change was requested in the software to alter the state of charge limits, allowing the state of charge manager to be used as a charge-based import override. This effectively meant that the peak lopping algorithm could be left to function, but if by a certain time the battery wasn't charged above a specified SoC percentage the battery would charge at a constant rate ignoring the level of solar output. It is considered that should the battery have a contract for network support, it would ensure the SoC was sufficient regardless of whether it had to take the energy from the solar park or the network (though solar park energy is significantly cheaper).

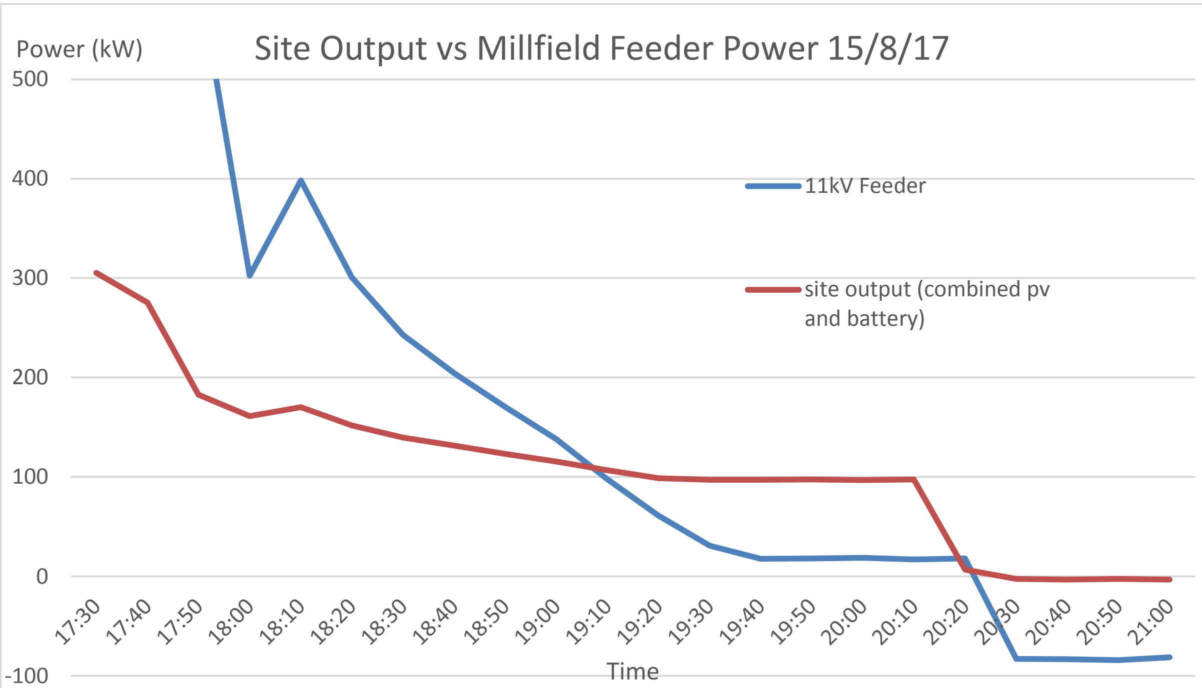

Figure 22: Site output vs Millfield feeder

Exporting into the network below a primary substation will reduce the load on that substation, so the aim of this use case was to prove the reliability of the service and to investigate the potential costs of a standalone battery and a solar-connected battery in providing this service.

Figure 22 shows the evening output of one of the test days. The Argand Millfield data was only available from 20/6/2017, which was after the lower evening export levels were being used, meaning the battery exported for an extended period at a lower level. On this specific day the battery was exporting at 100kW, with the tail end of the solar output also being visible. It can be seen that the export of the battery effectively removed the passive load from the two solar parks, Copley Wood and Higher Hill Farm, on the feeder. When the battery stopped exporting, the load increased to approximately 100kW. Obviously peak lopping at this level is not useful for a primary substation, but it does demonstrate the principle of masking load with battery export.

In the case of a grid connected battery, it is likely the goal would be to minimise the time required to complete the use case. This would involve importing between 3-5am and exporting between 5-7am. Based on the simulated PPA, the battery would have the same profit/loss as the evening arbitrage discussed previously. This gave a small profit of £270 in 2017, reflecting the fact that energy prices at this time varied considerably, sometimes creating a loss rather than a profit each day. The variation and unpredictability of energy prices means this profit will not be consistent year on year but does show that the battery is unlikely to incur heavy losses by importing and exporting at these times. This reduces uncertainty and therefore the price required by a commercial operator to offer this service to a DNO. However, the service as modelled will leave the battery operationally empty (10% SoC) at 7pm and, depending on the revenue stack of the battery, it may require charging afterwards to return to an operational SoC. If this is the case, the DNO/DSO payment will have to cover the increased import costs and time that the battery could have been earning money elsewhere.

The total opportunity cost of bidding out of the FFR revenue stack is £4,073 for this battery per annum, assuming no extra time is required for SoC management. Should an extra hour be required to return to 50% SoC, the cost rises to £5,092 per annum. These prices would need to be matched by the DNO, with an additional revenue for price uncertainties. This could be conditional on reported costs of power, avoiding the operator's exposure to changing prices. It should be noted that the FFR price decline means that year on year, the DNO price to match that revenue will also decline. The price paid could also be reduced further by offering a longer contract.

The price estimated for installing a WPD 33/11kV transformer is £300-500k, and upgrading a substation would require replacement of two to maintain firm capacity. Using the minimum cost of £600k, with an inflation of 2% and discount rate of 5%, it is possible to calculate the value to the DNO of deferring the upgrade for 5 years. Just over £65,500 would be saved by avoiding the upgrade, which then becomes the maximum budget available for a 5-year contract with energy storage/other controlled generation to peak lop the network.

It should be noted that in reality there would be a discount factor applied to the £65,500, as a DNO will need to cover the cost of administering flexibility services and will bear an additional risk compared with the robustness of traditional reinforcement. The new transformers would be directly under DNO control with proven reliability records, while the flexibility market is still developing. It has been assumed that if the flexibility service is 20% cheaper than the upgrade costs, this would be considered a significant enough saving to be worth contracting, leaving a budget of £52,500. This equates to a flat annual payment of over £12,000, well in excess of the £5,092 opportunity cost. Whether a battery of this size would be sufficient to delay a transformer upgrade for 5 years is however in doubt. Scaling up the opportunity cost to match the £12,000 would give a battery capacity of just over 700kW. It should be noted that this DNO contract for 5 years is significantly longer than any other contract currently available to energy storage, other than the capacity market, and would maintain its value while FFR prices are expected to decline sharply.

A co-located attached battery has access to different prices, as it can import power from the solar park at a cheaper rate than from the grid. However, this battery was limited to exporting on the solar parks PPA, which usually means lower export prices. On days where the solar generation is not sufficient between 3-5pm, the battery will have to import more expensive power from the network and may be unable to recoup its costs due to the lower export prices on the solar PPA. After running a simulation on the 2017 solar generation data, the profit was increased to £556 per annum, demonstrating the advantage of the cheaper solar power outweighs the less variable export rates compared to a grid connected battery

Millfield Primary has a mix of loads connected to it, with a sizable residential area, a Clarks Distribution Centre and the 240-acre Millfield School campus. It also has 3 solar parks connected to the 11kV side (Copley Wood, Higher Hill Park and Butleigh Park), totalling approximately 10MW of AC generation (11.5MW DC). This mix of loads create some unusual load patterns, which are not consistent with the average consumption profile of the UK.

Chart n : Chart e title

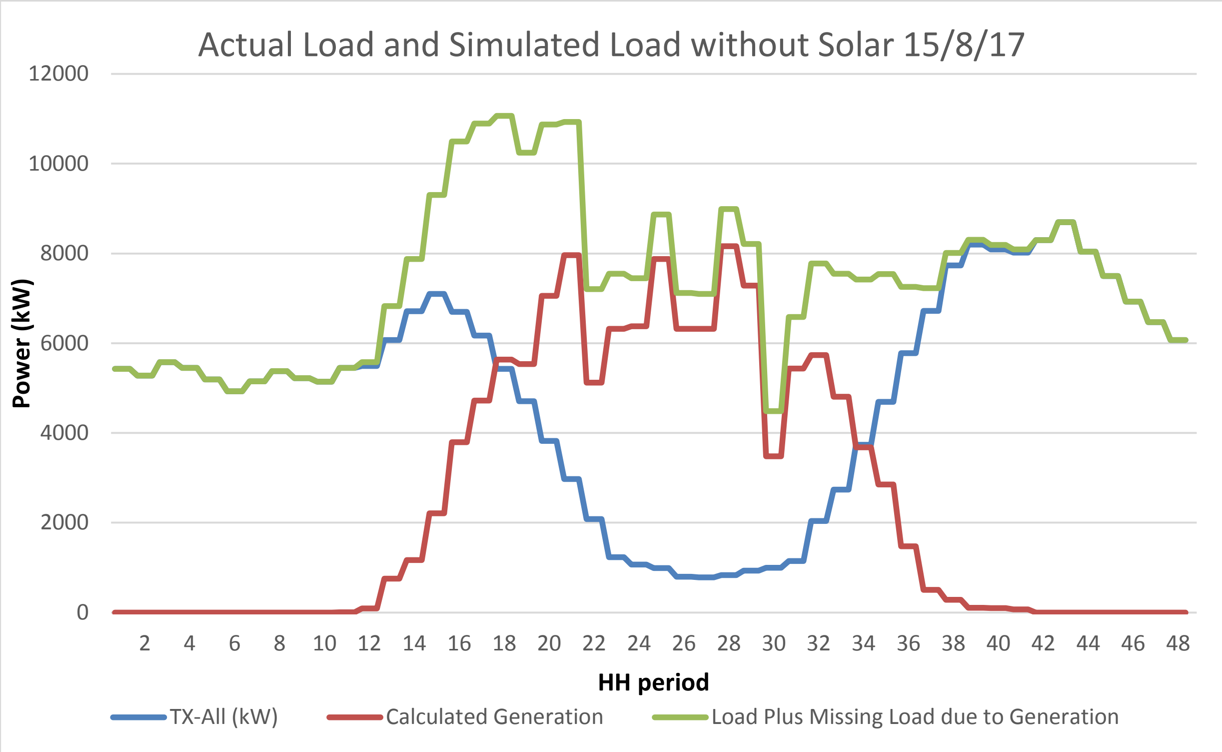

Figure 23: Actual and simulated Millfield load

Figure 23: Actual and simulated Millfield load above shows the full day of load on the 15/8/17, the evening of which is also in Figure 22. The blue line shows the load on the Primary from the analogue data provided by WPD. From the readings it appears that the Millfield substation has the standard evening peaks. However, the substation has three solar parks connected to it, two of which are on the monitored feeder (Copley Wood and Higher Hill Farm). The other park (Butleigh) is close by (2.5km) and is of a similar design to Higher Hill Farm (pitch angle, row spacing, DC:AC ratios), so it is likely that the Butleigh solar park will be experiencing almost identical conditions and output. In order to create a value of total generation feeding into the substation, the Copley Wood generation data has been subtracted from the 11kV feeder, and the remaining generation doubled, before adding the Copley Wood output back in again. This value does have limitations, for example the fluctuations shown above are likely to be staggered across two solar parks giving a more consistent output, but it does reflect a reasonably approximate value.

By adding the generation back in as 'invisible' load, the true consumption pattern of the Millfield substation is revealed. The substation on this day (15/8/2017) had a significant morning peak, declining around 11am and staying relatively consistent until 10pm. With the higher load occurring in the morning, the substation would not be suitable for peak lopping. However, this example day is from the summer, with extensive photovoltaic (PV) production and the lower load that comes with less heating requirements in the summer. The graph highlights the challenge faced by DNOs, i.e. they have no oversight of the output of distributed generation, only the drop in demand, and this can be hard to separate out from changes in the load profile of the local area.

A more useful metric to investigate the peak lopping viability is looking at the maximum load on the transformers throughout the dataset. The data runs from the 24th January 2016 to the 22nd January 2018, giving a two-year sample size. The highest readings for each half hour period has been plotted in Figure 24, illustrating that there is a peak that can be reduced by energy storage.

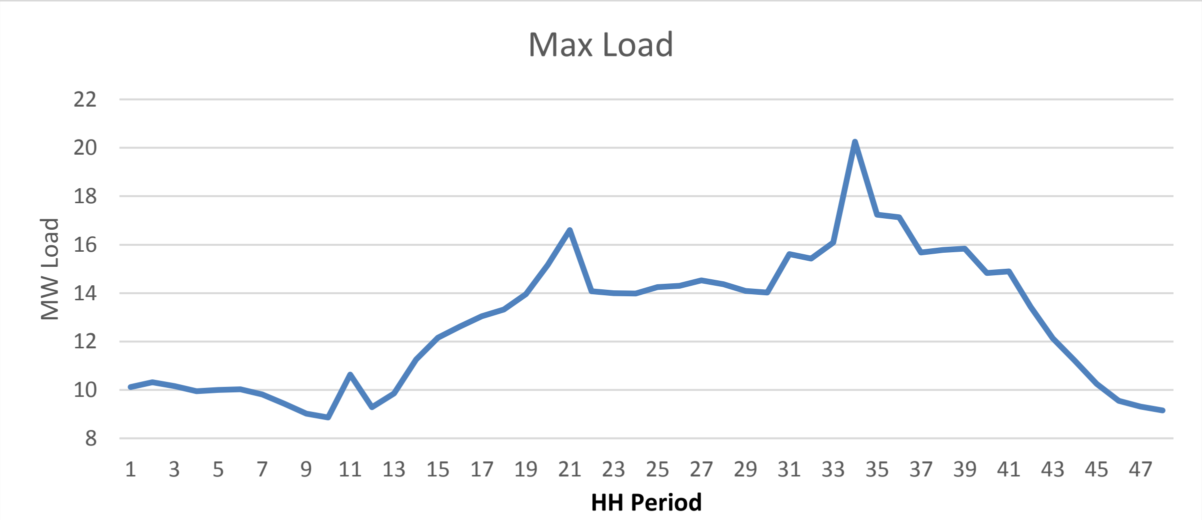

Figure 24: Maximum load over 2 years

It can be seen that the peak is only present for one half hour period, but it is significantly higher than the next largest data point. This kind of load pattern is ideal for peak lopping, as there is a short-term high load. If, for example, the maximum capacity of the substation was 18MW, served by two 9MW transformers, then a battery export of 2.25MW for half an hour would keep the substation under the thermal limits when presented with a load as detailed in Figure 24. Indeed, over the two years of data collection, a transformer with a limit of 16MW (2 x 8MW) would not have required additional support except for 9 half hour periods. All of these periods are in January or December, as expected with peak loads. Four of these half hourly periods were on the same day, meaning a maximum of 2 hours generation would be required. However, there is one half hour period where an overload would have occurred at 10:30am, outside the expected peak load window.

11.3 Conclusions

Even a rural substation with significant non-domestic loads could benefit from peak lopping services if it was nearing the limits of its capacity. The initial look at the effect caused by the existing solar generation showed a distorted load profile, with the solar only providing a reliable impact during the summer months. Realistically, all of the peak loads seen by the network are going to be in the middle of winter, as even if air conditioning was installed at a large scale in the UK as has been seen in other countries, it would be used during hot days when embedded solar is reducing the load seen by the network.

On the illustrated peak day, the minimum overnight load was at the same level as the peak load seen in the summer. This demonstrates that the requirement for peak lopping is seasonal, meaning a DNO would not need to pay for a year-round contract. Taking the theoretical limit of 16MW as an example, the contract would only be required in December and January, reducing the opportunity cost of providing this service by 10/12 compared to a full year contract against the FFR base case. It is clear however that a larger scale battery is needed - 4.25MW of load reduction would be required on the 11kV busbar to keep the load at 16MW for the peak half hour, far in excess of the power of the test battery.

The other cause for concern is the day that the overload occurred at 10:30am. It had been assumed that the higher loads would occur during the red DuOS period, so for a 2-hour battery 4 hours a day would be needed for the contract. The excursion outside this time suggests a requirement for a greater degree of flexibility, for example, there could be a contractual requirement that the battery should always be charged to 25% during these two months to provide ad-hoc network support as required. This would allow the battery to still operate in arbitrage while fulfilling the network requirements.

It would appear that the Copley Wood battery (310kW) was too small to have a significant effect on the network. This isn't surprising as it was not specified to fulfil any particular network requirements, but a short duration high power battery would be of more use at the Millfield substation. It is also apparent that the time-based export scheme used for the tests would be insufficient for commercial operation, but rather the load level on the transformers should be able to remotely trigger an export from the battery on demand. This would be a similar method to that used in the local customer support use case, section 12, which would have been triggered by a soft inter-trip ANM signal.

This use case also raises the question of whether any type of energy storage is the most appropriate technology to always provide this service. If there is a location with a predictable, regular overload then it can make sense for a battery to contract with the DNO, especially as a new-build battery is likely to have a restriction on import built into is agreement at the peak time, to avoid making the situation worse. However, for the high peak power requirement, and the longer and less predictable overload periods, the opportunity cost of restricting the battery (a flexible asset with access to many diverse revenue streams) may be more onerous than, for example, a rarely used peaking power plant, which only gains income with high system prices.

This kind of service is being looked at closely by a number of DNOs as they look to avoid expensive traditional reinforcement. At a much higher level, Scottish and Southern Energy Networks recently put out a tender for generation in 'constraint management zones'(CMZs), where during peak winter loads a circuit outage would leave the DNO unable to provide sufficient firm capacity. If load isn't predicted to rise considerably for several years, then avoiding the expensive reinforcement (in this case restringing several kilometres of 132kV lines and potentially upgrading some transformers) can make a worthwhile business case. The size requirements for the generation sites were low, with one site being less than 4MW, but with an availability window of over 5 hours. The business case for a battery of this depth, especially one that would be subject to ANM to avoid making the network congestion even worse, doesn't currently exist. The other standard revenue streams for energy storage currently favour shallower batteries.

Looking further into the future, if energy storage becomes viable as peak lopping for solar parks, these assets are unlikely to be fully utilised in the winter, meaning the asset owner could see this kind of contract as an upside while the DNO can take advantage of the potentially cheaper price offered by an otherwise idle battery.

12. Results – Use Case 3 – Local Demand profile matching

12.1 Introduction

When a customer requires an increase in load, this can be accommodated cheaply if all the network assets feeding their premises have spare capacity. The only potential cost would be payment for the study the DNO would carry out to ensure all statutory limits were maintained when drawing the extra power. However, there can be examples where, for instance, a factory wishes to open a new production line but a large network asset (such as a primary substation transformer) would require an upgrade to accommodate the increase.

For this example, a cost apportionment method that means the factory would not be responsible for the total cost of reinforcement could be applied. For example, if an increase of 2MW was requested and this triggered a transformer upgrade, and the next size up of transformer was an increase of 10MW in capacity, then the factory would be liable for 20% of the cost (2MW/10MW x 100%).

WPD estimates that installation of a 33/11kV transformer would cost £300-500k. A primary substation requires redundancy, i.e. a minimum of two transformers, doubling the cost. Taking the middle estimate of £400k x2, at 20% apportionment factor, gives a reinforcement cost of £160,000. This is a significant increase on any investment a business is making to expand its facilities. The factory could investigate an Alternative Connection option, which would send a soft inter-trip signal whenever the consumption of the factory exceeded the safety limits of the local primary substation. This would provide a signal for the factory to shut down parts of its machinery to reduce its consumption to within acceptable limits. While the trade-off is clearer when two customers are connected to the same feeder with the soft-inter-trip signal reflecting the combined load, there is the potential for trading between customers for the more complex Alternative Connection option of Active Network Management. In this case it may be harder to predict when a particular customer was likely to be constrained and ensure sufficient charge is held in the battery to support the service.

A potential, but more carbon-intense, alternative to this would be the installation of signal-triggered diesel generators, although this would require finding space on-site with an additional connection point, and factoring the cost of ongoing maintenance and fuel. If there is an energy storage system also downstream of the constraint, then a further option is negotiating a contract for the battery to provide network support when needed.

12.2 Method

The method for local peak lopping is effectively the same as for network peak lopping, i.e. to export to the network to reduce load. The difference is in the implementation, with the network support being operated over a longer period and being totally time-based. However, in this case the system must be compatible with an ANM connection, meaning the system has to respond to an external trigger.

Originally the project was to have a standard WPD inter-trip box installed, which would then trigger a volt-free contact that could be monitored by the battery control software. After the battery had been installed, it was discovered that the inter-trip box would have a long lead time and a high capital cost. Given that the signal required would be false anyway, and BSR required easy access to trigger the signal, an alternative was sought. A small Raspberry Pi based device, with remote login capabilities was installed so that the signal could be scheduled. The system also sent an email as confirmation when the signal was enabled and disabled.

On the Millfield substation, the highest peaks are between 4 -6pm according to WPD's internal analogues (data points), which doesn't match the standard peak DUoS times. This means that the additional battery export will only be needed during these peak times, similar to the network peak lopping. It is possible under abnormal network conditions that there is excessive load outside of these times, but to safeguard against any of these would require a full-time contract or part of the battery capacity being dedicated to this service, which would be significantly more expensive while the chances of being utilised are extremely rare.

This was the first time the RESolve combination method was used (Appendix C). This allowed the battery to operate in network and local support mode at the same time, demonstrating that two revenue streams could be accessed at once. As they are both export use cases, they complement each other rather than directly competing, and the effectiveness would be easy to view. This is why it was selected as the first combination method.

As this export would be condition-based, rather than for a pre-set time, it isn't possible to predict exactly how much energy is required. While network support with export at 100kW for two hours is simple, requiring 200kWh, the export on request scenario means potentially no energy is required at all. The network data for Millfield wasn't available during the design of this test, but an assumption was given that the highest peak (causing a network overload) would only occur for a maximum of an hour a day, within the period 5-7pm. Subsequent data showed the peaks occur between 4-6pm, but the test was based around the DUoS times. A nominal value of 100kW was selected as the support level required to avoid an overload. The level of support could actually be up to 310kW with this battery, but it wouldn't have demonstrated the ability to fulfil two use cases at once.

During the earlier part of the day the battery was programmed for solar peak lopping, but if there wasn't sufficient generation to charge it then the contract for local support would be breached, if called upon. This set a minimum SoC that the battery had to enter the contracted hours at, to ensure it had sufficient energy. After discussion with RES, a modification to the State of Charge Manager was implemented, allowing the minimum to be set at 60% rather than 20%. With 10% being the minimum SoC used during the tests, this gave 50% of the batteries capacity available for both network and local peak lopping.

A capacity test wasn't carried out in the middle of the test regime, so it was unknown how much capacity the battery still had. However, with the nameplate capacity of 650kWh and commissioning tests demonstrating an original capacity of over 700kWh, it was assumed that the battery was still capable of storing over 600kWh. With 50% available for export, this equalled a minimum of 300kWh, with two hours of 100kW network support and 1 hour of optional triggered 100kW support. The final capacity tests proved the battery had well in excess of 600kWh remaining capacity, meaning this assumption was correct.

There was a limitation within RESolve that meant the same use case type couldn't be used twice within the combination method. The ideal scenario would be using the export control mode, with a two-hour timer at 100kW, and a second export control set at 200kW if the external signal is switched to on. The workaround that was selected was that the Auto SoC mode provided the base 100kW export while the Auto Export mode responded to the external signal.

This compromise meant that during the tests the 100kW export continued past the two hours, if the battery was charged above the 60% minimum. This is considered to be an artefact of the testing rather than a real ongoing limitation, as it merely requires a more tightly specified control system to avoid these issues.

An example day export is shown in Figure 24, with the full hour of additional local support being used in one continuous period. Conversations between the factory, DNO and battery asset owner would be required to establish minimum required support times after the signal has been triggered.

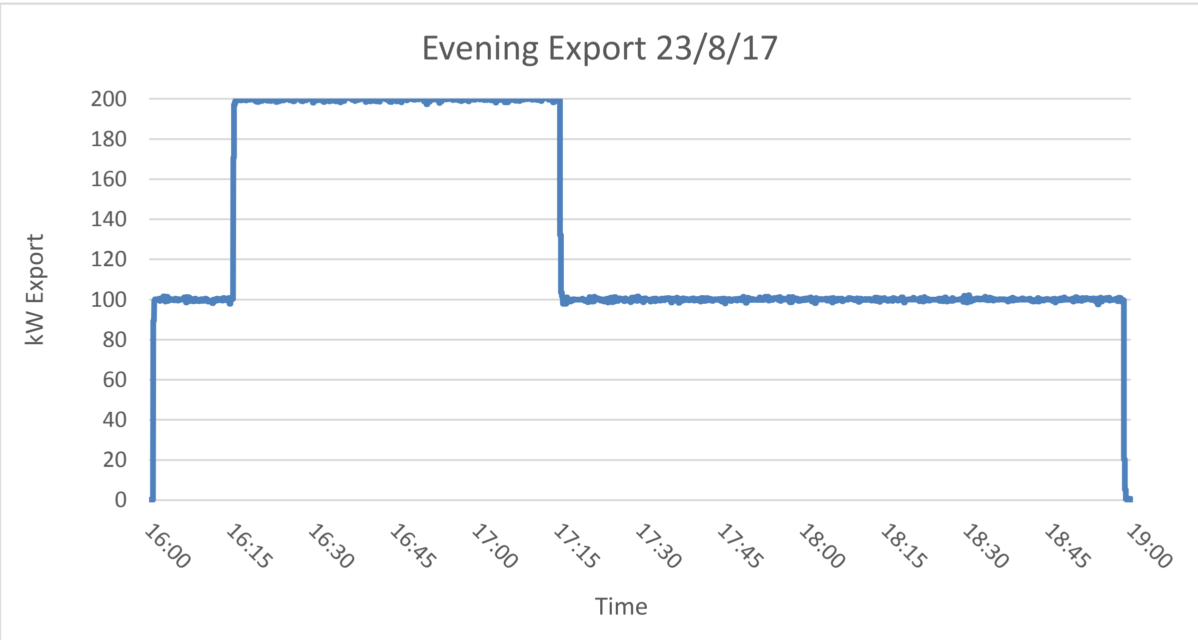

Figure 25: Combined Export

In this example, the signal was triggered by the Raspberry Pi at 16:15 standard time and released at 17:15. The combination method worked well in event drive mode, switching the control mode over to the 200kW export program.

Towards the end of the testing schedule, when the required number of tests had been completed, the opportunity to gain some additional learning was taken. The triggers were set to occur three times within the 2-hour availability window, though in total remaining under the hour allocation. This test demonstrated the reliability of response to the signal, and to simulate an alternative option of operating the network. In this scenario, the battery only exports for the duration the overload is detected, rather than for a set time after the initial excess load. Figure 26 shows the three increased export times, including a short two-minute window.

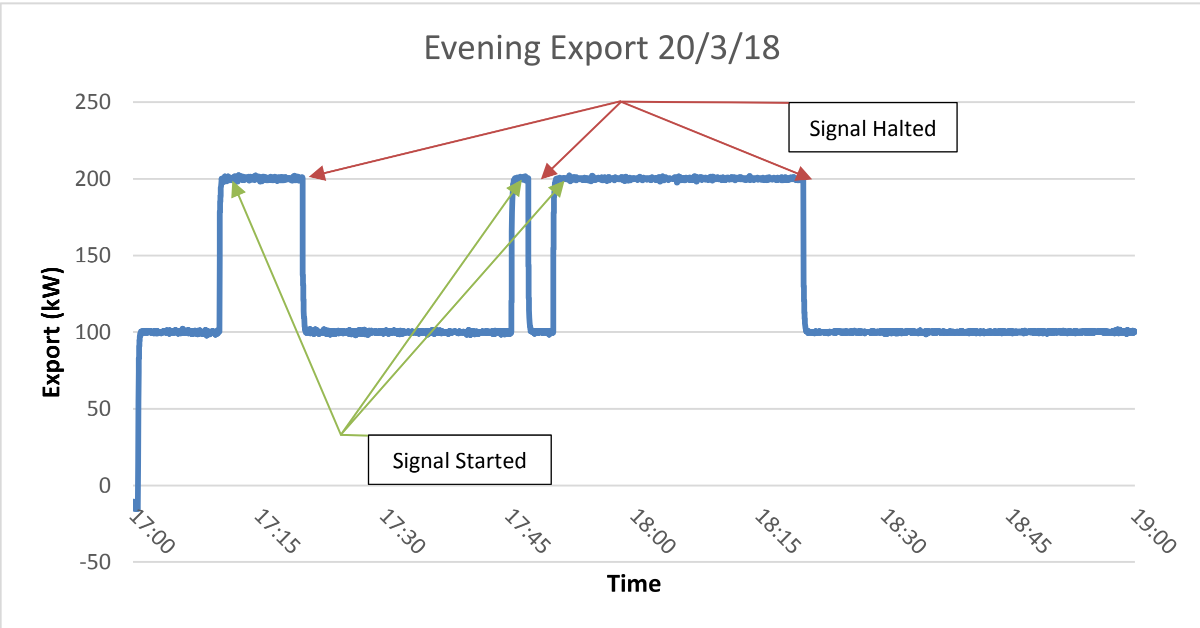

Figure 26: Multiple Export Triggers

12.3 Conclusions

The battery responded reliably to the trigger signal, increasing its output in line with the programmed schedules. There were no significant processing delays associated with reading the signal state and switching the mode, meaning the processing behind the detection is already system-ready. The combination method demonstrated that this kind of support contract was possible while allowing the battery to continue to operate in its standard mode. If the network market becomes extremely fluid, it could be possible to bid into other markets at a lower capacity for the duration of the support window, stacking revenues.

The opportunity cost is significantly lower due to this combination method, as there is still 210kW of battery power and total capacity less 100kWh available for other uses. The base case of FFR would only require bidding lower capacities into one of the six 4-hour 'windows' that National Grid offer. These windows run from 11pm in 4-hour periods, so total charging and discharge time for this case would easily fit into the 3-7pm window. Thus, the calculation gives a total of £3.60 per day[11]. Returning to the analysis of the Millfield substation data in section 11, it has been established that the highest loads occur between 4-6pm, rather than the expected 5-7pm. In this case, the level of support required (less than 50% of the 2-hour battery capacity) would mean the system still has time to charge within the hour between the start of the reduced FFR period and the beginning of the contracted support window.

It is assumed that the contract would need to run for four months in the winter, from November to the end of February. This totals 120 days, for a total opportunity cost of £436. The extremely low opportunity cost, coupled with the lower N2EX prices for the likely import time of 3pm, make this one of the smallest cost use cases. Not only this, but it is lower risk, as the battery is still able to continue the majority of its standard operations during the contracted time. The maximum value of the contract would be the cost of paying

5.25 The 'Security CAF' is applied, where the costs are driven by either thermal capacity or voltage (or both) as assessed against the relevant standard. This rule determines the proportion of the Reinforcement costs that should be paid by you as detailed below.

$$\text{Security CAF} = \frac{\text{Required Capacity}}{\text{New Network Capacity}} \times 100\% \qquad \text{(max 100\%)}$$

5.26 The 'Fault Level CAF' is applied, where the costs are driven by Fault Level restrictions. This rule determines the proportion of the Reinforcement costs that should be paid by you as detailed below.

$$\text{Fault Level CAF} = 3 \times \frac{\text{Fault Level Contribution from Connection}}{\text{New Fault Level Capacity}} \times 100\% \text{(max 100\%)}$$

Figure 27: WPD CAF calculation extract

[11] Opportunity cost at 2018 prices is £2.79 per hour. Effective opportunity cost = (full export capacity(310kW)/utilised export capacity(100kW)) x hours of operation (4) x opportunity cost (£2.79) = £3.60

for the network upgrade, minus any co-efficient[12] of perceived risk by the company relying on a battery asset vs having a full grid connection.

The above calculations are taken from WPDs South West statement of methodology and charges for connection[13], and demonstrate how customer contribution is calculated when a request for capacity requires network upgrades. The rules state that the "required capacity" is the total increase in capacity, even if only a small percentage of that is over the limit of the existing equipment. In addition, the calculation takes into account any relevant capacity increase from that customer within the last three years. This avoids customers requesting up to the available capacity limit, and then a month later requesting the remainder, "gaming" the calculation.

With a battery contract, it would be possible to avoid some of the reinforcement charges. At smaller capacity increases this isn't as viable, or if the required capacity is well above the current network capacity, as the payment to the battery would have to be larger. An example where the variables are in favour of this type of contract are: an existing capacity of 15MW provided by 2x15MW transformers (2 required for redundancy). The existing peak load is 12.7MW, and a company on the 11kV side wishes to increase its demand by 2.5MW. Traditionally, this would immediately trigger an upgrade, requiring two new transformers of the next available size, e.g. 20MVA. The upgrade would require both transformers, as the P2/6 regulations[14] require redundancies and the restoration of load customers within specific timeframes. If the installation of each transformer costs £500,000, then the calculation will be applied to a CAPEX of £1million. 2.5MW/20MW x 100 = 12.5%, or £125,000. However, if only 2.3MW was applied for, or a 2.5MW ANM connection, no reinforcement would be required. The battery is then contracted to provide the additional 200kW when the DNO transformers are overloaded. The company could then in three years' time apply for the final 200kW, triggering the reinforcement. This time the payment due would be 0.2MW/20MW x 100 = £10,000. Crudely speaking, if the three years of payments to the battery equalled less than the difference between the two grid payments, i.e. £115,000, then the battery contract is more economically advantageous. Even doubling the required power to 200kW only gives an opportunity cost of £872 for the battery, while there is a budget of over £35,000[15] to break even with the grid costs. This kind of cost avoidance is likely to be prevented by DNOs as soon as it begins to be exploited, as the upgrade costs would be spread among all customers rather than targeting the cause of the requirement.

If the factory was able to defer this CAF cost by 3 years, at a discount rate of 5% with 2% annual inflation, it would give a saving of £6,500. This would provide a flat rate of approx. £2,400 per annum with which to contract a battery to provide support, far in excess of the

---

[12] A company would calculate their own coefficient of perceived risk that would dictate how much they're prepared to pay vs the actual cost of reinforcement.

[13] https://www.westernpower.co.uk/docs/connections/Charging-Statements/Connections-South-West.aspx

[14] http://www.dcode.org.uk/assets/uploads/ENA_ER_P28_Issue_1__1989_.pdf

[15] (£115,000 / 3 years)

opportunity cost of £436. However, larger savings can be gained by utilising longer contracts. At 5 years, the saving is £12,600, and at 10 it is £26,400. The annual cost of these would be £2,900 and £3,400 respectively. The opportunity cost is £436 per 100kW, meaning these costs would cover 650kW and 750kW of peak demand. Even in this scenario, where all of the savings are paid to the battery each year, the factory is able to maintain the grid connection for a small annual fee rather than a large up-front payment, which could be easier for their cashflow.

This kind of more permanent arrangement would occur if the substation was space-limited, requiring extensive building replacement, extension and potentially additional land purchase. This can cause costs to rise out of control, meaning a company would stay contracted with the battery. When the cost of upgrade is so high, the alternative is usually to limit consumption rather than pay the grid costs, so the equation becomes the comparison between the payment to the battery and the cost of restricting the expansion of the company. Over time this requirement and payment could grow significantly, becoming the main use case of the battery during the winter months.

Financial implications are not the only factor when a third party is deciding whether to contract with energy storage in this way. Even if the quoted reinforcement costs are reasonable, there are many locations in the UK where grid connection offers now have a target date several years in the future. Despite not necessarily being charged for this work, there are often lists of 'Enabling Works' at either the DNO or even National Grid level, preventing additional capacity being released. Some quotes are being returned with a 3 year delay, which can leave businesses at a huge disadvantage against their competition in fast moving industries. In these cases, the price available for a battery contract is not calculated against the reinforcement cost, but against the cost of being unable to increase load for the next several years and potentially losing a competitive advantage. These contracts are likely to be the most lucrative.

This use case is very geographically dependent, requiring a series of specific network conditions coupled with a large load customer that wants to increase their consumption. The exact level of support required, duration each day, and amount of the year the contract is required, would be subject to specific negotiations with the company in question and are unlikely to follow similar patterns in separate locations. The key part of this testing was to prove that the battery could be "called upon" rather than always needing specific pre-programming. It can react to dynamic events.

This was the only use case within the testing that demonstrated the battery responding to an external signal. This is significant as it shows that these assets can leverage existing data collection systems and more complex analytics done remotely (e.g. WPD's SCADA network) and respond to a signal sent by these systems, rather than relying on the in-built computer being capable of all processing. This is how commercial batteries currently receive their setpoints for frequency response, and a hybrid system of local logic and remote setpoints is most likely to be effective. For example, solar peak lopping would require readings from a local sensor, and logically would be processed within the battery for speed of response, while an arbitrage system linked to a price-prediction algorithm could run remotely. This also allows a single central algorithm to control multiple assets, while the localised contracts such as network support would be handled by override logic within the local controller.

13. Results – Use Case 4 – Low demand grid voltage support

13.1 Introduction

Several areas of the network are set up to cope with heavy daytime loads, with transformer taps set up to avoid excessively low voltage at the end of feeders. However, overnight the demand drops, causing a voltage rise at the substation. Previous solutions have included automatic tap changers, which cut in and out extra turns on a transformer to adjust the voltage conversion ratio, and voltage reactors which adjust the power factor to avoid voltage rise.

Other issues due to light loading are caused by the Ferranti Effect. This effect manifests when low current and long cable length on some feeders act like capacitors, resulting in a voltage at the far end of the distribution line which is significantly higher than that at the substation end. This issue is often solved by the installation of voltage reactors (effectively inductors, cancelling out the capacitive effect of the empty lines). This scenario could occur when a factory feeder, that operates heavily loaded during the day, has no other loads on it during the overnight shutdown.

One of the simplest ways of overcoming the variability of the day-night load cycle is to ensure that the system is still loaded overnight. It was predicted at the beginning of the project that the battery was likely to be empty overnight after taking advantage of the higher evening peak prices. Therefore, it could charge up, providing a load on the network while charging during the expected low overnight prices.

13.2 Method

This use case had the simplest programming requirements for testing, even compared to use case 1 (arbitrage). The battery was programmed to import overnight, during early testing for a shorter period at a faster rate and then later for a longer period at a lower rate. The voltage was recorded by the battery PQM and the Argand monitoring at the Millfield Primary, so any variations due to the battery should be captured. This energy was then exported during the expected morning peak to leave the battery empty for testing the remaining use cases throughout the day.

The overnight load on the feeder is low enough that the additional import from the battery was a significant percentage increase. It was expected that the drop in voltage caused by the active power would be visible at the substation, although the battery was relatively small.

13.3 Conclusions

The testing had two import options, at a higher (250kW) and lower (100kW) rate. This was expected to show the impact on the voltage of different charge rates, and the lower rate would allow the battery to operate over a longer overnight period. The results of the test were plotted, with the import setpoint and voltage reading overlaid. The voltage during the night appeared to spike up and down randomly, with no obvious influence being demonstrated by the battery during the import events. These voltage changes are likely due to larger load fluctuations connected elsewhere, coupled with auto tap changer operation, and are likely to be seen at most primaries across the network. Several overnight operations, for 9 days in a row (between 10pm and 6am) can be seen in Figure 28. The voltage reading from the Millfield substation was increased by 150V to align it with the voltage recorded at the Point of Connection (PoC), to improve readability. The graph demonstrates the changes in voltages, not the voltage level itself, so this correction isn't considered important.

The voltage doesn't drop noticeably during the longer import windows, although there is closer correlation with the shorter export and voltage peaks. The actual effectiveness of this active power import and export is concealed by the randomness of the voltage overnight, which constantly alters based on external factors outside of the visibility of the project.

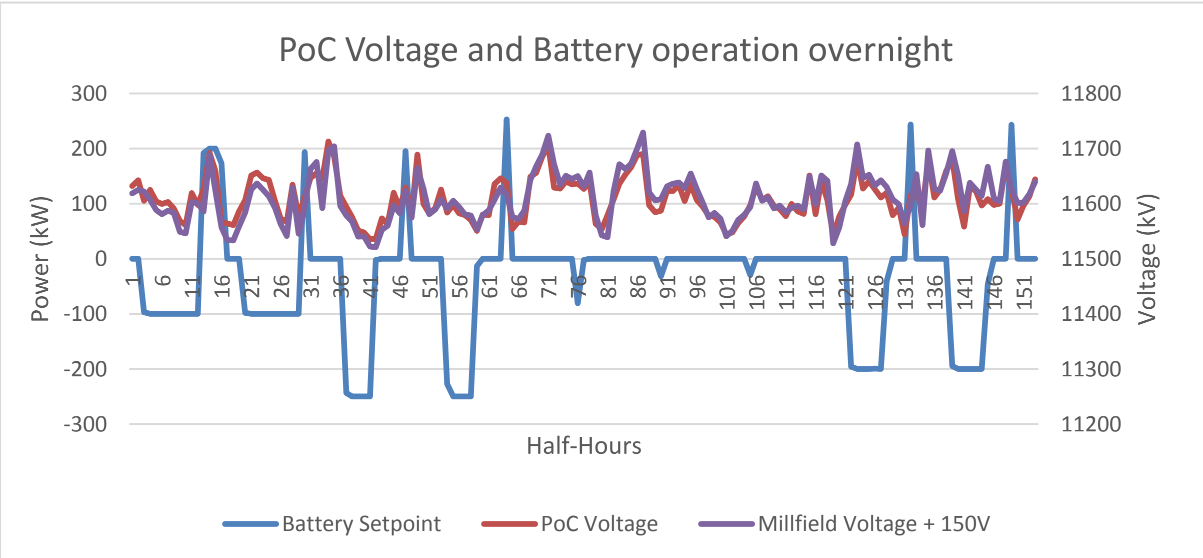

Figure 28: PoC voltage and battery setpoint

The reason that the battery had such a small effect on the voltage is simply a matter of relative size. Figure 29 shows the minimum, average and maximum loads of Millfield Primary over an almost two-year period (with missing data removed). The minimum load in this time overnight was approximately 5MW. (310/5000 x 100) 6.2% was the maximum potential influence the battery could add to this. A more average day gives the overnight loads at approximately 7MW, reducing the potential influence to 4.4%. A larger battery or a substation with less overnight demand is needed to see the effects of overnight import on the voltage.

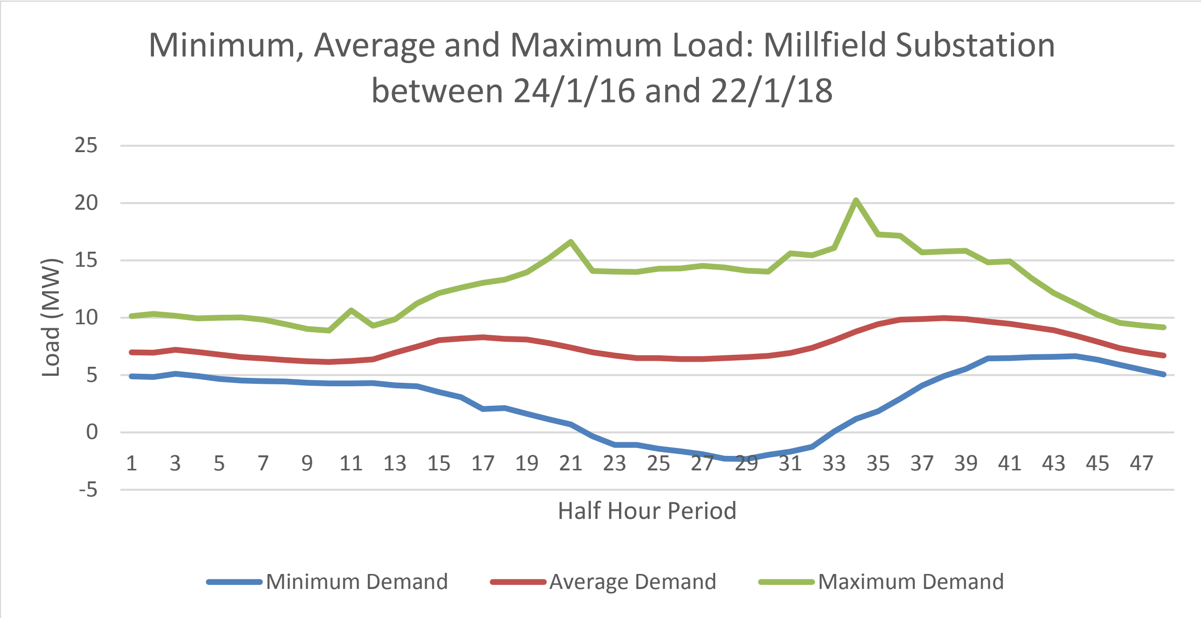

Figure 29: Minimum, Average and Maximum Loads

By definition the existing network is designed to resist voltage changes and stay within statutory limits, with tap changers set to automatically adjust based on voltage changes. It appears that the battery import doesn't have an effect, but if it was extremely effective this would trigger a tap change which could mask this. This kind of change isn't visible at half-hourly resolution. Another factor which impedes the research is the time limitation on the import versus the amount of power that can be imported. The battery could only sustain a 310kW import for two hours, whereas voltage control is likely to be required for a more significant proportion of the night. A more targeted research project investigating just this role could be carried out with a load bank, which while wasting power, doesn't have a time limit on its use.

Figure 28: PoC voltage and battery setpoint demonstrates a very close correlation between the changes in voltage observed at the PoC and that observed at the Millfield Primary. Having observed this correlation, it is possible to investigate the second by second data to observe the effect on the PoC, with the expectation that this should be mirrored at Millfield Primary The battery power, 310kW, is very small to be creating a large influence on an 11kV network. The SoC manager issues (detailed in section 7.1.1) created some unusual use patterns when the battery reached the set charge limits i.e. where the SoC manager would take control and export for a short period, before returning to the programmed schedule which immediately began importing again. While this was an issue to be worked around in other tests, it did create some data which demonstrates that the battery can affect the voltage.

Figure 306 shows second by second data at three points during the day, while the battery was running other tests. The battery oscillated between full import and full export quickly, giving an effective swing of 0.62MW

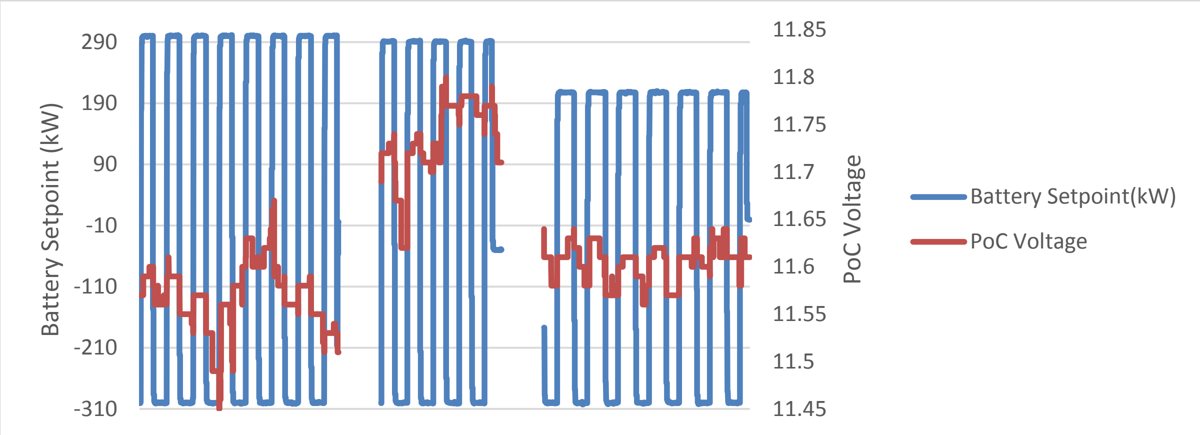

Figure 30: Voltage Fluctuation from change in power flow

The data was recorded in seconds, meaning these changes are occurring every few minutes. It causes approximately a 50-Volt variation each time it changes, proving that the active power flows do cause a change at the PoC (point of connection with the network). Due to the lack of second by second data at the substation it isn't possible to state with certainty that the voltage fluctuations continue right back to the substation, but given the correlation between voltage changes in the half hour data it is likely that they do.

The cost for this particular service is high, when calculated using the £2.79 opportunity cost. There are no mitigating factors that can reduce this price: if the active power is potentially required for 12 hours then the resulting opportunity cost is £33.48 each night the service is required. If the main focus for the battery is FFR then this could be even longer, depending on whether the start and end times align with the four-hour blocks in the auction process.

It is considered unlikely that this was an effective use of the battery. The voltage rise issue is likely to be prevalent in the summer, as winter loads are higher at all times reducing the voltage. In the summer months, there isn't enough time in the morning to take advantage of the morning peak prices before the solar park has started generation, blocking the batteries export route and preventing it from peak lopping during the day. The reactive power voltage control mode has been calculated to be much more attractive for a battery developer, although exact effectiveness for the DNO has to be quantified on an asset-by-asset basis.

The risk of high voltages may be reduced by other options which can have a more durable effect such as lowering voltages across the network in general, as has been trialled in South Wales following on from learning from the LV Network Templates project[16]. Another

[16] https://www.westernpower.co.uk/projects/network-templates

potential alternative is the System Voltage Optimisation method under investigation as part of WPD's Network Equilibrium project[17].

14. Results – Use Case 5 – Voltage control by reactive power

14.1 Introduction

The Reactive power is a necessary component of the energy transmitted through the transmission and distribution system. Active power is the value that most people will have encountered, with measurements in kW. For example, a 3kW kettle would require 3kW of active (or real) power, and if left on for an hour would consume 3kWh of energy. The heating element inside the kettle is effectively a huge resistor, called a linear load, and so only requires active power to operate.

An inductive motor, such as those found in many washing machines, not only requires the active power to do the work of making it spin but relies on the rotor and stator of the motor to interact by using a magnetic field. The power that creates the magnetic field is reactive power (kVAr). It isn't used in the same way that active power is but is required to constantly flow in and out of the stator coil to create the magnetic field. The same is true of a transformer where the reactive power creates the magnetic field that allows the primary and secondary coils to interact.

The ratio of reactive power versus active power is given by the power factor, which has values between 0 and 1. At 1, or unity power factor, there is no reactive power component at all. At 0, the only power flowing through the system is reactive power, e.g. magnetic fields in inductors are being maintained, but there is no real power flowing and so no real work is being done. If power reaches a transformer at unity power factor, the requirement for reactive power means the power flowing from the other side would be at a lower power factor, by passing through an inductor the synchronisation between current and voltage has been shifted.

Although reactive power isn't used in the same way as active power, inductors are noted as 'consumers' of reactive power. They cause the alternating current to lag behind the voltage rather than occurring at the same time. To compensate for the 'consumers', 'producers' are required, such as capacitors. Capacitors cause the current to lead the voltage, hence creating a leading power factor. The types of power factor can be seen in Figure 31.

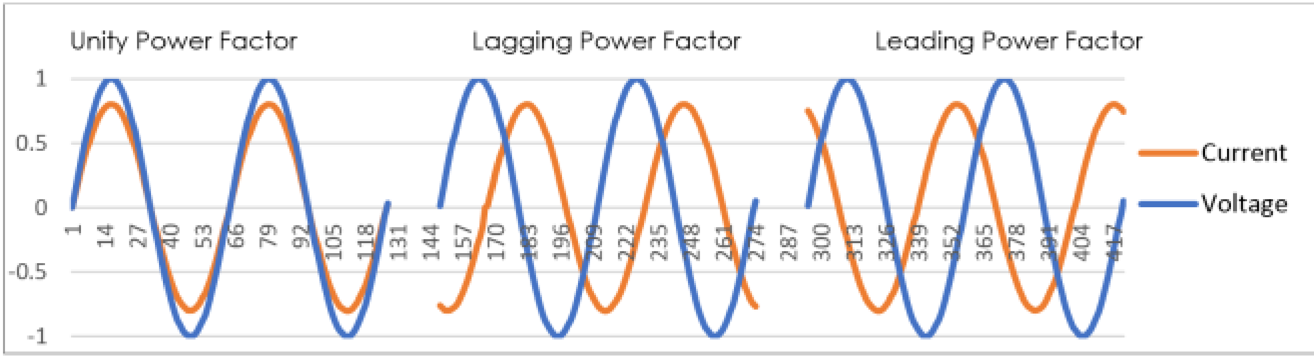

Figure 31: Leading and Lagging Power Factor

The power factor also influences how power networks operate, as with a power factor below 1.0 there is a requirement to increase the apparent power delivered in order to deliver the same levels of real power. Low power factor effectively wastes energy through higher losses and requires network assets to have a higher capacity than for unity power factor. Injecting leading or lagging reactive power onto a network alters the power factor and hence alters the voltage drop across the various elements of the network. Therefore reactive power can be used for voltage control on the network, although there are limitations due to the efficiency problems encountered at low power factor.

A battery inverter is able to both generate and absorb reactive power and operate over a wide power factor range. The battery inverter is effectively equivalent to a STATCOM, a device used for power factor and voltage regulation on the network. By generating reactive power when the voltage is too low, and absorbing when it is too high, the voltage can be maintained at a specific level.

This kind of compensation is used at large scale in rural parts of the country. National Grid have a relatively weak network within Cornwall and have installed capacitor banks to generate enough reactive power to support the voltage. In recent years the situation has reversed due to large amounts of distributed generation being installed. The new cables on these sites act as capacitors, creating a leading power factor. STATCOMs or energy storage installations have the potential to adapt to these changing requirements, while the capacitor banks are no longer useful.

A significant advantage of this use case is that the capabilities of an inverter to provide active and reactive power are linked, but non-linearly. This contrasts with all the other use cases, where any capacity dedicated to one use case lowers the capacity available to another use case proportionally.

14.2 Method

RESolve had a voltage control mode, which allows the user to set a target voltage and a voltage slope. The voltage slope set how much the system responds to deviations from the target voltage, and maximum limits on the amount of reactive power the system is allowed to respond with.

Setting the target voltage was a little more challenging than first envisaged, as the battery PQM voltage readings were imbalanced. Initially it was thought that this could be due to a

Figure 32:Battery and Solar phase voltages

problem with the inverters of the battery, but there was no trace of imbalance from the readings provided by the Copley Wood PQM. Thus, the imbalance appears to be a zero error on the phases. The magnitude of the changes in voltage was accurate, but the voltage control mode always takes its set point from the highest of the three phases. This resulted in needing to set a target voltage that was higher than the statutory legal limit, as the real voltage was several hundred volts lower.

The difference between the solar park and the battery phases are clearly visible in Figure 32. While there are imbalances on the solar phases, these are within 50 Volts of each other, while the difference between the highest and lowest battery phases is approximately 550 Volts. However, the analysis on the effectiveness of the battery was expected to be recorded at the Millfield substation so the error was not considered a major barrier. In a commercial system that relied upon this technology, the issue would have been resolved as a matter of priority.

Anecdotally, from other research projects carried out by WPD, it was expected that the battery would actually be too small to have a significant effect on the voltage at 11kV. Trials with larger STATCOMS than 310kVAr had demonstrated that bigger generators and absorbers of reactive power were needed to have the regulating effect required. This, coupled with the unpredictability of the overnight voltages, altered parts of the trial plan. For instance, while the setpoint for some of the tests were set at the actual target voltage, allowing the battery to try and regulate its generation and absorption of reactive power, other tests were run with the target voltage deliberately set lower than normal. This meant the inverters were working at their peak rating all the time and should represent an average drop in voltage overnight, i.e. when the system is operating versus a standard night with no control mode active.

The next part of the testing was to combine the reactive power provision with the active power use cases. This was the most exciting of the use cases as it represented a non-linear trade-off, whereas allocating 100kW of active power to one use case meant 100kW was unavailable to another, allocating 100kVAr of reactive power to voltage control only ‘costs’ 17kW of active power. This results in a 0.945 power factor, close to the 0.95 lead/lag power factor usually offered for energy storage grid connections. If the battery is operating at a lower active output while maintaining the reactive 100kVAr the power factor will drop.

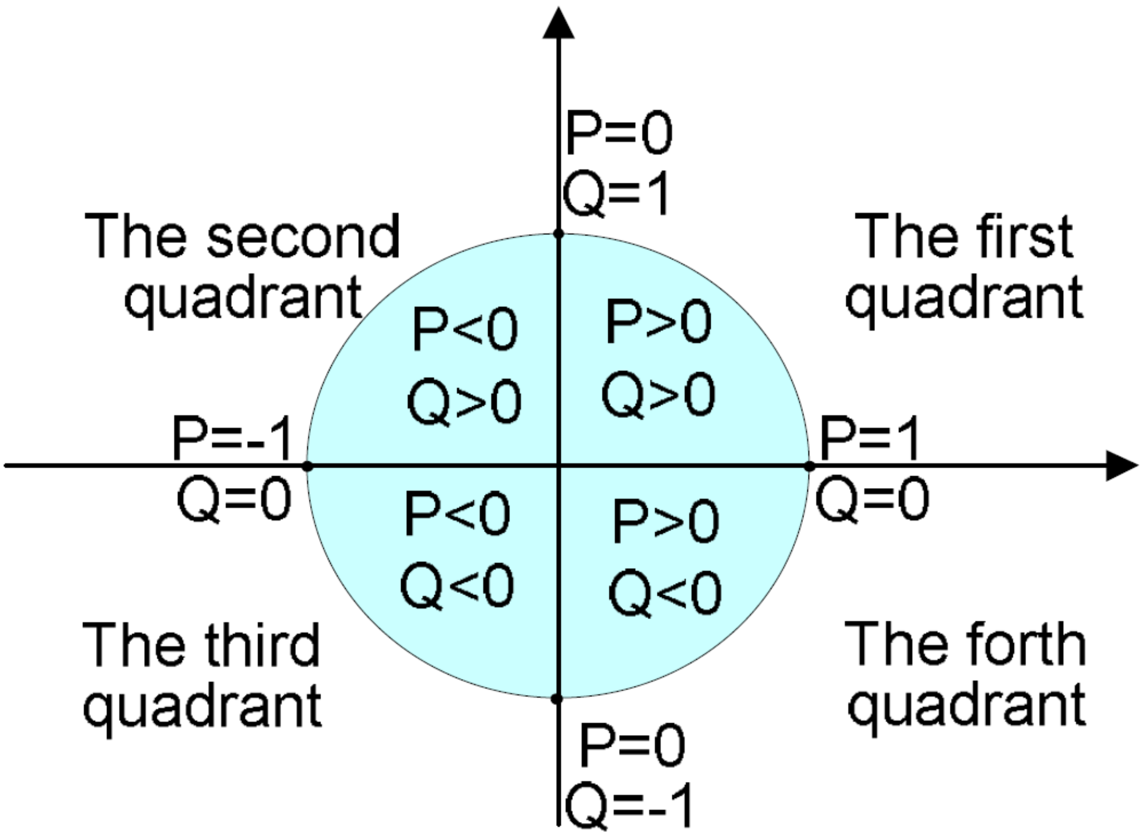

Figure 33: BYD P-Q Curve

Figure 33 is taken from the RESolve manual and demonstrates the relationship between active and reactive power capabilities of the battery inverters. The shape is symmetrical, meaning performance should be identical at leading or lagging power factors, or while importing or exporting active power. Figure 34 demonstrates the non-linear relationship between reactive and active power. The maximum active and reactive power output available from the battery inverters is 310kW and 310kVAr respectively. Currently generators and batteries can't take advantage of this flexibility as connection agreements

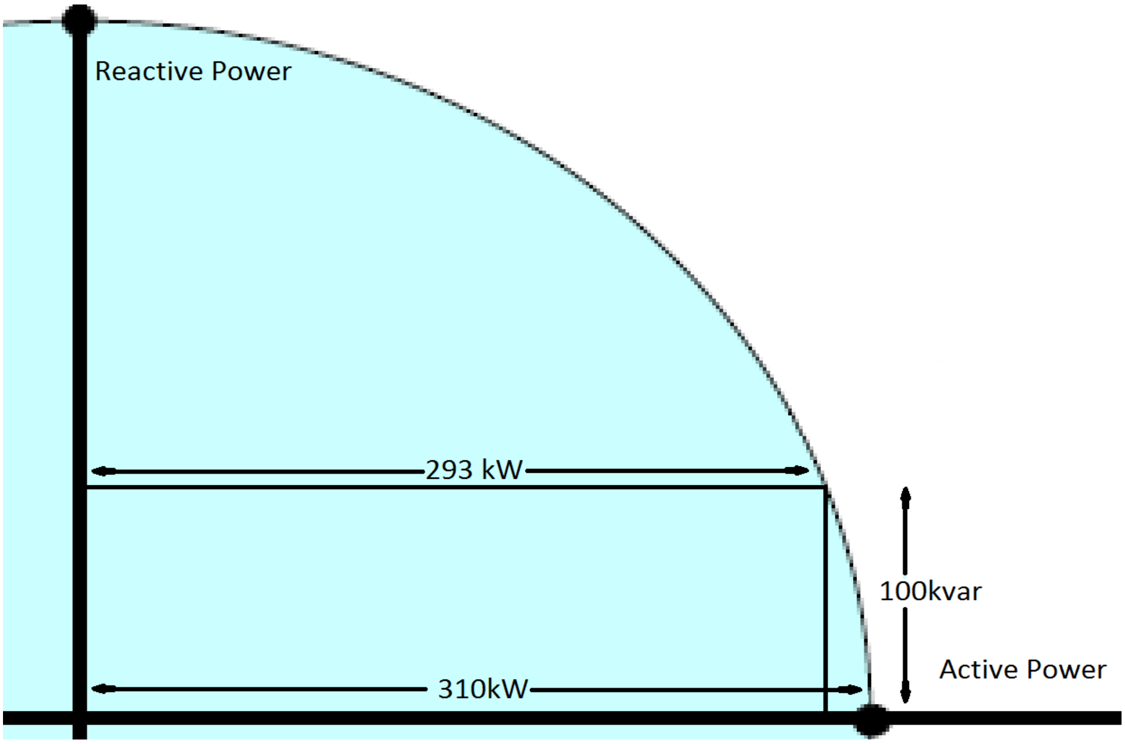

Figure 34: Active vs Reactive Power

limit the operational power factor to 0.95 leading and lagging. This means that for an active power output of 310kW, the maximum import or export of reactive power would be 100kVAr. However, any lowering of the active power output below 310kW would require a lowering of the reactive power output, to maintain the ratio. For the purposes of testing, the Copley site was given special permission from WPD to breach the standard connection limits. In some cases this could lead to a power factor of almost zero, as the inverters are capable of producing or absorbing reactive power with no active power flow.

The issue of offering less restrictive connection agreements is that it could allow poorly performing equipment to be connected at low power factors, creating more issues with grid stability. It is proposed that the best solution would be to add a clause into future connection agreements, stating that for the purposes of network support and only when under instruction from the DNO/DSO, the batteries can operate outside the standard power factor limits. This allows DNO's to keep standard generators and loads within the current power factor limits, while still opening up reactive power support as a potential market.

To try and evaluate the effect of the reactive power on the network, the overnight voltage recordings at the PoC were recorded and averaged, with standard deviation calculated. This was then compared to overnight sections when the reactive power mode was active. This was again looked at in isolation, with the battery doing nothing else, to make the effect easier to see.

14.3 Conclusions

After comparing the average, minimum and maximum voltages, and the standard deviation, there was no significant alterations during the tests. The following four graphs (Figure 35-Figure 38) demonstrate the target voltages set for the schedules in red, as well as the variations from the control voltage from the untested days.

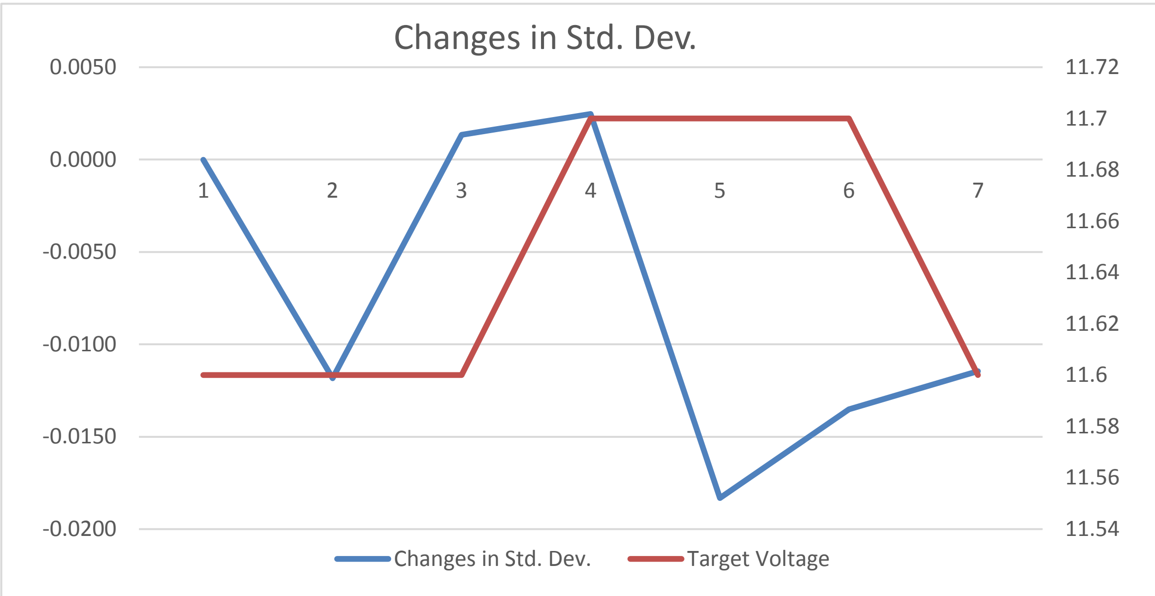

Figure 35: Changes in Standard Deviation

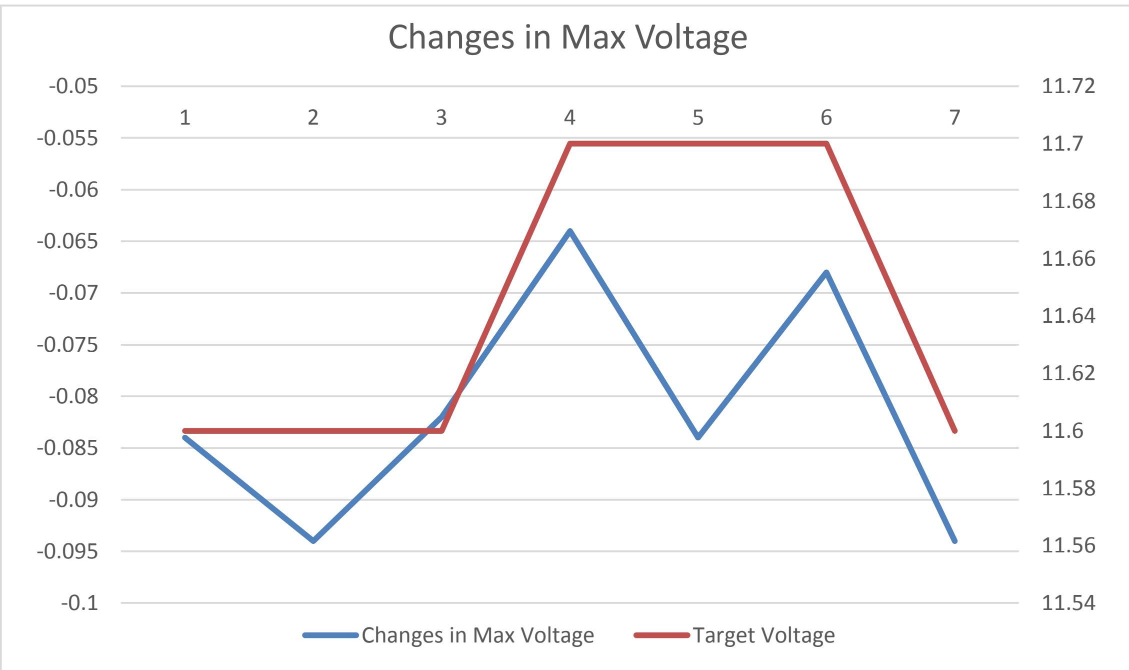

Figure 36: Changes in Maximum Voltage

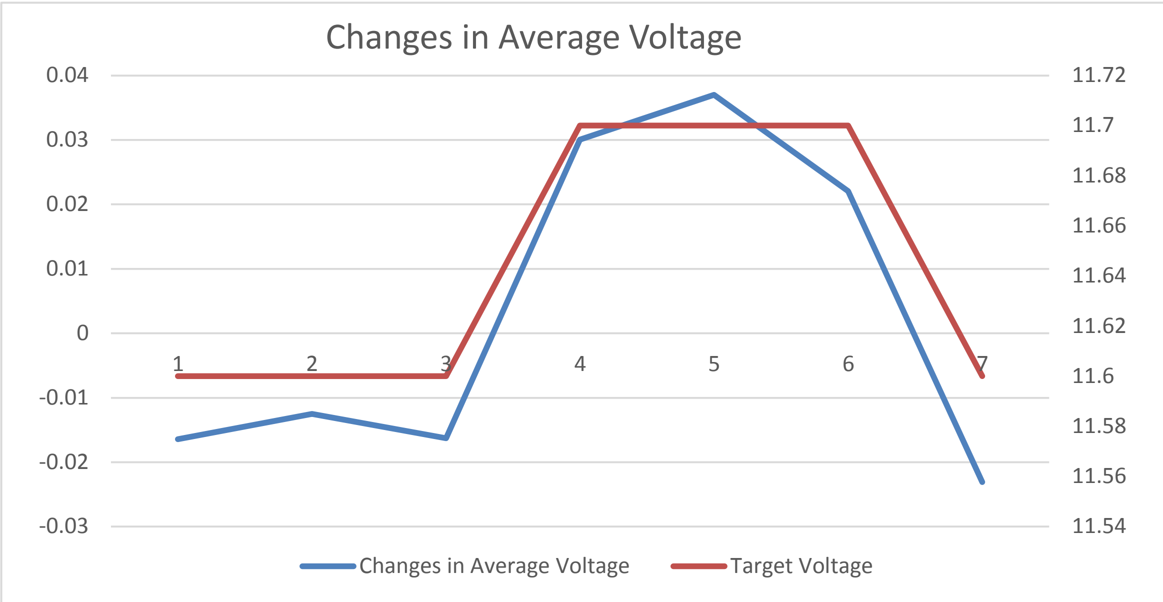

Figure 37: Changes in Average Voltage

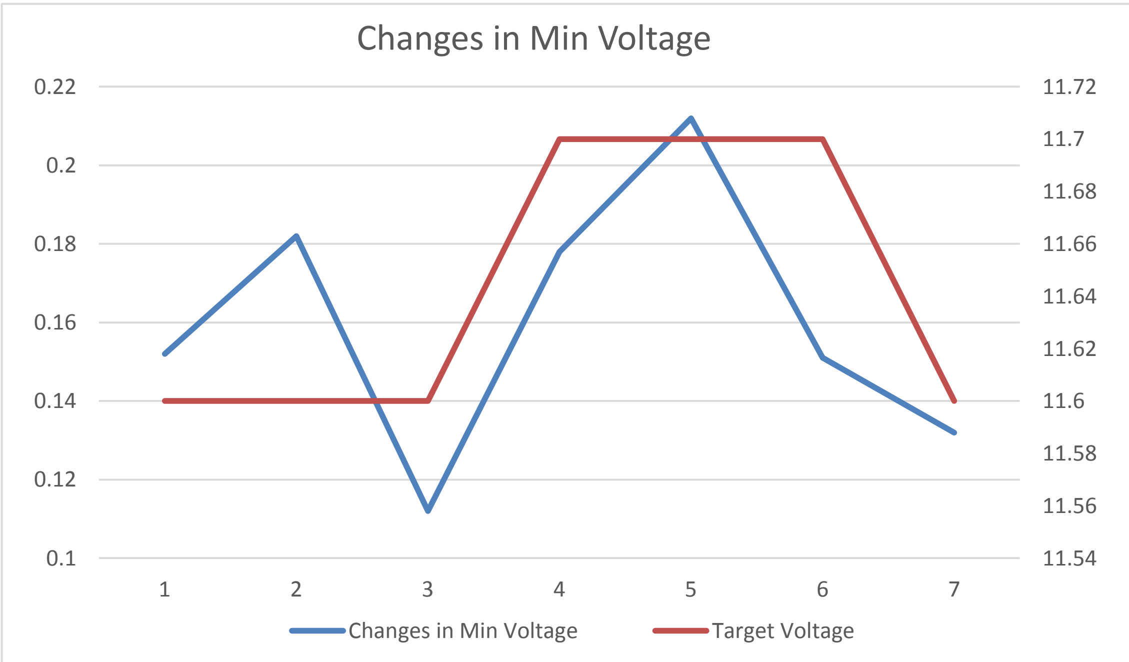

Figure 38: Changes in Minimum Voltage

It can be seen that the changes in average voltage matched the changes in the target voltage. In addition, the maximum and minimum voltages were decreased and increased respectively, which points towards the battery being able to curb the excesses of voltage within the system and affect the overall average. However, the changes in standard deviation aren't as consistent, with some test days showing increased deviation versus the control data. All of these tests were run with maximum reactive power available and no active power flow.

It is difficult to draw clear conclusions from the minimum/ maximum voltage graphs, as the control data is taken from a whole month of nightly data versus each individual nights testing. This could mean rare excesses of voltage are present within the control data. The less consistent change in standard deviation points towards this being the case. Overall the changes are very small, even when the setpoint was set artificially high to view the full effect of the battery. This backs up conclusions from other trials that a 310kVAr battery is

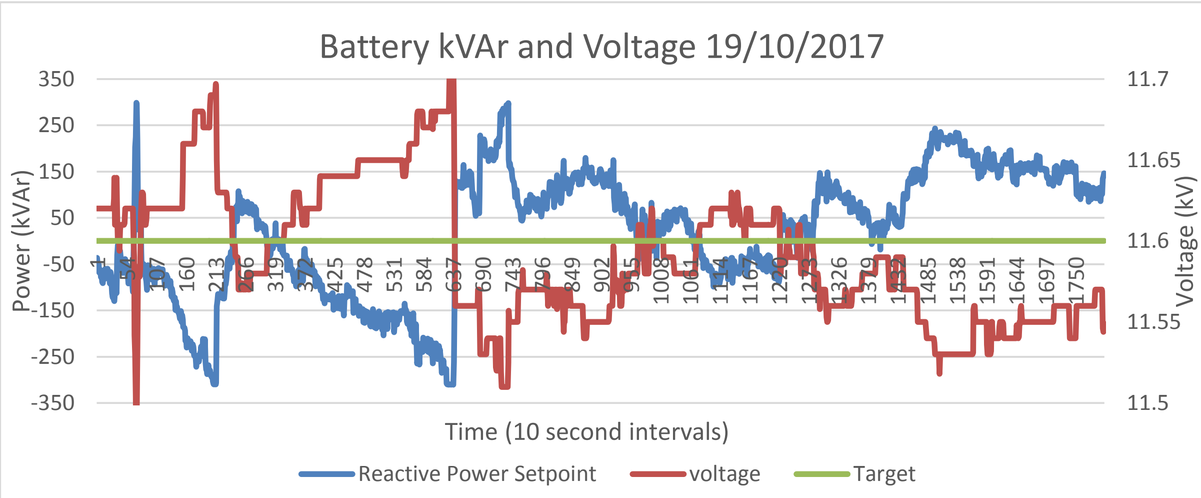

Figure 39: Reactive Power Control (kVAr vs kV)

too small to have an effect at 11kV.

The results suggest that the control algorithm functioned as desired, steadily increasing response as the grid voltage deviates further from the target voltage. The control slope can be set so that a full response is given when certain voltage thresholds are reached. The axis have been adjusted so that the target voltage is in the centre, and it can be seen that the battery reactive power flow is equal but opposite to any voltage changes away from 11.6kV.

This use case algorithm, fitted to a larger battery, should have the desired effects on the voltage. The fact that even this small battery was able to change the average voltage suggests that a larger system should be effective under automated control. The system could potentially be used under the new Power Potential scheme[18] being trialled in the South East by UKPN and National Grid. The trial is demonstrating the potential for existing

[18] https://www.nationalgrideso.com/innovation/projects/power-potential

and/or new generators and energy storage systems to provide active and/or reactive power on demand. The reactive power providers are expected to operate in an automated voltage droop control mode when not receiving specific signals from central control. This test battery clearly demonstrated the ability to operate autonomously and has set the groundwork for further software development. It should be noted that several installations running this software, working together, should be able to achieve a more obvious voltage stabilisation effect. This co-operative mode with many assets working together is likely to be more economically advantageous, as each generator would only have to take a small de-rating of active power rather than having an installation provide all of the reactive power with no capacity to trade in active power services. It also makes the service more attractive to DSO customers, as the risk of service failure is spread across many assets.

The opportunity cost of this service varies, up to the full £2.79 per hour if the full reactive power is committed. In addition to this, there is the cost of the reactive power. There are no special provisions for reactive power within the Copley Wood export agreement, meaning the solar park will be charged the default rate according to the WPD 2018 charging methodology[19], which lists 0.096p/kVArh as the charge for intermittent high voltage (HV) generation. For the Copley Wood battery, that means a maximum of £0.30 per hour[20]. The reactive power cost is almost negligible, and indeed these charges could potentially be zeroed if engaged in DSO support services, in a similar fashion to the way that power used for National Grid services is zeroed. A more realistic use case is the option to slightly de-rate active power to offer reactive power support, again using the 17kW/100kVAr exchange. This results in a de-rating of 5.48%[21] or, in monetary terms, £0.15[22] per hour. At this point on the active/reactive power curve, the battery needs to sacrifice less than 6% of its active power in exchange for 32% of its reactive power capabilities: an extremely efficient exchange.

It appears that the best way to take advantage of this has already been adopted by the UKPN Power Potential project where each DSO-connected generation asset offers as much reactive power as is viable, without having a major impact on the main revenue streams. In this way, it would be possible for the DSO to get this type of support at extremely cheap prices, although it does rely on contracting with multiple parties. This in itself could be considered an advantage, as a single unexpected outage won't cripple the reactive power support. It is also likely that this service can be procured at cheaper rates than installing capacitor banks or voltage reactors on the network, due to the opportunity to leverage existing assets. In addition, the flexibility of being able to absorb or generate reactive power means that there won't be any grid-owned stranded assets, unlike the capacitor banks on the National Grid in Cornwall.

[19] https://www.westernpower.co.uk/charging-statements 2018 Schedule of charges and other tables
[20] ((310kW*0.096p/kVArh)/100)
[21] (17/310)*100
[22] £2.79 x 5.48%

15. Results – Use Case 6 & 7 – PV export limiting and variable PV export limiting

15.1 Introduction

Please note this section covers two use cases, solar peak lopping (use case 6) and the glass ceiling scenario (use case 7). This is because throughout the research the peak lopping level has been adjusted to try and reflect the generation. This means that both use cases effectively use the same test and the same data.

Solar peak lopping/shifting is the 'headline' use case whenever combining solar and storage is discussed, as it appears the most logical way to increase the revenue of an existing park with a storage asset. The logic is that the system price is being driven down by the availability of solar power during the middle of the day, while the peak evening price is unaffected as solar plants have stopped producing at that point.

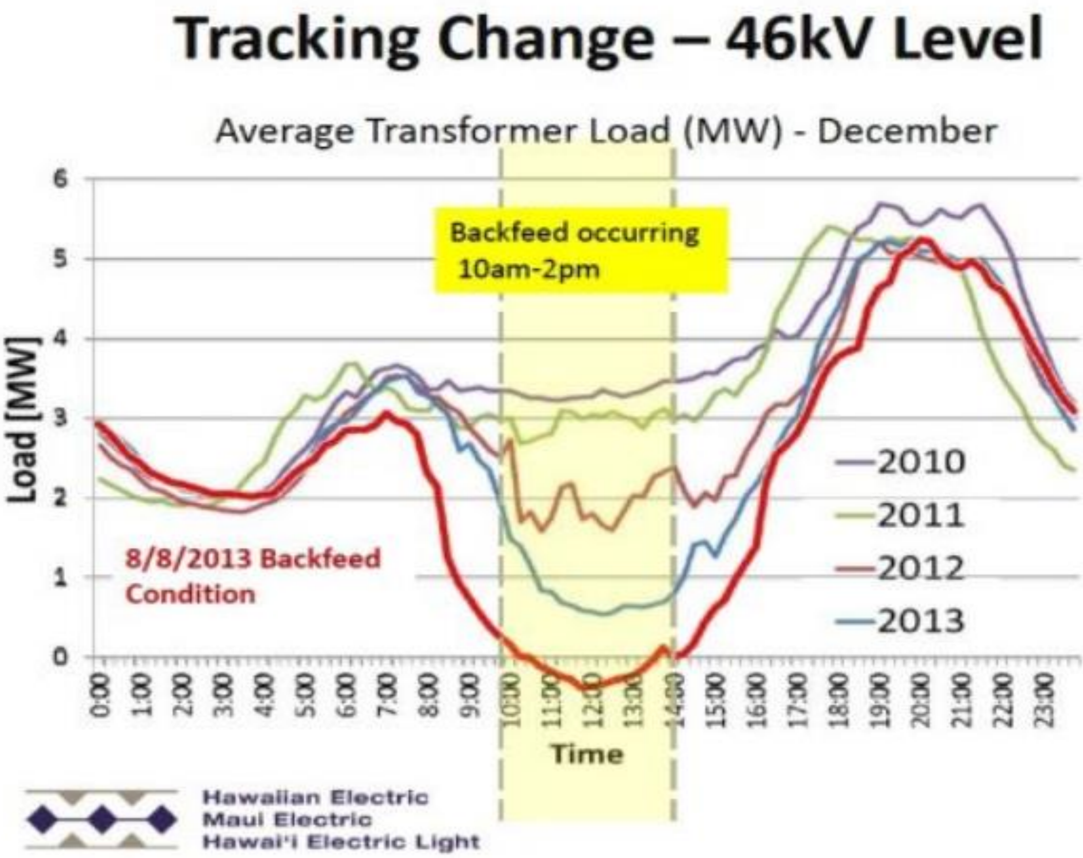

Figure 40: Hawaii grid load at 46kV, showing extreme duck or "Nessie" curve

In fact, higher levels of solar production could even increase the evening peak price, as the steeper increase in demand requires faster ramp up from generators. This puts greater strain and requires more fuel than the traditional increase.

An excellent example of the potential effects of high levels of solar on the grid can be seen in Hawaii where there are a series of island grids with a good solar resource. Figure 40[23] demonstrates the reduction of load in the middle of the day, demonstrating that on average in 2013 Hawaii was having to increase generation by a factor of 5 from 2-6pm. As shown on the graph there was a point where the system was backfeeding at the 46kV level, yet the peak demand at approximately 8pm was unchanged. This excess of power generation followed by high demand will be mirrored by system price.

In the UK, the ratio of solar to fossil fuel generation is much lower, but nevertheless there is a marked decline in price towards the middle of the afternoon. This could be attributed to a decline in mid afternoon demand rather than excess solar on the grid, but the specific cause is irrelevant as is remains an opportunity for increasing solar revenue.

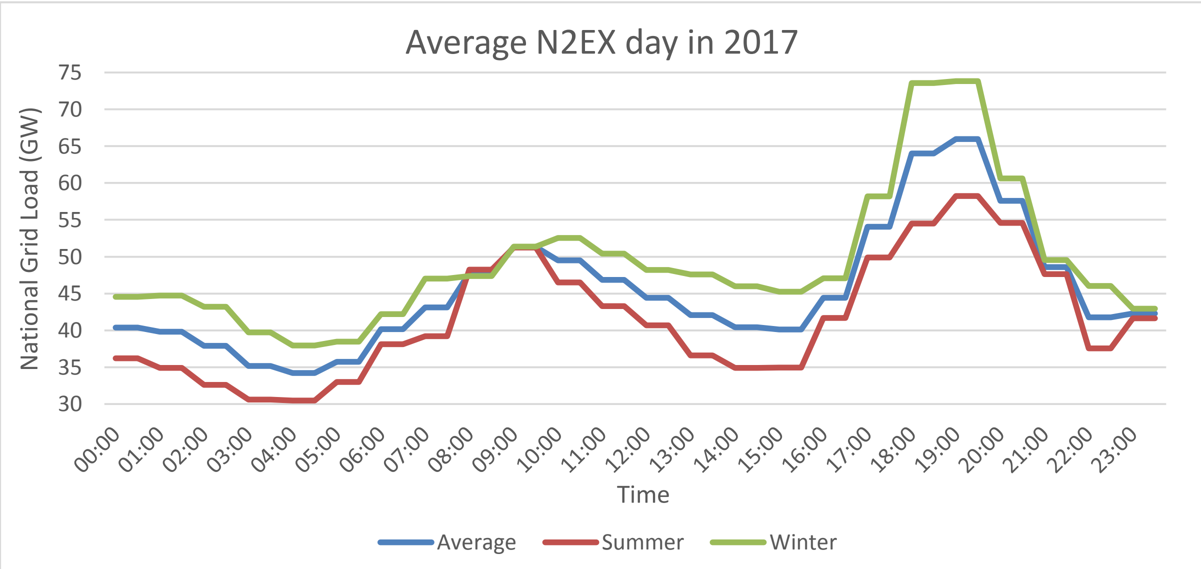

Figure 41: Average N2EX summer and winter days in 2017

[23] https://www.greentechmedia.com/articles/read/hawaiis-solar-grid-landscape-and-the-nessie-curve#gs.je9N978

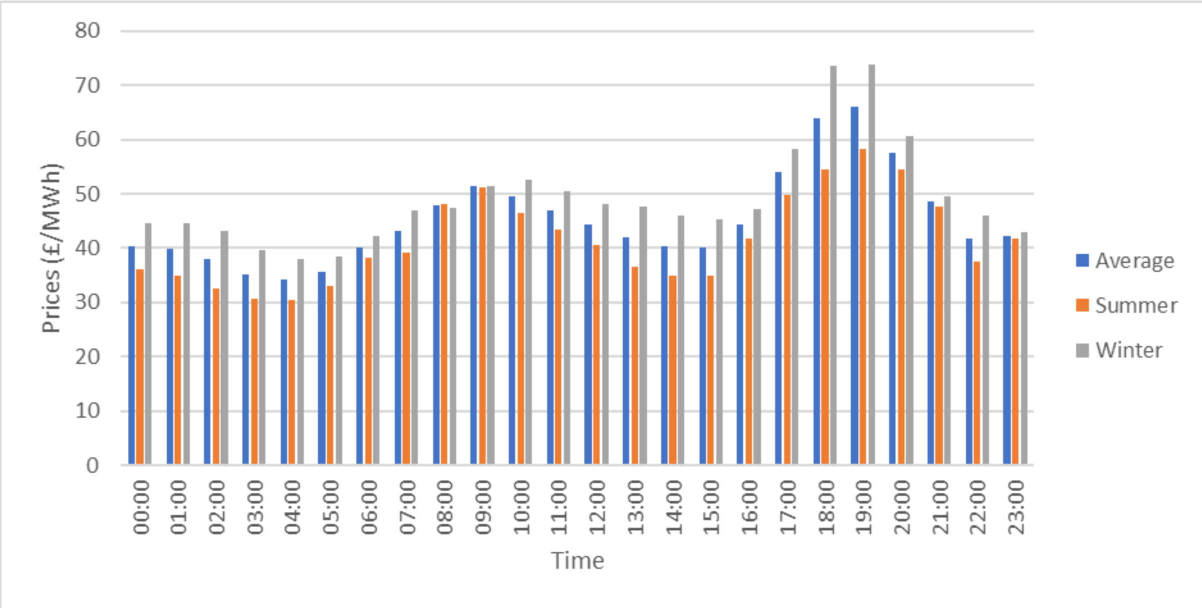

Figure 42: Average, Summer and Winter N2EX bar chart

Figure 41 and Figure 42 demonstrate the average N2EX price (the export price the Copley Wood Solar Park receives) throughout 2017, separated into summer, winter and annual average values. It can be seen that the mid-afternoon decline is noticeably steeper in summer, while morning prices at 9am remain consistent throughout the year. Compared to the shallower decline in the winter, it looks likely that some of this price variation is due to increased solar power during the summer.

It is interesting to note that the cheapest and most expensive times of day barely change between summer and winter, meaning a lot of the value should be able to be captured by a simple time-scheduled control system, rather than one requiring external variables for decision-making.

There are two types of peak-lopping that can be investigated, which will be termed ‘voluntary’ and ‘mandatory’. Voluntary is when a solar park that has sufficient grid capacity to export at full power, elects instead to store that power for more favourable prices later in the day. Mandatory would be for a solar park that is oversized for its grid connection and would be wasting power if the excess couldn’t be stored somewhere. Mandatory peak-lopping can generate more potential revenue than voluntary, as the lost-opportunity cost of charging the battery is nil. In the case of voluntary peak lopping, the lost-opportunity cost is the price of export at the time the battery was charging.

15.2 Method

The solar park has no constraint, so technically all peak lopping was voluntary. However, the income will be modelled under both conditions. The peak lopping algorithm detected the generation of the park from the Copley Wood solar PQM and, if it exceeded the user-set level, imported power into the battery. The variable user-set level means that this use case covered both standard peak lopping and the glass ceiling scenario.

Typically, the operator of the solar park does not have clear knowledge of the day-ahead generation, as knowing this wouldn't result in any commercial advantage or increase of return on investment. This data is available, if purchased, but it was outside of the scope of this project. Therefore, the local weather forecast[24][25] was the guide for setting the peak lopping level on a day-ahead basis. On occasion, if it was noted that the threshold was too high or low shortly after the schedule started, a second schedule would be programmed with a corrected threshold.

Throughout the testing time 800kW has been a standard peak solar generation threshold, as on sunny days with clouds the output would often pass through this level multiple times and the effectiveness of the battery controller could be reviewed. In addition, a 20% oversizing of solar park AC inverter power vs. the grid connection export capacity appeared to be the smallest level at which a battery would be installed to offset the loss. At a lower level of oversizing, it would be logical that the developer would simply not build as large a solar park or accept that the inverters would be software locked to a lower level, avoiding the extra cost and complexity of the battery.

15.3 Findings

Right at the beginning of the testing period, a small issue was found in the control software that caused the response from the battery to be exactly half of what was required to maintain the peak-lopped level, when reading the value from the solar PQM. This was quickly isolated and fixed after being reported (see section 7.1.1) . Close to the end of the testing period, it was discovered that a similar issue existed when using the site export as an input rather than the solar generation. This was not fixed as, with only one source of generation on-site other than the battery, the required response was the same so the Copley Wood PQM could be used without any impact. It is noted that the fix required is likely to be identical to the previous one and was only not implemented as it was deemed unnecessary for this project, as this functionality would only be needed if there were multiple sources of generation. Approximately midway through testing in 2017, some optimisations were made to the algorithm which increased the speed of response of the battery, reducing overshoots when solar power ramped up quickly.

[24] https://www.metoffice.gov.uk/public/weather/forecast/gcn4dy73f#?date=2019-01-10
[25] https://www.accuweather.com/en/gb/butleigh/ba6-8/weather-forecast/716133

Figure 43: Solar Peak Lopping second by second

Figure 43 shows that the site export is being limited to approximately 450kW despite the solar generation being significantly higher. The battery/control algorithm appears to struggle with the speed that the generation changes, leading to excursion above and below the target line. While this isn't harmful to any equipment, in a real-world situation it would lead to over-export, and according to the DNO specifications would result in the park being tripped off.

There were a few occasions where the generation was more than 310kW greater than the export limit, so any excursions during these times can't be attributed to the control system, as the battery PCS was too small to absorb the required power levels.

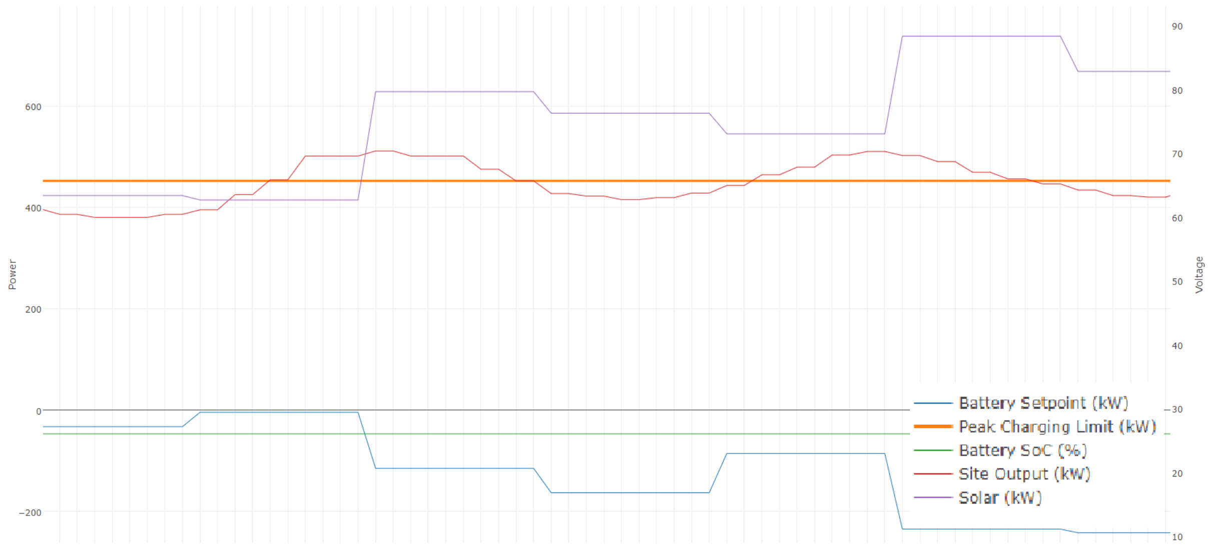

Figure 44: Detailed view of site export for one minute

Ignoring these excursions, and enhancing the graph, Figure 44 shows the extent to which the export exceeded the target. The two overshoots shown below peak at 61 and 60kW of deviation, with the time taken to bring the power back below the limit taking 13 and 16 seconds respectively. The WPD standard for export limiting devices is stricter than this, stating that the system needs to restrict the power within 5 seconds. Further to this, any excursion beyond 5 seconds will result in the backup system disconnecting the generation. These limits are also stated in the ENA G100 engineering recommendation[26].

In both cases shown above the park would have been disconnected for over-exporting, meaning the control system didn't respond quickly enough. One potential solution for this is setting the export limit 50kW below the actual target, as this reduces the deviation time significantly. This would result in importing more power than necessary.

This control method was only trialled for this project and not written to comply with any export specifications, therefore the lack of speed of response is not considered an issue moving forward. Were the system to be used commercially, the speed and accuracy of the control loop would be optimised to comply. Battery systems in the frequency market have sub-second response times, proving that 5 seconds is not a challenge for the technology.

The battery export limit was kept constant throughout the day, as the use case was based around a solar park that had a limited grid connection. An alternative use would be retrofitting a solar park with a battery and importing energy during the price decline that occurs towards the middle of the day and exporting it during the evening peak. Both scenarios have been looked at below.

Despite the slower than statutory response, the control system was reliably able to calculate a setpoint based on the generation, and the battery responded appropriately. The system reliably targeted the export limit until the end of the schedule or until the system was full, whichever was the sooner. The testing was done based on a constant export constraint each day, but there are also opportunities to create income from retrofitting batteries on existing sites.

The 800kW limit on sunny days has been used because the 199kW difference between the export limit and the grid limit of 999kW is well within the inverters capabilities. This avoids any potential anomalies resulting from operating at the peak of the inverters rated power. By setting the limit this high it also increased the likelihood that a drop in generation would be enough to cross the 800kW threshold, as the high DC:AC ratio of this park means a drop in irradiance has a reduced impact on AC output.

A simulation based on the second by second solar data has been performed by BSR for the whole of 2017, as if the battery had been programmed to always peak lop at 800kW. This gives 264 days when the solar park exceeded 800kW, and therefore 264 peak lopping

[26] http://www.energynetworks.org/assets/files/ENA%20EREC%20G100_amnd_1_final.pdf

opportunities. In total, this would have caused the battery to peak lop 85MWh of power over the year, an average of 0.32MW per peak-lopping day, although this is heavily skewed towards the summer.

The simulation was designed so that the battery would charge when the production was above 800kW and discharge when the generation fell below 800kW. The size of the battery was set to 310kW, with infinite capacity. This allowed the depth requirements for a 20% peak lop to be investigated. Figure 45 shows the required battery capacity for each of the 264 days.

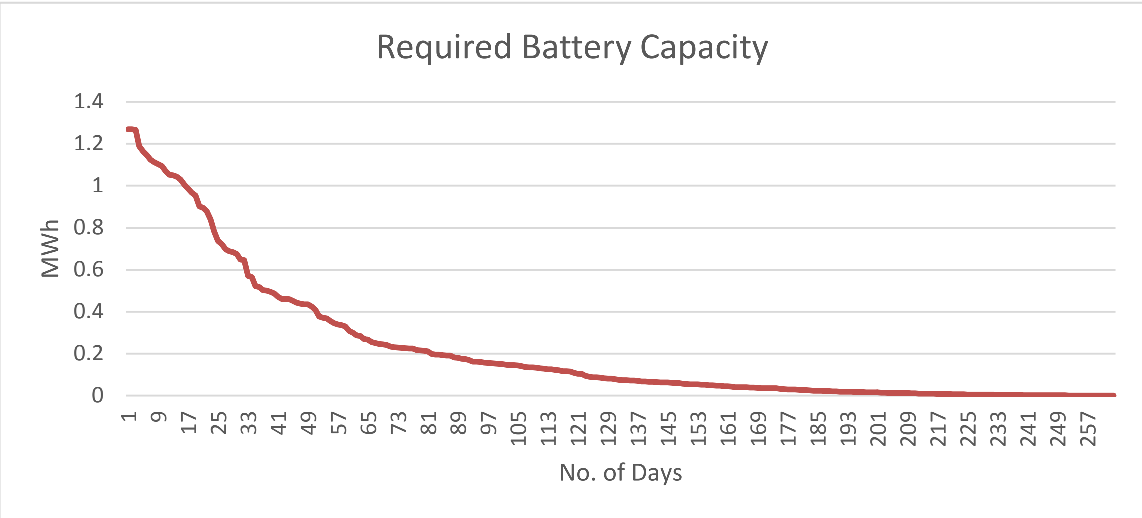

Figure 45: Required Battery Capacity for Peak Lopping with Export

On the highest production days, the energy required to be stored is in excess of 1.2MWh, equivalent to over 6 hours of maximum peak lopping 1.19MWh[27]. However, the requirement for this level of storage declines quickly: a battery of 600kW would be sufficient storage for all but 33 days of the year. This is based on the presumption that the battery would attempt to discharge every time the power dropped below 800kW. Rerunning the simulation, on the assumption that any peak lopped power would be stored until the higher price of the evening, yields the results as shown in Figure 46.

[27] (999-800 = 199kW) x 6

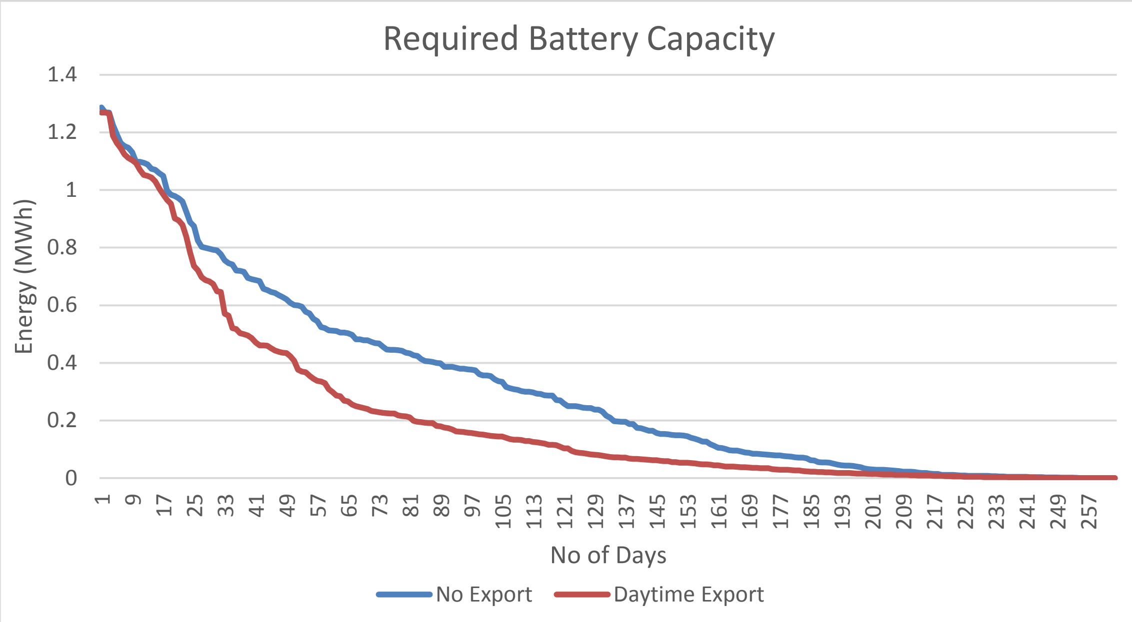

Figure 46: Peak lopping energy with and without export

The peak requirement is practically identical, pointing towards there being little opportunity for daytime export on the sunniest days, but the drop off in requirement is much shallower. Here, the 600kW battery would be insufficient for more than 50 days.

This raises the question as to what to do with the energy once it is stored. If it is used to 'solar firm'[28], with export during the day, this lowers the size requirement of the battery but reduces the amount of energy available to export during the evening peak. If the battery is only peak lopping and aiming to export in the evening peak, then a large amount of energy will be wasted on the sunniest days, but the storage can take advantage of the higher price.

Had the solar park been constrained at 800kW as shown above, and the battery had stored as much energy as it could for the evening peak export of between 6-8pm, the battery would have earned just over £3,750, using an assumed 88% round trip efficiency. Sized at 620kWh the battery would have been able to capture 83.6% of the energy generated over 800kW, equal to 71MWh[29]. This is compared to the baseline FFR case, £2.79 from 7am to 11pm (due to the 4 hour blocks that FFR is procured in). The revenue from FFR in the same time period would have been £16293.60: 2/3 of the annual FFR revenue[30]. If the setpoint was 900kW, the income drops to £1,720, but 100% of the energy can be captured. This would indicate an oversizing of the battery for the application, as the maximum peak lopped energy is 580kWh, meaning the battery only reaches 93% SoC throughout the year. The difficulty in making a scenario such as this work is that there are additional costs associated

[28] Solar firming: the battery will export power to maintain a stable site output when solar generation drops due to cloud cover etc, making site output more stable and predictable.
[29] 83.6% x 85MWh
[30] £2.79 x 16 x 365 **or** using the total from Table 5, £24,440.40 / 6 x 4

with oversizing a solar park (physically building more, increases in rental area etc.) plus the cost of the battery, which then operates in a single use case throughout the day without access to a flexible grid connection.

A simulation for the year of 2017 was run based on the battery importing between 2-4 pm and exporting between 6-8pm. This meant a daily uplift for the solar park with the simulation showing up to 620kWh could be stored to take advantage of the increased price, at the 88% round trip efficiency recorded in the decommissioning tests. In the summer the battery can't absorb all of the generation available and during the winter there isn't sufficient generation to fill the battery within the two-hour charging window. This simulation is based on a simplistic time-based scheduling system, with the peak lopping target set to zero. This is in line with the capabilities of the control system on the test battery.

If the battery had been programmed with this peak lopping schedule for the entirety of 2017, it would have created a profit of just over £2,750. Using the final degraded capacity of 676kWh, this would have resulted in 208 cycles per year, or 200 cycles using the original capacity of 702kWh. It is likely that the battery could operate for longer on this regime, as despite operating at high and low state of charge, this is only half of the cycling predicted to be used by FFR batteries. However, each chemistry and manufacturer is different and would want to run simulations of operation to give a definitive answer for a particular use case. The average income uplift per MWh peak lopped from the solar park is £17, a huge increase over the average export price during the afternoon.

The battery was filled to between 619-620kWh (effectively fully charged) for 58 days of the year and over 600kWh for 93 days of the year. During the winter, the battery often went almost unused due to low solar generation. During December to February, the generation only reached 800kW 27 times, and only 3.35MWh of generation was available to be peak lopped. This is less than 4% of the 85MWh peak lopped over the year, despite being a quarter of the time available. There is potentially a greater upside to be captured in the winter due to greater price volatility and greater price spread, so the peak lopping could be set for a longer period, but this would involve more advanced automated control systems predicting whether the battery will be fully charged between 2-4pm or not. The opportunity cost of this peak lopping mode is easily calculated for 6 hours between 2-8pm each day. In 2016 this opportunity cost, at £7.19 per hour, is £15,746.10. In 2018 figures, at £2.79 per hour, the cost is £6,110.10.

This opportunity cost seems high, certainly well in excess of the income from peak lopping. However, the battery would be unable to access the FFR markets during the day while the solar park is exporting. There was no excess grid capacity available for the battery to be able to operate independently. This then demonstrates a viable income stream for a period of the day when nearly every other use case can't be accessed. The advantage of this use case is that the income is expected to increase over time, as solar further pushes down midday prices and the increased ramping required from traditional generators in the evening increases the peak cost.

If the operator of the solar park/ battery purchased day-ahead solar irradiance predictions, they could dynamically adjust the peak lopping window. As the price remains below the evening peak for most of the day, an extension of the peak lopping window, starting earlier in the day, would still provide extra income. The simulation was adjusted so that if the battery wasn't fully charged between 2-4pm, then the day was re-run with the window starting half an hour earlier. This was continued until either the battery was fully charged or the window started at 8am. At that point, the algorithm recorded the day as being too low on generation to fully charge the battery. The battery was then discharged at its maximum discharge rate between 6-8pm, continuing for the full two hours if fully charged or stopping early if not. This gave a potential income of over £3,800 after round trip losses, a 38% increase in income. It required 323 cycles at 676kWh or 311 at 702kWh. This is a 55-56% increase in cycles for a 36% increase in revenue compared to the more simplistic 2-4pm peak lop. It depends on the terms of the guarantee and technology as to whether this increase in income is worth the use of the extra cycles. If the battery is being underutilised vs. the agreed cycle use within the guarantee, then it would be prudent to use them to bring in additional revenue, but it is unlikely to be worth oversizing the battery to allow additional capacity decline from this 36% increase.

15.4 Conclusions

Even with simplistic time-based control systems an energy storage system can generate income from peak lopping, increasing the value of the power from the solar park. Assuming that the battery and solar park are one entity, this provides a boost to the returns of the solar park and futureproofs the solar park against further downward price pressure during the middle of the day. Should the duck curve (the reduction of power prices in the middle of the day) become more pronounced, then the economics of new-build solar parks will be significantly reduced. Typically, existing solar parks are more protected from the effects as their income is partially made up of guaranteed subsidies, reducing their exposure to the open market power price.

It is interesting that £2,700 can be generated with no external price signals or solar prediction, and only a two-hour import window per day. This is bordering on guaranteed income, as long as demand in the middle of the day is significantly lower than the evening peak, the price differential will be maintained. In fact, every prediction of the future energy mix suggests that the price variation will increase, with a shortage of available capacity in the evening peak. This means the income of £2,700 should increase over the lifetime of the battery, but currently would not justify the installation of energy storage aimed at solar peak lopping. It is a good 'standby' use case though - if a battery can generate income from other services during the evening and night then it can also provide this uplift to the solar park without impacting the grid connection.

The more profitable use case is peak lopping generation above a grid export limit. As this is energy that could not otherwise be generated, the 'cost' is nil (the cost to the solar park of generating extra units is negligible, it doesn't wear the equipment unlike, for example, a gas turbine), so all revenue generated by exporting this power later in the day is extra revenue, and 'saving' the solar park from being restricted. The limitation of this use case, and the reason it only generated £3,750, is on the days the solar park doesn't reach the peak lopping limit then the battery sits idle. It is considered that a commercially operated battery would combine two options, slightly oversizing a solar park (or running an inverter without clipping) and allowing the battery to absorb the excess and, on days where this isn't achieved, then importing solar power between 2-4pm to take advantage of the arbitrage. This works the battery harder, but with guarantees of 10 years, at approximately 800 frequency response cycles (or 400 arbitrage cycles) being available from various manufacturers within the industry, using the battery at less than this amount results in underutilisation of the asset. This should ensure a blended income of the two figures and mean the battery is working every day.

The peak lopping extension simulation revealed that a 36% increase in revenue could be realised from a more advanced control system, although this results in 55% more cycles. If the extra cycles are covered by a manufacturer's warranty then this can still be a viable benefit, but in real terms, the increased damage to the battery's capacity and operational life is unlikely to be matched by the income increase. The cycle number is still below the standard industry offering noted above. In addition, there are several days where the battery is not charged to 100%, meaning each cycle does less damage than a flat-full-flat arbitrage cycle.

Realistically this use case on its own is not going to be sufficient to create investment opportunities for new build batteries. It does prove that the control systems today are capable of generating income from solar peak lopping, as well as providing export limitation without wasting the excess energy. It demonstrates that there is potential for improving the revenue of solar parks. If Copley Wood had been a subsidy-free installation, it has the potential to generate revenues of approximately £60,000 based crudely on the N2EX power price in 2017. The £3,800 income from the advanced peak lopping simulation would represent a 6.3% increase, which could be further enhanced with the minimal oversizing suggested above.

16. Results – Use Case 8 – PV power quality improvement (Ramp Rate Control)

16.1 Introduction

Ramp rate control is aimed at reducing the variability of solar output, and therefore reducing the voltage fluctuations and increasing the power quality on the network. The ramp rate control generates when there is a large fast drop in solar output and imports when there is a significant increase.

16.2 Method

The ramp rate control method is relatively simple, with only three options available to set, i.e. the maximum increase and decrease ramp rate, and whether the 'Auto SoC' adjustment is enabled. The minimum ramp rate available in RESolve is 1kW/s, which would only be exceeded at the Copley Wood Solar Park on high irradiance days with variable cloud cover.

The ramp rate of a solar park is the change in generation levels, or variability of output. For instance, a cloud moving over a solar park can cause it to reduce output by up to 80%, before increasing just as quickly. By controlling ramp rate, a much smoother site output can be achieved, reducing voltage fluctuations. A greater ramp rate setting would have minimal effect on the ramp rate and so wouldn't provide useful results for the tests.

Initially the test schedule didn't have a provision for testing the ramp rate control independently. Instead, it was to be combined with other use cases such as use cases 6 & 7. However, the combination method's internal calculation caused interference between use cases, most notably solar peak lopping, meaning ramp rate control had to be tested on its own. This was only identified as an issue late in 2017, making it challenging to satisfactorily complete the required testing, i.e. ramp rate limits are only triggered on sunny days, which are relatively rare in winter.

To increase the chances of capturing sunny test days, without delaying the rest of the testing programme, one day a week was selected that would be dedicated to ramp rate control. This meant that the 8 tests would cover a two-month period, extending the test times to a time of year more likely to experience sun. The rest of the test programme was operated on the remaining days as usual, without ramp rate being included in the combination methods. This allowed sunny days to be captured, demonstrating the use of ramp rate clearly.

A ramp rate of 1kW/s (assuming infinite battery size) means the output of the solar park would take over 16 minutes to reach peak export from zero[31]. This provides sufficient time for the network tap changers to respond and prevent them from adjusting up and down every time a cloud reduces output.

The Auto SoC adjustment feature was left disabled for these tests, instead the battery was charged to 50% SoC before the ramp rate schedule began so the amount of energy required for controlling ramp rate could be measured.

16.3 Findings

The battery response to changing solar production was both fast and accurate, and it is clearly visible in Figure 47 that the output is much smoother. The fluctuation was only significant enough to trigger the battery response on sunny days with variable cloud cover, but this is not considered an issue as typically low levels of solar export do not cause large voltage fluctuations on the grid.

[31] 999kW / 60seconds = 16m39s

Despite the effect of the control method on site output being easily plotted and observed, it is very difficult to quantify the impact on the voltage from the ramp smoothing at either the PoC or at Millfield substation. Despite the batteries responses, the voltage continues to rise and fall in line with solar production, rather than site export. Higher Hill Farm Solar Park (a significantly larger solar park to Copley Wood) is connected to the distribution network at the same point on the 11kV feeder. Higher Hill Farm's ramp rate was unaffected by the battery as it was not monitored by the battery system, and there were no other loads or connections on the 11kV feeder, meaning the voltage fluctuation at the PoC was entirely based on these two solar parks output. The Higher Hill Farm Solar Park is likely to experience very similar amounts of cloud over to the Copley Wood Park, as the Copley Wood Park was installed in two sections on the edge of the existing park as illustrated in Figure 48. The green highlighted area is the Copley Wood Solar Park, which the battery was linked to.

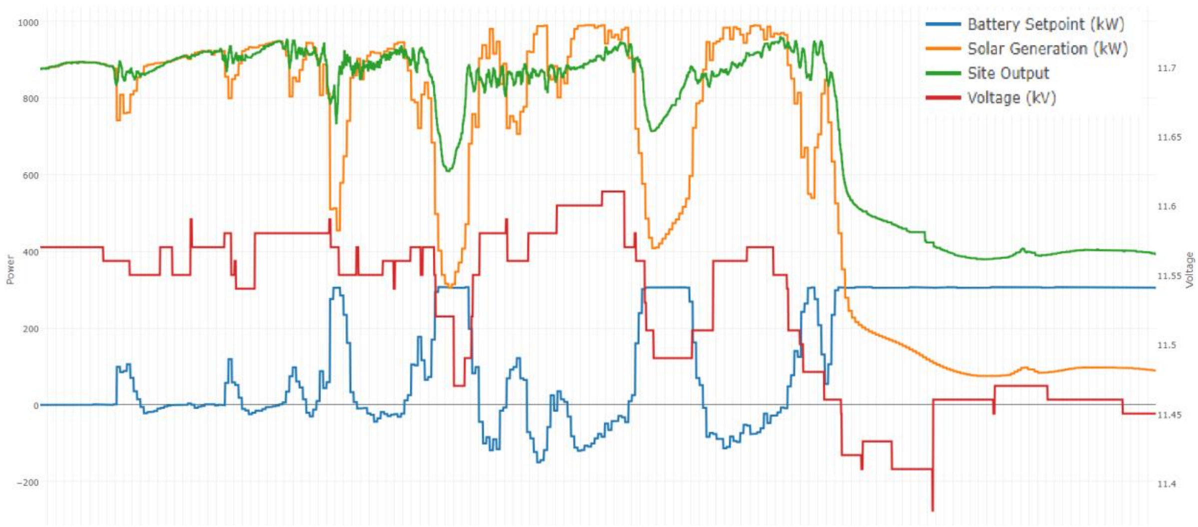

Figure 47: Output affected by ramp rate control

Figure 48: Map of Higher Hill and Copley Wood solar parks

The original Higher Hill Farm Solar Park has a DC rating of 5MW and an AC rating of 4.5MW, while Copley Wood has a 1.5MW DC and 0.999MW AC rating. Taking both parks into consideration, means that the battery had a rating of 5.6% of the total joint output at peak power[32]. As a result, when the solar is generating less, the battery could affect a proportionally greater amount of output.

[32] 310 / (4500 + 999) x 100 = 5.6%

The voltage at the PoC is driven by the amount of generation, as can be seen in Figure 49 with increases in generation driving up the voltage. It shows the Copley Wood solar generation in isolation, at second by second granularity. Figure 50 shows the PoC voltage, the Millfield substation voltage, and the kW of solar generation from Higher Hill and Copley Wood flowing through the 11kV feeder, averaged over ten minutes due to the limitations of

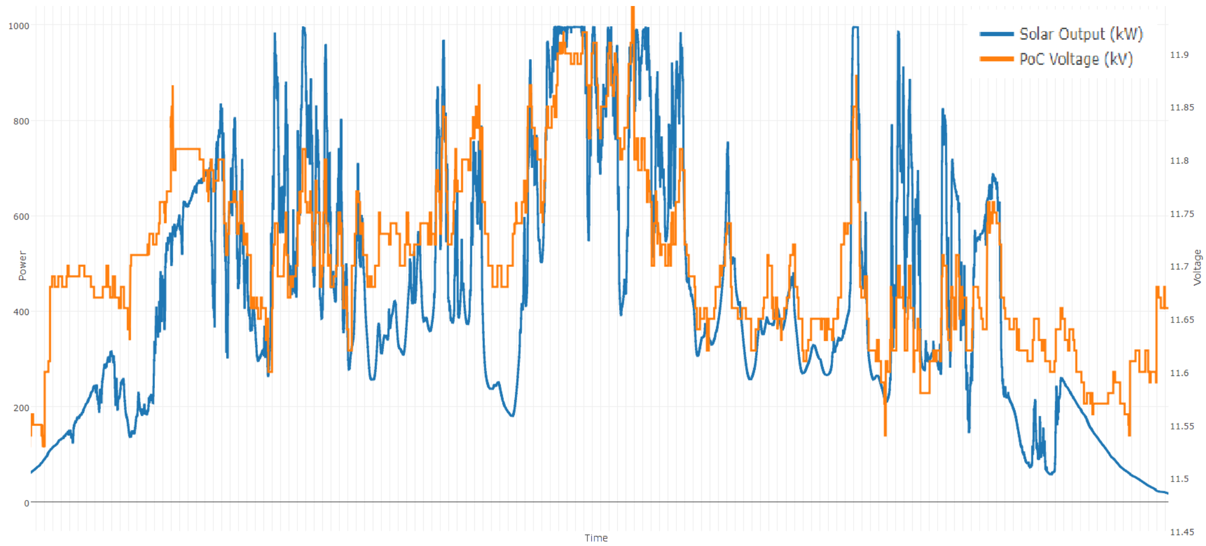

Figure 49: PoC Voltage vs Solar Output

the monitoring equipment installed at the substation. The voltage reading from the PoC has been reduced by 200 volts, as it was reading excessively high due to a configuration error. This adjustment puts it in line with the Millfield Substation voltage overnight, allowing easy comparisons, although the exact split between the configuration error and the voltage drop from the feeder isn't known. The change in voltage is the key comparable point here however, so this isn't considered an important issue.

It is interesting to note that the voltage rise during the day at the PoC is significantly higher than that experienced by the Millfield Primary. The peaks and drops in generation are clearly having an effect at the Millfield Primary, but not to the extent that would be expected by monitoring at the PoC. This demonstrates the grids smoothing effect on voltage fluctuations, likely caused by load on the primary substation. It also appears to reduce the need for ramp rate control, as the biggest excesses are already smoothed out. However, Figure 49 shows that the speed of voltage fluctuation is still very high (this is not as easily observed in the Figure 50 due to the averaging process in the Argand data).

Figure 49 illustrates a test day when the ramp rate control was in operation. It shows a section where the ramp effect has smoothed out large drops in solar production, making overall output significantly smoother. There are still dips in the green line (output) due to the solar variation being in excess of the maximum battery power.

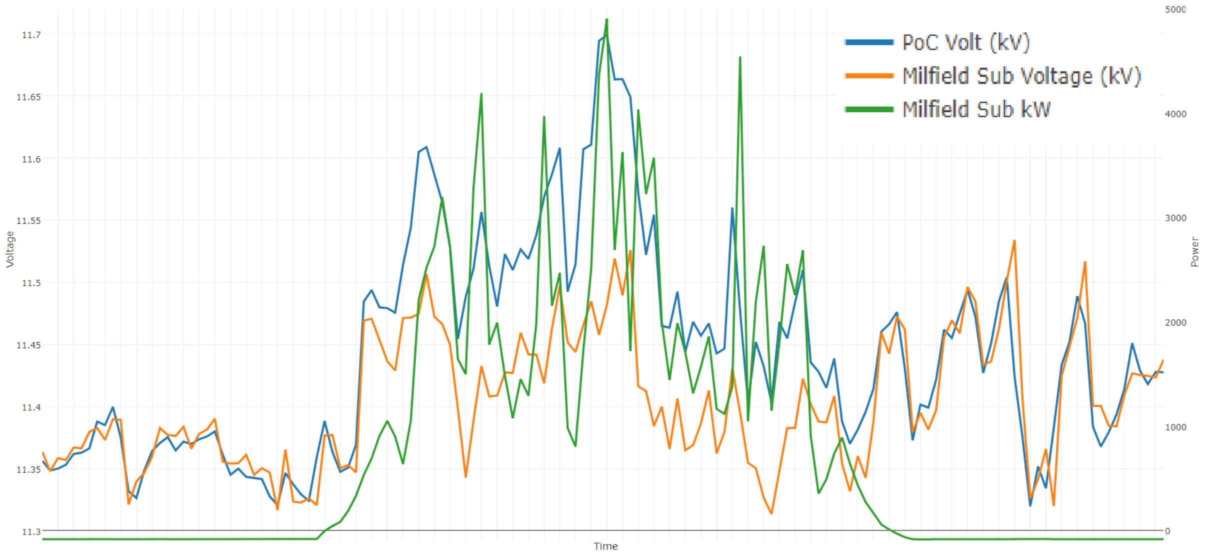

Figure 50: PoC Voltage, Millfield Substation Voltage and Millfield Substation Generation

This use case very clearly shows the limitations of the size of the battery. The energy storage system could only control ramp rates up to 31% of the solar output[33], and the solar park variation often exceeds this. As can be seen in Figure 49 and Figure 50, even with the fast response time, the battery was unable to fill the drop in generation which, as a result, would lead to voltage changes on the grid.

A battery size of close to 100% of the solar park's maximum AC output is required to guarantee the ramp rate control, as frequently the power ramps from below 10% to over 80% in 1 minute. This could be combined with solar firming to ensure a more reliable output and potentially realising a higher PPA price as a result. However, a blanket rule of 100% of a solar park's AC output is quite high, and realistically the battery would be sized to be the minimum AC rating and depth required to avoid excessive voltage spikes and drops on the network. The sizing then would be a product of discussion with the DNO, who would set the maximum voltage variance allowed on their specific section of network. Solar firming could help justify the business case for a larger scale of battery.

The voltage was still fluctuating in line with the variation in solar generation rather than the site output, which shows that smoothing the smaller site (Copley Wood) was not enough to affect the voltage. While this is disappointing, as the direct effectiveness of the ramp rate can't be isolated and measured, the reliability of the ramp rate control is encouraging. It is also easy to see that the voltage at the PoC is heavily driven by the solar generation during the day, with obvious rises and falls.

It can also be seen that although the load at the primary reduces the excesses of voltage seen at the PoC, it still experiences fluctuations driven by solar generation. It demonstrates

[33] 310kW/999kW X 100 = 31%

that a smooth solar output, whether by peak lopping and solar firming or by ramp rate control, would have a positive effect on voltage stability.

Ramp rate control can be viewed as a more cost-effective alternative to solar firming, in that it smooths solar output but doesn't provide a predictable flat generation profile. The cost of operating this use case is very low, as the requirements of solar ramping are often symmetrical in nature (i.e. a sharp drop in generation is matched by a sharp increase when the cloud clears), and on the test days the energy absorbed and released by the battery results in costs ranging from -£3 to +£3. This doesn't take into account the lost opportunity of FFR and Capacity Market, as a battery used for ramp rate control would be dedicated during the day to that task.

The cost savings of ramp rate control compared to peak lopping are due to the small amount of energy storage required for the system to work effectively. Even with SoC management disabled (when enabled, the SoC manager would bias the setpoint slightly to charge or discharge the battery back to 50% whenever it wasn't ramping) the biggest change was 17%, a total of approximately 113kWh over a day (albeit this value is inaccurate as the battery was too small to control the ramp rate). This was for a whole day of ramp rate control. However, it is possible to calculate the likely required depth of battery. If a 1MW AC solar park is connected to a 1MW AC battery and the DNO states that a 15-minute ramp rate is required to ensure network stability at this location, the worst-case scenario for the solar park is that the central inverter trips off and generation drops to zero when it is generating at peak power. At this point the battery tries to control the ramp rate from 1MW to zero output over 15 minutes. The average output over this time would be 500kW (125kWh in 15 minutes). Ramp rate systems will ideally sit at 50% capacity, in readiness for ramps in either direction, therefore in this scenario a capacity of 250kWh is required for a 15-minute battery.

Systems may be over-specified slightly to allow for batteries not being at exactly 50% when this condition occurs. 15 minutes is an extended ramp rate as most network tap changers respond within two minutes. At a 5-minute ramp rate the battery would only need to be 5 minutes deep (a 12C battery). This opens the space up for other technologies that do not suffer the cycling limitations of lithium, such as super-capacitors.

16.4 Conclusion

The battery needs to have a higher power capacity to appropriately control the ramp rate on a solar park. Many of the best instances of the control system responding to the sudden change in generation are also instances of where the battery had insufficient power rating to control the ramp rate properly. The obvious effect of the solar generation on the system voltage demonstrates that smoothing the output could be a major component of improving power quality on the network. The battery also clearly demonstrated much quicker response times and tighter responses to fluctuating generation than a tap changer could.

The picture of effectiveness is confused by the presence of a second, independent solar park (Higher Hill Farm). In addition, Higher Hill Farm has over four times the grid output of Copley Wood, dwarfing any alterations to its output profile. Unfortunately, WPD were unable to provide statistics on the increased frequency of tap changer replacement and maintenance that is caused by the solar parks. Without this value being quantified, the exact potential payment from avoiding placing stress on these systems cannot be calculated. On the other hand, the cheaper cost of such a short term energy storage device compared to standard lithium batteries means that a more stable grid is achievable with significantly lower capex. It has also been noted that, were DNO's able to quantify the cost of replacing tap changers due to solar installations, this cost could be passed on to the solar sites.

The project has been able to provide voltage readings from the PoC and at the primary substation supplying a clean feeder (no other loads or generation connected) with only solar generation affecting it which is in itself an interesting study on the impact renewable energy has on the distribution system. It demonstrates at a granular level how solar power can cause fast fluctuations that the distribution system cannot easily absorb and is a demonstration of why there is a limit to the proportion of renewable energy that can be connected without additional stability being provided.

It is clear that a much larger system than the current battery is needed to compensate for ramp rates, but it is also clear that stabilising the smaller parks output is ineffective. In a location where there are multiple renewable installations, the system should either monitor all of them to generate a set point or be attached to the largest generator. Otherwise the effect is too small to be useful, or indeed visible.

Another potential installation case is placing these short-term storage devices actually at a primary substation. This would take advantage of the existing smoothing effect of the grid, meaning the storage has to smooth out smaller variances. Building one larger ramp rate control unit in front of the meter is more efficient and cost effective than multiple units independently connected to solar parks and would effectively future proof that section of the network as more distributed generation was connected. This is one of the innovations that could be a product of the 'smart grid', with communication links to each generator informing the ramp rate control unit how to respond. Placing these at the substation also opens up the possibilities of smoothing out variable loads as well, as these also cause voltage fluctuations. This gives a 'two-for-one' effectiveness of the storage, meaning the developer should get more revenue while the DNO can pay less for each service. Finally, spare/unused substation space can be used for the storage asset, potentially making planning and electrical connection easier and cheaper. The regulatory framework for energy storage is likely to limit DNO battery ownership to being the developer of last resort and so, while locating batteries at primary substations may have technical advantages, many of the commercial issues would be the same as if the DNO were buying services from batteries located at a solar park. However, from a developers view there would be significant cost and planning advantages from being able to leverage existing substation space.

Ramp rate control is potentially very exciting. The effectiveness of very short-term storage opens up the market to a range of different technologies other than lithium, avoiding the degrading performance due to cycle issues. With parity solar on the horizon bringing ever higher percentages of intermittent generation to the UK grid, controlling that intermittency could prove vital to staying within voltage limits. This kind of storage is too small to take advantage of arbitrage, but it could be connected directly to the DC strings, further driving down costs due to not needing separate inverters. This can't be done with other energy storage use cases as they rely on the battery being able to charge independently of the solar park, to access the arbitrage market.

The contract for this service (or, potentially, the requirement for improved ramp rate control) would have to be long term for this market to develop. The energy storage would not only be dedicated to ramp rate control, but, crucially, it would also be designed specifically for that service. This means it can't swap to other income streams later, so is vulnerable to market changes. However, this doesn't present as large a problem as it initially appears to, as the requirement for ramp rate control is likely to increase rather than decrease. The cost of installation of this type is also heavily driven by 'Balance of System[34]' costs, which is a mature section of the market and unlikely to achieve large scale cost reductions, therefore being an early adopter isn't such a disadvantage on price as it is with battery times of an hour or above.

In conclusion, the ramp rate control market could be one of the lowest capex of any of the use cases, which brings forward its viability. Currently, battery cell cost is approximately 50-60% of a 1-hour lithium battery build, so reducing the storage requirement reduces the costs hugely. However, the market for this service doesn't currently exist, with DNO analysis specifically aimed at keeping the voltage fluctuations within tolerance levels.

17. Results – Combining use cases

The original expectation was that the battery, by fulfilling multiple use cases at once, could generate more income than by only sequentially switching between them. A matrix was created at the start of the project with use cases that were expected to complement each other (for example, network support and arbitrage both require export in the evening period). By stacking the revenues in the same time period, it was expected that the business case would be more viable more quickly, beginning the change from DNO to DSO as soon as possible.

With this in mind, the RESolve control software had a combination method included, which would allow two use cases with different set points to operate at the same time. The theory was that this would also allow the same asset to enter multiple markets at the same time, reducing risk. However, the research of this project has shown that usually one use case is always going to dominate as the 'base case'; previously being FFR contracts and now shifting towards arbitrage. The rest of the use cases are usually just providing additionality, rather than receiving equal credit for stacking the business case.

There are some scenarios where combining use cases can make sense: peak lopping could be combined with ramp rate control, with peak lopping preventing export above a set limit and ramp rate control activating as soon as generation dropped below that limit. The test to

[34] Items such as wiring, inverters, foundations, and potentially the grid connection although this is sometimes accounted for separately. It can be explained as any part of a solar park apart from the panels themselves, and the term has been extended to energy storage as referring to any component other than the battery cells.

prove this was attempted but unfortunately the ramp rate control mode was unaware of the set limit of the peak lopping, meaning that when there was a sudden change of generation the ramp rate control set point activates, even if the output was already smooth due to peak lopping. This is a simple logic problem that can be easily rectified and would allow a battery to prevent over-exporting, whilst improving the voltage characteristics of the network for the solar park, within the same asset.

The only other obvious successful active-mode combination is that of local support and arbitrage. This is because it is unlikely the local customer support would require full power output of the battery and could leave most of the inverter capacity available for any other use case. This only works successfully when the local support is required at a time when power prices are likely to be high (remembering local network conditions don't necessarily reflect national demands for power). There would have to be a priority-based system to prevent any other use case telling the battery to import during the time that the network support export was active, otherwise the battery would be causing increased demand while also trying to reduce it.

Other active power use cases can effectively split the battery into multiple markets at once. This reduces the revenue available from merely selecting the most valuable market for that time period and the only advantage is that it can reduce risk through diversification. This adds significant complexity to the control system for little reward, while the same diversification can be achieved by operating a portfolio of batteries with different target markets, connected to either the same feeder or substation. This solution also reduces the administrative overhead that would occur by entering parts of a single battery into separate markets. This technique has been used by some providers in the monthly FFR auctions to increase the likelihood that at least part of the asset is able to access this additional revenue.

The most valuable combination method available is that of the reactive power control mode. The non-linear trade-off between the active and reactive power makes this the most efficient combination, and as the DNO networks become more congested and more challenges appear, the asset will be able to both absorb and generate reactive power as required. This adaptability, available from a very small de-rating of the asset, could be a benefit to DNO's struggling to adapt to quickly changing network conditions and help increase network stability and reliability. The biggest advantage is that these set points can be changed with minimal interference to the main active power activities of the battery: the reactive power set point up to a limit can be placed in the control of the DNO, whose control room would alter it as required in response to real-time changes on the grid.

The combination of many uses cases at the same time on the same battery requires a great degree of complication and detailed studies of the network it has to operate in. The operation mode of switching between 4 or 5 use cases is already relatively complex, with concurrent operation adding huge numbers of potential interactions. At this point the bankability of the business case starts becoming an issue. Exploring the electrical and technical possibilities is one area, but being able to present it in a way that is able to attract investment capital becomes significantly harder, with the combination method layering more complexity on what is already an emerging business case.

18. WPD Flexible Power

WPD have launched their flexible power services for four CMZs across their licence areas. The target audience is for business owners with BTM assets that can provide flexibility (e.g. DSR, energy storage). The Copley Wood battery was on a larger scale than most of these existing assets are expected to be, but it still provides an interesting comparison for investigation.

WPD's flexible power map shows the expected levels of utilisation that are expected within each CMZ which shows significant variations between months. There is an expectation that Flexible Power would be an additional revenue stream for existing assets rather than a stimulant for new assets to be installed. This may change as the energy market evolves: several pieces of information suggest that the revenue structure and amounts will change in the future.

For this assessment it has been assumed that the availability of the Dynamic version of services would be required each day for two hours between 5 and 7PM. This has an hourly availability fee of £5/MW. Given this is the peak time that the DNO is most likely to need assistance with in the case of any fault, it appears appropriate to mitigate against this. However, the Dynamic service, with a utilisation charge of £300/MWh, is not guaranteed to be called upon.

As an estimate, it has been assumed that the service will actually be called upon approximately every 10 days, or 37 times a year. For simplicity it has been assumed that each event will require the full capacity of the battery over the two hours. It has also been assumed that the battery would need 2 hours to charge up fully before the event, giving a total of 4 hours required per day for this use case.

The times selected coincide with the afternoon arbitrage previously calculated in use case 1 – Arbitrage. This gave an annual profit of £270, without any opportunity costs taken into account. The availability payment per day is £3.10[35], with the FFR base rate would have generated £5.58 over the same period. Per year this equals £1130

However, the revenue from the utilisation payment also has to be included. This generates £6770 per annum[36]. This means the total income from the 4 hours of Flexible power use over a year is £8170[37].

This nearly matches the idealised arbitrage income for the entire year, while using only 4 hours of each day. It clearly demonstrates how lucrative the Flexible Power payments can be, although this revenue can only be generated within specific flexibility zones within WPDs

[35] £5 x 2 hours x 0.31MW
[36] £300 x 0.61MWh x 37 occurrences
[37] £1130 availability payment + £270 arbitrage profit + £6770 utilisation payment

licence area. In addition, there are no guarantees of how often the services will be called upon, or even how long the availability windows will be. It is thought that this is an excellent opportunity for existing assets in the right areas to benefit from a significant extra revenue stream, but unless more details and firmer long-term contracts can be offered, then this move to DSO may not be able to stimulate new installations of batteries.

19. Techno-economic learning.

19.1 SRI Technologies

Early in the project, some scene setting techno-economic analysis was undertaken to consider the potential for arbitrage and how battery size may affect the financial returns available. The analysis, carried out by SRI Technologies, used the Meteor add-in for Excel along with market price data and solar PV production data from BSR to optimise operator revenues for arbitrage. It determined what the “perfect” operating regime would have been for a 300kVA, 300kVAh battery and the value associated with this. As price forecasting is unlikely to allow for perfect battery scheduling this over estimates the benefits of arbitrage but does allow for comparisons to be made when other variables are altered. This found that there was limited potential for value from arbitrage and that if the round-trip efficiency fell below 70% then the costs were likely to outweigh the benefits. Other services might be more lucrative than arbitrage, such as managing the risk of imbalance charges for the provider of the PPA to the solar park. Providing such services might allow for a more favourable PPA to be negotiated.

The analysis was repeated for different values of battery power and battery capacities e.g.

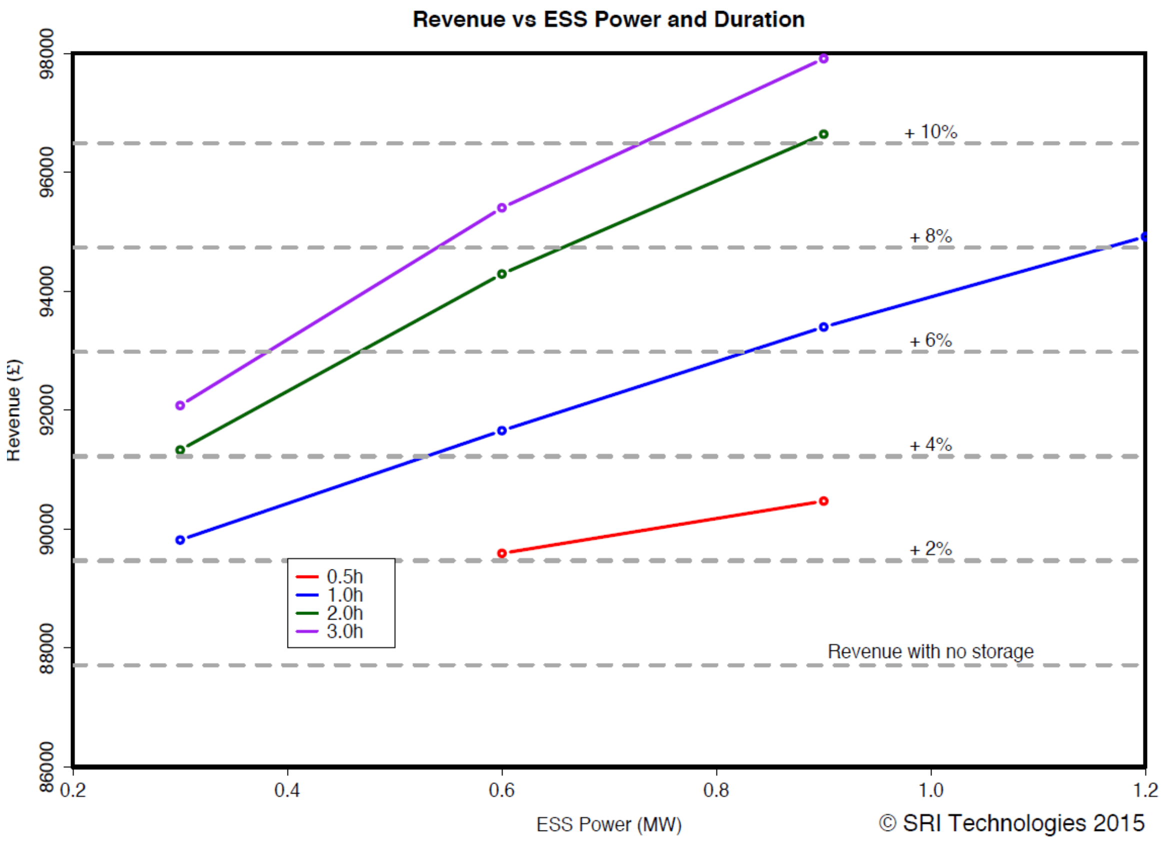

Figure 51: ESS Power and Duration

batteries that could sustain full power for ½ hour, 1 hour, 2 hours or 3 hours.

Figure 51 shows the diminishing returns for adding greater battery duration. For example, if a 0.3MW battery with 1hr duration is added, then the increase in revenue is slightly over 2%, whereas adding a battery of 2hr duration would add 4% to the revenue and a battery duration of 3hr would add 5% to revenues. To get better value for money the cost of the additional battery units to support the longer duration would have to show greater benefits of scale. It seems unlikely that battery durations longer than 2 hours would be cost effective.

In all cases, increasing the size of the battery added increases to the revenue benefit. It can be seen that the increased benefit from larger battery power also diminishes as battery power is increased e.g. the increased revenue from replacing a 2 hour 0.3MW with a 2 hour 0.6MW battery is greater than the increased revenue by adding another 0.3MW. This reflects that in modelling the perfect arbitrage scenario, that assuming many of the largest variations in energy price would have already been captured and additional trading would be at lower and lower price differentials.

The second phase of the analysis modelled the use of batteries to manage a transformer constraint.

Figure 52 shows the three worst winter days displayed as if they were consecutive to each other. The transformer support is triggered when loads exceed 16.1MVA to maintain the load on the transformer under the threshold of 17MVA. It can be seen that the peaky

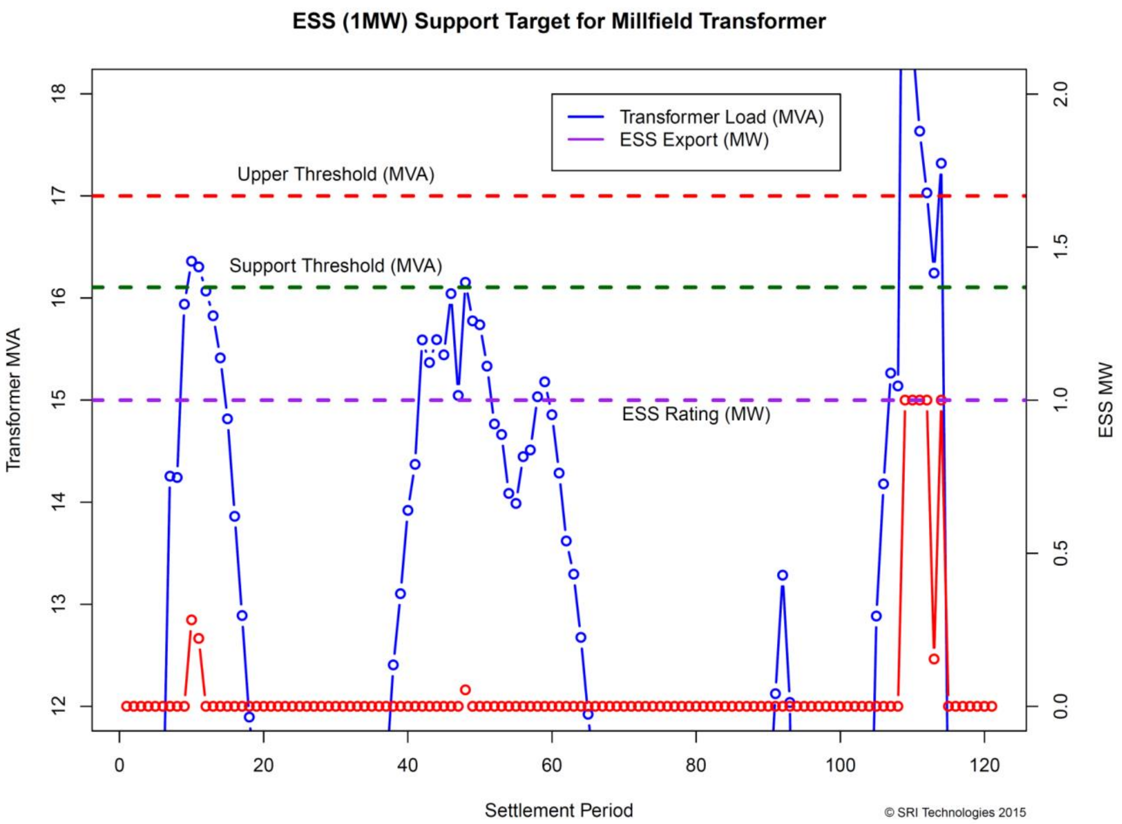

Figure 52: Support Target for Millfield Transformer

nature of the load profile (which varies considerably between days) would present a challenge for battery sizing. A battery sized to manage the requirements on the very worst day would not be required for a large part of the year.

Having set a threshold for the Millfield Primary transformer, the required battery power was calculated for each half hour and the ability of the battery to that support requirement was modelled for batteries of different sizes and capacities. For example if in a particular half hour period the transformer was over the threshold by 1MVA, a 0.5MVA battery would be insufficient even if fully charged. In addition, the modelling determined when the battery would not be able to supply the required power, as it was depleted by discharging in the preceding half hourly periods.

Figure 53 shows the 300 half hour periods with the highest target power values for the Millfield transformer modelled as being supported by a 2MW, 3MWh battery. The target power is shown in blue and the red circles represent the power delivered. For the majority of the half hour periods the battery is capable of meeting the target power and the red dots cover the blue line. However, there are a number of periods where the battery is depleted and cannot meet the target power. Additionally, where the target power is 2MW, this reflects the battery sizing and the required power to prevent transformer overloading would be greater.

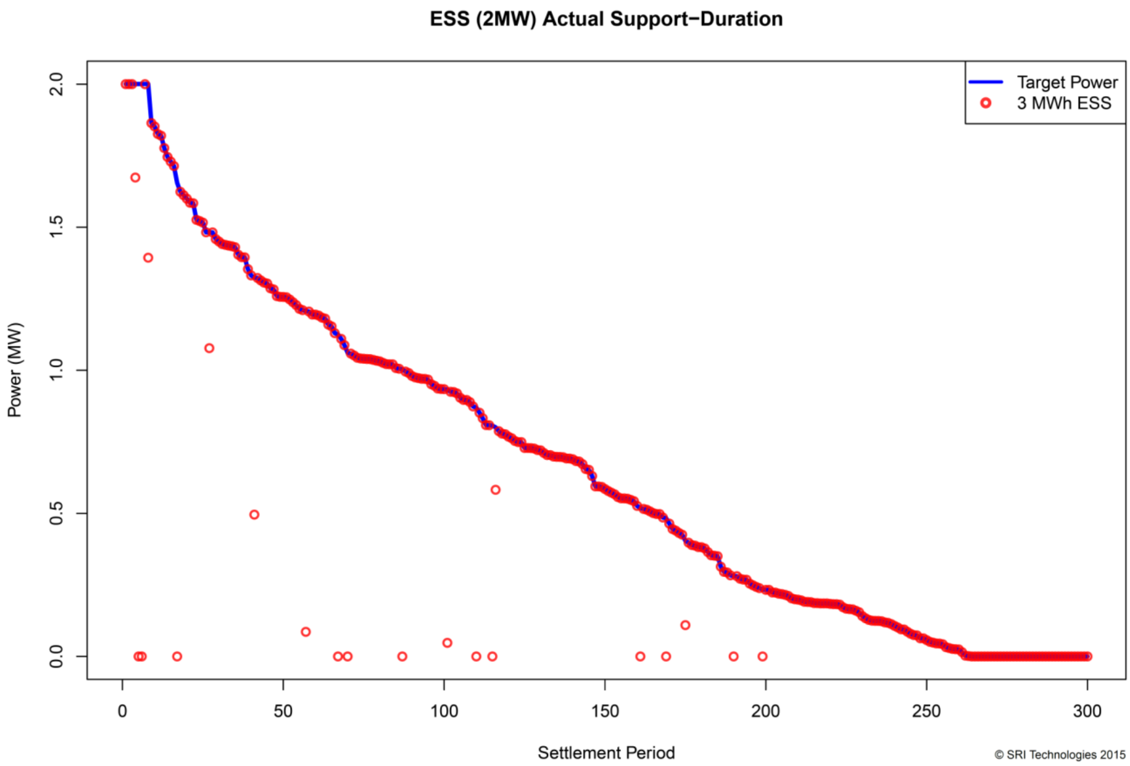

Figure 53: 2 MW Actual Support Duration

Once again there are diminishing returns from increasing the battery size and capacity. While smaller capacity batteries may have a relatively low performance, the financial impact of overloading the transformer for an additional number of half hour periods will often be less than the cost of increasing the battery capacity. This suggests the need for multiple options to manage network issues, such as, combining battery export with demand side response and load reduction via voltage control.

19.2 Utilities Insight

Utilities Insight carried out an assessment on the impact of the existing regulatory regime in July 2016, such as the Balancing and Settlement Code, on the potential development of the storage market. It considered different ownership models and found that the regulatory issues were more significant for the business models which entail DNO ownership and operation of the storage asset. This stems principally from the concern that DNO activity in storage projects could distort competition in generation and supply activities.

The report concluded that DNO led and owned development of smaller scale storage projects is therefore possible within the regulatory framework, but ensuring that such activity avoids distorting competition in generation and supply is a major factor which appears to block operation (though not ownership) of the assets by DNOs under the current framework. The regulatory framework does not prevent DNOs from procuring storage services from third parties, however third party owned storage may be less likely to be located at a site where it can provide a useful service to the DNO, than DNO owned storage. For example, where storage is expected to provide services to National Grid that may be required at any time of day or year then developers are more likely to select a location where an unrestricted connection can be provided rather than where network constraints exist.

In broad terms, the UK demand level runs at an average of 34.42GW. Storage capacity of circa 1.6GW has been predicted by 2020 with 500MW relating to DNO and Grid Services (Eunomia, 2016). The level of storage capacity in 2016, stands at just 24MW, comprised of projects that would be considered exclusively as lighthouse or proof of concept in nature. Additionally, connection applications for hundreds of MW of storage have been received, indicating an expected change in capacity of several orders of magnitude. If DNOs are limited in their ability to own and operate storage there is a risk that the storage will be installed at locations that are not beneficial to DNOs, reducing the overall benefits from storage and increasing the requirement for network investment.

At present the trade-offs between the potential for extra income from DNO services and the possibility of reduced income from National Grid services are hard for storage providers to evaluate. DNOs may be unwilling to commit to long term contracts with storage operators rather than looking for shorter term options that involve storage as part of a wider market for flexibility services.

The report suggested possible activities to promote DNO ownership and operation of storage in line with market requirements below:

- Clarifying/modifying treatment of electricity storage within the framework, including classification and requirements for licences to operate;

- Enabling DNO operation of electricity storage assets for balancing or constraint management purposes in a transparent and non-distortionary manner, delivering consistency with unbundling requirements;

- Considering the potential for Great Britain (GB) DNOs to trade in a non-speculative manner under a model similar to that under which National Grid fulfils its system operator role; and

- Including storage investments appropriately within price controls. This would need additional consideration. For instance, it could come in the form of an 'investment allowance' or equally it could provide the justification to reduce their overall investment after allowances have been set, thus funding the storage from these benefits.

DNOs have an opportunity to reduce the investment required to manage constraints via periodic and timely storage discharge and recharge.

In the short term DNOs should continue to deliver these practices up to a de minimis cap[38] of 2.5% of their revenue (approximately 15 projects per DNO licence). Beyond this de minimis cap it is considered by the Regulator to be distortive to the market. The following are therefore recommended positions for consideration:

1. Apply for a change to the distribution licence condition that increases the 2.5% de minimis cap (Section 6.1.1 in the Solar Storage - Impact Study report[39]) on supply and generation activities to a suitable number based on an assessment of market demand for storage (DNOs to provide market forecast evidence).
2. To remove this value associated with this de minimis cap so that this activity is unlimited in volume but critically is more heavily governed by a specific set of 'scenarios', a basis for which would be the usage cases in this project. The governance arrangements could include providing the DNO with the role of storage provider of last resort, such where a positive business case could be demonstrated that could not be provided by the market. This would allow the market to flow

[38] A regulated limit on non-core activity spending
[39] https://www.westernpower.co.uk/innovation/documents Solar Storage - Impact Study

unconstrained, with DNOs being able to participate fully, and also provide the protection against concerns of distortion of the supply and generation markets.

3. A combination of the measures in 1 and 2.

After this report (Solar Storage – Impact Study) was produced in December 2016, Ofgem and BEIS issued a joint call for evidence for “A smart, flexible energy system”, where the role of storage was considered. Ofgem have subsequently indicated that they remain in favour of DNO’s only owning storage as a provider of last resort and, even then, on a short-term basis with periodic reviews to determine whether the ownership could be transferred to a third party.

20. Battery Roadmap

20.1 Market size and competition

In 2018, in addition to energy storage already installed, there are approximately 9GW of batteries at various stages of development. Half of these, approximately 4.5GW, have secured planning consent and presumably must have grid connection offers. Whilst batteries had been seen as a panacea for grid services, there are a large capacity of established technologies already in the space, competing for all the same revenue streams and able to out-compete batteries currently e.g. small-scale gas reciprocating engines, DSR etc. In addition, a lot of the assets currently connected at transmission level have an obligation to provide grid services such as reactive power and frequency response at a set price, so even if technologically they are less suitable the service is still obtained more cheaply from them than from new distributed assets.

Previously, battery investment cases have been predicated on stacking all possible revenue streams including FFR, Capacity Market, Triads, Arbitrage (wholesale and balancing markets) etc. However, it has become apparent that not all items can be stacked simultaneously e.g. a battery cannot be chasing Triads whilst under an FFR contract. (Note, Triad income is less likely to be a factored revenue stream now that National Grid has split Triad Demand and Triad Generation whereby generation no longer receives the reciprocal of the demand charge, but a payment predicated on the avoided GSP (grid supply point) cost. In fact, many DNOs are due to go negative (capped at zero) in their Triad Generation payments).

It is well known that the FFR market size is around 600MW-800MW; a fraction of the battery development pipeline and an even smaller fraction of the total GWs that could perform FFR. In other words, the market is all but saturated.

In fact, National Grid are not procuring daytime (0700-2300) FFR capacity until at least April 2019 and then only 100MW per month rising to 200MW by the end of 2019.

This lack of available long-term frequency response contracts is cited as one of the contributory factors in the failure of Camborne Energy Storage, who co-located a Tesla Powerpack battery system at another PV plant in Somerset in 2016. While the company had established a pipeline of projects, they were not able to secure sufficient investment.

The arbitrage and balancing mechanism markets are estimated to be approximately 8-10GW, with some assets already competing in this space. It is expected that more energy storage assets will move across to this as they are forced out of FFR. The price fluctuation and therefore the market size is expected to grow as the renewables: fossil-fuel ratio continues to increase, assuming parity projects start coming online. However, given the lack of FFR market, all the batteries and other technologies will be competing, and cannibalising the revenue from this larger market.

Additionally, there are a number of companies trying to develop the market for storage for domestic PV owners to maximise self-consumption of PV generation. Depending on the degree to which these batteries are open to control by aggregators, this could offer further competition in the market.

The saturation of National Grid services creates an opportunity for the DNOs/ DSOs to be able to procure services at a significantly lower price than was originally predicted in 2016, as a National Grid contract is no longer assured. This is likely to remain the case until the price of batteries declines to the level where they can offer grid services at a level where other technologies are unable to compete. However, the declining revenues are expected to influence the deployment of energy storage, with less developments going ahead. The energy storage that is deployed is likely to be on unconstrained grid connections, by definition in areas that do not require DSO grid services. This creates a compromise: DSOs could offer cheaper prices but the batteries may not be located to usefully alleviate constraints: higher prices would increase the geographical diversity of storage.

This could mean that DNOs procure services cheaper than previously thought because the competition from national services is largely saturated, certainly until the cost of batteries declines sufficiently to the point that other technologies can no longer compete. Predictions show that it is highly unlikely before the mid-2020s. One potential option is that batteries should be subsidised, or given favourable charging regimes, in order to compete sooner, if there is a desire for cleaner flexibility.

If batteries are more beneficial to the DNO, compared to other technologies, and assuming the DNOs are in competition with national services, the DNOs need to offer something more favourable. One key factor to getting battery projects off the ground is funding and bankability; being able to leverage funding off the back of longer term, more fixed revenue streams such as the CM. EFR offered a 4-year contract but will not be auctioned again. FFR – National Grid are largely only procuring monthly contracts and with the intention (or at least trial of) moving towards weekly auctions. Currently there's not enough equity interest and only up to around 50% of costs can be met by senior debt. There is a lack of interest in funding as there is a lack of long term, contracted revenue. This is the key to unlocking the battery potential. Although there is a large pipeline of shovel-ready projects, it seems likely that a large number of these will be shelved due to declining FFR prices and poor returns from contracted revenue. Even if some of these batteries are sited at solar parks, there doesn't appear to significant revenue gains available from this co-location: FFR income heavily outperforms solar peak lopping, as does arbitrage to a lesser degree. If they are built out, then DNOs would be competing with a lower revenue base case for the batteries time, but if they aren't built at all then this could force the DNO to more expensive measures. A possible third way is DNOs offering longer term contracts, increasing the bankability of the revenue streams and getting investors on board.

While DNOs are aiming to improve the long-term signals to investors as to which locations are likely to require flexibility services in the future, without actual contracts this may not be enough to trigger battery investment. There also seems to be a gap in storage developer and DNO expectations, with DNOs assuming that their services would not, in most cases, provide the base case investment but rather offer additional income to service providers that already had a business case for investment. DNOs will not be able to justify flexibility service payments that exceed the cost of traditional reinforcement and therefore it is possible that if DNOs are the sole source of income for battery development than the business case may not be viable.

20.2 Cost

Li-ion costs have come down approximately 10% in 2017. Cost reduction is already slowing, showing that Li-ion has already leveraged the big price falls as it moves into being a mature technology. Prices are forecast to fall a further 8% per annum on average over next five years.

Full system costs are expected to fall by around 4% per annum over the next five years. This overall cost reduction is slower than the rate of cell cost reduction due to the slower decline expected in the relatively mature balance of system costs.

20.3 FFR/EFR

Between August 2017 and February 2018, the accepted dynamic prices tendered for FFR declined by 66%, from approximately £15.MW/h to just over £5/MW/h. For the remainder of 2018, the average price was approximately £6-7. Forecasted FFR rates have generally seen a 20%+ downward adjustment reflecting the increased supply of energy storage technologies.

FFR rates could be 60% lower by 2030 due to market saturation and related competition. Rates could stabilise at a similar time to costs stabilising and the opportunity cost of arbitrage is on a par with ancillary service income. Other than the cost of cycling, assuming a difference between operating for frequency and operating for arbitrage, battery owners could therefore be indifferent as to revenue stream targeted, thus supporting business cases that incorporate the necessary flexibility to flip between markets.

20.4 Frequency

The frequency market could be considered to be largely saturated. National Grid is currently mainly procuring overnight frequency response, where the lower opportunity cost depresses overnight rates by around 15%-20% compared to blended 24/7 rates. National Grid is not forecasting any daytime frequency requirement until 70MW in April 2019. This slowly increases, before stabilising at 200MW by the end of 2019. 600MW has already been procured forward until then.

In addition, National Grid is predominantly only procuring monthly FFR contracts. There is also no indication of any plans to repeat the EFR auction; currently the longest-term contract created in this space. This creates a total absence of long-term predictable revenue for bringing new assets onto the network. This is despite the frequency response requirement being expected to rise throughout the 2020s, driven by coal plant closures, reducing system inertia.

With declining costs, batteries may force pumped hydro out of the market by 2025 before levelling out with a circa 90% share of a 1.5GW FFR market at around 1.4GW in the FFR space. With approximately 4.5GW of batteries having already obtained planning consent, batteries securing FFR contracts will be a relatively small proportion of the current battery pipeline, clearly pointing towards the need for batteries to adopt a business case flexible enough to incorporate distinct revenue streams.

National Grid has been discussing carrying out closer to real-time procurement of frequency services. This has the potential to vastly increase the competition (and liquidity) for frequency contracts as it lowers the barriers to entry currently faced by the likes of solar and wind which struggle to perfectly forecast generation beyond the immediate horizon. Again, this gives existing assets access to an extra revenue stream but won’t be helpful in gaining investment for new assets.

Weekly trial auctions will commence in December 2018 with the hope that these become daily going forward. In conclusion, front of meter batteries can no longer assume a purely National Grid-contracted frequency response business case and increasingly need to turn to arbitrage. This may change as DSOs increase their portfolio of procured distribution network services.

20.5 Capacity Market

The two last auctions, T-1 and T-4, held at the start of 2018 saw a huge surplus of capacity, including interconnectors and demand side response (DSR), driving clearing prices down.

Current trend in battery depth, 1 hour, saw its de-rating factor drop from 96% to 36%, whilst DSR (including behind the meter batteries) was subject to a factor of only 86%. DSR was one of main winners in the capacity market partly due to its lower capex and partly due to its favourable de-rating factor, however, if Scottish Power’s proposal[40] for a change to the Capacity Market Rules is successful, whereby DSR assets are assigned minimum durations and corresponding de-rating factors, DSR will be much more closely aligned to front-of-meter batteries in terms of capacity market price it could tender.

[40] https://www.ofgem.gov.uk/publications-and-updates/scottish-power-capacity-market-rules-cp353

Going forward, as coal plants are taken offline, new build capacity, including CCGT, reciprocal engines (both gas and diesel) and batteries, will be increasingly required, however with the duration limited de-rating factors, batteries may need to start considering greater depths to help offset the low factors. As the only long term contracted revenue, assuming new builds continue to able to secure 15-year contract, the CM is significantly more valuable to a business case than the face-value revenue it provides.

20.6 Requirements

For the rollout of batteries to occur as originally predicted, costs need to decline significantly more rapidly. This may occur from an entirely different technology, such as vanadium redox-flow systems, but it must occur at an accelerated rate to avoid a delay in energy storage projects. Developers also need to investigate deeper batteries, as these boost arbitrage revenues and, more importantly, give increases in CM revenue which as noted is important to investors. There are opportunities for new technologies to move into this space, if they offer better economies of scale or alternative properties than lithium, and there are opportunities for hybrid systems to combine the best features of two technologies. A potential future area of research would be the interoperability of different storage solutions and the effectiveness of hybridised systems. It is expected that the biggest issue facing the hybrids is the control system complexity required. The issues faced in the control system for this battery suggest that combining two systems and ensuring that each are used to their greatest effectiveness could be difficult.

As well as targeting wholesale markets, if PeakGen's modification proposal, P355[41] is successful, and assuming not all batteries can be aggregated into Balancing Mechanism Units (BMUs), almost all front-of-meter batteries (>1MW) will be able to access BM revenues. This modification proposes to introduce a new 'BM Lite' classification to allow smaller generators to offer energy to the System Operator for energy balancing, in competition with the larger BMUs already in the market.

The potential size of the balancing market could be around five times that of the frequency response market, especially if relatively volatile renewables such as solar and wind take off and dominate new build penetration.

20.7 Business case summary

The future FFR market is small compared to the capacity of batteries in the pipeline, let alone the wider flex technology market. For the business cases to succeed, developers should consider ignoring frequency as the base case on which to secure investment, as only a small handful of projects with the very lowest grid costs will be able to compete and lock themselves out of other contracts. The key focus of new build projects should be arbitrage.

Even with this shift of focus, for batteries to succeed in a more liquid market, the processes and operation of batteries needs to be finely tuned in order to efficiently move between markets, and both inter and intra-day. However, a critical consequence of almost fully

[41] https://www.elexon.co.uk/wp-content/uploads/2017/06/P355-BM-Lite-Proposal-Form-v1.0.pdf Proposal for introducing a BM Lite balancing mechanism

merchant revenue is the lower appetite of banks to provide project finance. A project is unlikely to raise more than 50% of the capex.

Battery developers should be happy to forego potential merchant/arbitrage upside for surety of contracted revenues and improved gearing. However, currently only capacity type markets offer any contracted revenue and capacity requirements are ultimately finite and fall well short of supply.

Perhaps going forward, traders/suppliers/generators etc. may be willing to procure capacity availability in advance (and related depth), eliminating a battery's upside whilst promoting certainty.

Given the current saturation in frequency and capacity markets, in front of the meter batteries can only turn to DSOs whom in many cases are some years away from launching any product procurement, although WPD is already rolling out flexible power. The battery business case that will come to the fore the quickest will be the behind the meter model.

If DNO charges are not removed for batteries, the DNOs will be double counting this revenue from batteries and end users. However, batteries are still using the grid and should arguably pay something for it. Perhaps, however, such DNO charges could be removed for assets only while they are providing DNO/DSO services. Whilst DNOs are not permitted to own generation assets, they may be permitted to own their batteries if not anti-competitive. It is considered more likely that offering longer term contracts to projects, that are on the verge of being built, is likely to be more economically advantageous however.

For this system to work to the DNOs advantage, if a DNO recognises it needs services specifically required by the DNO (and not National Grid) e.g. reactive power, power quality, curtailment, ANM etc., then the DNOs must assess and advise developers where these locations are and what grid capacity is available to support the local grid. WPD has consulted with stakeholders on the issue of signposting[42] and is developing a Network Flexibility Map[43] to complement the existing capacity maps for load and generation.

National Grid

The revenue streams from National Grid services are forecast to decline, due to simple over supply of National Grid's requirement. This is unlikely to change due to simplification as batteries are currently perfectly able to access current services. If anything, simplifying services could make it easier for other technologies, and less experienced battery operators to access the services, adding to the saturation.

[42] https://www.westernpower.co.uk/signposting

[43] https://www.westernpower.co.uk/network-flexibility-map

However, if there is a well-defined 'service duration' hierarchy, this may allow battery developers to fine-tune their strategy and optimise their batteries more effectively to target certain products on offer.

The new, more frequent, FFR auctions are likely to benefit National Grid as they will have a bountiful supply of assets to draw on for its services, whilst ensuring continuous lowest cost, as opposed to locking into [up to] two-year contracts.

For existing batteries, it will certainly give operators more optionality compared to monthly auctions, however, if coupled with shorter contracts and greater uncertainty of securing contracts, more frequent auctions can only be negative for getting new batteries funded.

For the current cost of batteries, it would appear that they have to be fully utilised, whether this is FFR 24/7, cycling around once per day, or arbitrage cycling approximately twice a day (noting the impact of degradation versus revenue capture).

Currently, any strategy that only employs the battery for part of the day e.g. test cases, the revenue per unit of test case will have to be sufficiently high to offset the forgone opportunity of operating for more hours (noting the potential benefit of less degradation).

It will certainly be interesting to see how many MWs out of the 4GW+ with planning consent, and up to the 9GW in pipeline, will be deployed in the next two years and during the time it takes for DNOs to transition to DSOs.

However, given that a relatively small proportion of the battery pipeline will be built out in the next 12-18 months, funding and development will be looking to its next preferred route of secure revenue, namely that from DSOs.

In the meantime, many projects and MWs will not be built and not be sufficiently progressed/progressing invoking the DNOs' new regime of 'slow process[44]', and developers will find it increasingly difficult to come up with new and valid reasons such that grid capacity is likely to be lost and handed back to DNOs. This, in itself, may support the current infrastructure with less need to provide local network services simply by having less storage connected on the network.

DNOs are now carrying out the analysis to support early signposting of their future requirements. It may be that they are able to pay more per unit of service across fewer assets i.e. targeted battery requirement such as those under consideration in SSE's CMZ tender[45] and UKPN's Power Potential.

How can DSOs help?

Although the market is still developing, the industry is aware of the services that are likely to be required by DSOs. Details such as location are key, as well as which service, number of MWs or MWh or speed of response, how many hours of the day its required etc. The

[44] If projects are not submitted to planning by certain dates / hit other milestones, then their offer and capacity is revoked. While this doesn't often happen 'on the day', if it is highlighted to DNOs then the offer enters the slow moving process.

[45] http://news.ssen.co.uk/news/all-articles/2016/12/ssen-opens-constraint-managed-zone/

number of years the service will be required is also critical, as this will need to be built into any investment model. The Open Networks project has considered how DNOs can best signpost this information to potential service providers and this type of information is now available for WPD regions via their flexibility map.

Realistically, if the investment case can't be built on services already procured e.g. FFR/CM by National Grid, then developers will be reliant on DSO's to provide the certainty that will allow debt to be raised to bring new projects online. Currently the only long-term contract is the Capacity Market auction, which despite huge derating factors and rock bottom prices, continues to see every new project bid in to get a small part of their revenue stack locked down for the longer term. This gives an indication of the appetite there is for fixed revenue certainty.

Battery cell costs

Whilst battery (cell) costs are due to fall, there is little scope for balance of system costs to do the same, such as civils, inverters etc. i.e. technology and works that are some way along the learning curve.

Whether cell costs (lithium) do decline in practice, will reflect factors similar to the pricing of oil, or any other finite resource. Those countries/ parties that control the supply of lithium can, to a certain extent, manipulate the price via simple supply management versus a globally growing demand. Furthermore, lithium demand is not just limited to grid applications whether they are international, national or local grid use or behind-the-meter, applications, but rather the demand from vehicles is significantly higher. This is borne out by the extended waiting times for batteries that several in the industry are now experiencing. Several cell / battery manufacturers are often 'vehicle-first', meaning their stationary offerings are using spare capacity rather than dedicated production lines to fulfil orders.

Behind-the-meter applications, whether half hourly metered or not, have the potential to make significant savings on DUoS and TNUoS charges, however DNO and National Grid expenditure needs to be recovered somehow. It is understood DUoS energy costs are flattening i.e. the difference between Red and Amber, making less of a case for this type of arbitrage. However, grid charge arbitrage is not sufficient to make a battery viable. The battery needs to be 'working' 24/7 in the other markets and outside the peak price times.

Domestic and small-scale rooftop could benefit from batteries by allowing more self-consumption however there are a couple things to consider here. Firstly, as with solar park peak shifting, for this to be most beneficial to homes, the battery should be deeper than one hour, for example, five hours comprising 2-3 hours in the evening and 2-3 hours in the morning, as this more closely matches standard network demand and the amount of spilled solar power. Less than this capacity is unlikely to reduce dependence on the grid during the peak times. Another, more retrospective, flaw is given that solar PV systems up to 30kW are deemed to export 50% of its generation, and (if registered before April 2019) are paid the corresponding Export FiT, if more solar energy is consumed onsite, even up to 100%, the Export FiT and LCF budget is in a sense funding the domestic battery i.e. the Export FiT is paying for energy not exported. This clearly benefits the household and may help the deployment of domestic batteries in the short-term, but it would be expected that Ofgem would require metered export if this became a significant issue.

New housing developments on a shared network could benefit from a centralised battery for the estate, helping the estate network minimise its reliance on the main grid and distribution networks almost being 'off grid'. Again, this is likely to require a deep flow battery. A collaboration with a large number of companies in the energy and storage space are collaborating together on a project such as this, known as Flatline Energy.[46]

The magic ratio

For the solar-storage hybrid to stand any chance of having a successful business case, further investigation is required to derive the magic ratio of the three key elements:

- Solar capacity
- Energy storage power and capacity
- Grid connection (and inverter) capacity

In the traditional solar PV model, a solar park capacity is typically oversized compared to the grid connection capacity by an AC:DC ratio of around 1.2-1.4. The rationale for this is to be able to export more of the non-summer, non-midday generation. In other words, whilst the very peak of solar generation may be capped at the inverter/ grid capacity i.e. for 1-2 hours during the peak summer months, this loss of output is more than compensated for by the increased production and export at other times of the year that is as a result of the higher DC:AC ratio.

Taking the peak shift case as the direct correlation to the typical oversized solar park, with storage combined there may be the opportunity to greatly oversize the solar park's DC capacity compared to the grid (export) AC capacity whilst still capturing the much larger, and otherwise capped/ lopped, peak. For example, (and the ideal ratio has in no way been tested) the solar park may have a 5MW DC capacity, a grid export capacity of 1MW AC and a flow 1MW/ 5MWh (5 hour) flow battery. The flow battery would capture several hours of solar generation during the middle of the day for export at other times. The benefit is potentially two-fold: the stored energy could be sold at times of higher prices, whilst the capital cost of a 1MW grid connection is significantly lower than the typical 4MVA associated with a 5MW DC solar park.

However, this doesn't work on a couple of other fronts. Firstly, the capital cost of a flow battery is 5-10 times more expensive than a lithium battery, albeit it doesn't degrade as such, and secondly, whilst the peak captured in summer might be 5 hours, fully utilising the depth of the flow battery, the lopped peak might only be 1 hour, if that, during winter, underutilising the flow battery for peak shifting purposes.

If looking at parity solar parks, i.e. new build solar and new storage not retrofitted, it may be the case that instead of the solar park exporting to the grid in winter it charges the battery

[46] https://www.flatlineenergy.co.uk/

for export at higher price times. This may help solar reach parity sooner but at the expense of a battery that needs to pay for itself.

Alternatively, retrofitting one-hour batteries to subsidised solar parks with the current typical oversizing, the solar park could capture an element of summer peak generation that would otherwise be curtailed whilst receiving energy revenues and subsidies, net of any payment to the battery. Further, if the battery was only called on by the solar park in this way for a month or two, the battery would then be free to perform in other markets for a greater period of the year.

In a recent case, Ofgem stated, by way of blog posting, retrofitted batteries do not infringe or curtail a solar park's subsidy.

Whether using one hour lithium batteries or five hour flow batteries, for the battery to operate effectively co-located with solar, it will be critical to get the magic ratio right.

21. Limitations and future learnings

The brief for this project was relatively large, with 9 use cases expected to be investigated. Although PV export limiting and variable PV export limiting were combined, as they were tested in an identical way, this was still a considerable range of operating modes to be studied.

It was partially for this reason that the use case investigating multiple storage system control (use case 9) was omitted from the project. As there was only one battery involved in the project it would not have been possible to do real world testing, but the original plan was to have a virtualised RESolve battery operating alongside the real one and investigate the interaction in this way. Given the existing operation of Active Network Management schemes does not operate by harnessing individual control systems together for the various assets , then it is a fair assumption that multiple isolated storage systems would also operate by way of some centralised control by either the TSO or DNO. This has been borne out by the work progressing on the Power Potential project, which involves one control system in charge of the whole region, sending specific setpoints to specific assets. In this way, the battery control systems are simpler and just have to receive and process a setpoint, rather than calculating the effect on the grid locally.

The battery size was selected so as to be large enough to be able to affect the network, while minimising the size and therefore cost and drive the most value for bill payers. While most of the active power use cases were able to drive valuable learning, the reactive power capabilities of the inverters were too low to show any results. Similarly the low demand grid voltage support did not show any obvious results, but the battery was able to demonstrate reliable performance in all the other use cases.

One large aspect that could be improved moving forward is the retention of data. Despite checks that the server was set up to retain all information, the SCADA system was actually deleting information after one year. As this issue wasn't noticed for quite a long period, several months of data from the beginning of the project have been lost. However the rest of the time has a complete database.

Although PV export limiting was investigated, solar firming (the act of continuing export when generation dips) was only briefly tested as it wasn't detailed as part of the core test program. The export limiting and solar firming mode could be the subject of a more advanced and detailed study: working towards an algorithm that meets the requirements of G100, effectiveness of peak lopping on increasing PPA prices (predictable export windows should then improve the negotiated PPA sell price), and optimum size of the energy storage (from this report it is expected to be 80-100% but a more accurate figure would be valuable).

It is possible for the project to deliver additional learning from third party analysis. The dataset that includes all operational data of the energy storage system is recorded and available, meaning further, more focused analysis can be performed by anyone with a more specific interest in a particular aspect of the project.

22. Conclusions Summary

Arbitrage was originally considered to be a simple base-case, and it was expected that the battery could take advantage of inter-day price variations to generate revenue when no other opportunities were available. It transpired that the PPAs and import agreements entered into by commercial solar parks effectively removed any opportunity for the battery to access those variations. The price stability and certainty targeted by solar parks are the opposite of the aims of the battery, which favours as much exposure to price fluctuations as possible.

To create a more realistic scenario, a virtual PPA for 2017 was created, based on the system price and grid charges. This reflected the likely arbitrage opportunities the battery would have had as a standalone installation. The theoretical revenue generated was approximately £8,000, with future years likely to generate more, due to expected increases in price volatility. This in stark contrast to the FFR revenue, which has declined from £64,000 at 2016 prices to £24,000 at 2018 prices. The 2016 price hasn't been used to compare any other use case as it is irrelevant to the current business case, and is merely used as a benchmark to demonstrate how quickly the energy storage landscape changes and the downward trend of the revenue from grid services.

With FFR contracts becoming extremely difficult to win, even with providers continuing to cut prices, arbitrage has rapidly become the most likely business case for energy storage. The 'sweet spot' of battery depth is different for arbitrage than FFR, with a deeper battery being able to access a higher proportion of the profitable half hours.

In 2017, arbitrage was being ignored as so low value compared to FFR that it was practically dismissed from business plans, but going forward it is expected that arbitrage will be the new base case that other use cases are compared to. The lack of reliance on third party contracts makes this an attractive and reliable revenue stream.

Other revenue streams are able to provide additionality, for instance, the network peak demand limiting (use case 2) would be viable as long as the hours that were 'booked' for this need generated either a higher or more bankable revenue than standard arbitrage. The difficulty in predicting whether this is likely to occur is that, unlike FFR with its flat rate, arbitrage opportunities can be grouped around certain times of day. If the network support is needed at a certain time that prevents the price arbitrage being taken advantage of, then the cost of the service would be higher, and this would be passed on to the DNO/DSO. An example of this would be if the network demand peaks between 3-4pm, due to unusual load profiles, while the high export price doesn't occur until 6-7pm. The energy that would have been sold at a higher price is instead effectively being exported at a discount, which the DSO would need to 'top up'.

The local demand profile matching (use case 3) is so location-dependent that calculating the returns is almost meaningless. The main outcome of this use case test is that the energy storage system is able to respond appropriately to the soft inter-trip signal and therefore could fulfil this need, if required. The calculation of value to a company would be based on the cost of interruption to their operations caused by a blackout, balanced against the costs of actually triggering the reinforcement process. If the value between these two scenarios is higher than the base case, then there is potential for a deal between the battery operator and the factory/company.

Low demand grid voltage support (use case 4) is expected to be the least useful and used use case. While the drop in voltage caused by increased load is well-known, the requirement for multi-hour import at a level high enough to impact the voltage means a large, deep battery would be needed. In addition, the dedication of the full capacity for so many hours would come with a large opportunity cost. It is possible that a storage device dedicated to managing constraints on a single high voltage line, e.g. with double requirements of voltage control overnight and network peak management during the day, could be viable compared to reinforcement at the 132kV level, but this will be the exception rather than the rule.

The more likely type of active power-based voltage control would be ramp rate control (use case 8). By smoothing out the sudden changes, the battery can improve power quality and reduce wear and tear on the DNO's tap changers. The effectiveness of the test battery was limited by its size, while an 80-100% power sizing to the solar inverters would be more suitable. The battery was only able to influence the smaller of the two solar parks, giving a total effectiveness of 5.6% at peak output across the two parks. This use case may suit a different technology to provide its effectiveness, with supercapacitors offering low amounts of storage but at high capacity. A 10-minute supercapacitor connected to the DC side of a solar park, with a control system linked into the solar inverters, would be the most suitable system for this use case. It would be a waste of a more versatile asset to use a lithium battery for this purpose.

The real voltage control is likely to be provided by reactive power systems. In this way, the voltage is able to be boosted or decreased on demand, with a large reactive power response available from a minimal de-rating of active power. Because of this, Use Case 5 lends itself most strongly to multi-asset systems: instead of one installation de-rating heavily to provide enough reactive power, several assets can all de-rate slightly, and working together, will have a sufficient influence on the reactive power flows of the network. This would be especially valuable at night, when power factor is often poor on the network, as the set point of reactive power could aid the power factor and control voltage rise at the same time. Reactive power voltage control is the use case that most strongly demonstrated an upside from trying to get an asset to operate in two markets at once.

The evolving energy storage market is making it more complex to create a business case with reasonable returns. Regulatory change is coming, but not all of this will be positive for batteries, with other user's needs being prioritised in some cases. Moving forwards, a more bespoke system looks increasingly likely, with detailed grid requirements potentially replacing the 24/7 capacity access currently granted to most connections. This has begun with ANM connections, but if a more detailed profile can be provided to DNOs then further grid costs could potentially be avoided. This would also allow storage to begin to be embedded in constrained areas of the network, which is exactly where the DSO contracts are most likely to be offered. Currently energy storage is only installed on 'clean' connections, at the strongest connection points with the most capacity. This works perfectly for National Grid contracts, but won't aid the DSO transition.

Combining solar and storage initially appears to make perfect sense, and gains a significant amount of interest across the industry. However the reality is less clear-cut. The battery can benefit solar parks by increasing solar revenue, but this currently is at the expense of more lucrative revenue streams. With FFR contracts reducing in both price and availability, the base case is becoming arbitrage. If solar-related revenues can't compete with arbitrage, then financially a behind-the-meter solar battery is worse than a standalone grid-connected battery. On the other hand, it is possible that some savings could be realised by sharing grid assets between a solar park and an energy storage installation, as long as each asset had its own meters and connection agreement.

A big advantage of combining solar and storage is it allows clustering of generation assets, which could ease the process of gaining land owners consent, reduce grid costs or shorten connection times, and even potentially aid planning permission. Solar parks are connected in rural areas, which are often the areas that have slower increases in growth and could benefit from network support services, as well as being weaker with less resistance to voltage fluctuations, meaning reactive power support could be useful.

It should be highlighted that these conclusions represent a static snapshot of the state of the market, and with continuous regulatory change and technological advancement some assumptions will quickly become out of date. For instance, a higher penetration of solar power would force the wholesale price lower during the middle of the day, potentially even negative. At this point the arbitrage between midday and evening peak price could be significant enough that a behind the meter solar battery would make the most financial sense. Realistically there are benefits to allowing behind the meter storage and there should be an overall goal to remove any regulatory barriers that could prevent this, as even if in 2019 this isn't the most efficient use of storage, it could become so in the future.

DNOs have now all committed to exploring flexibility options, and help ensure that RIIO-ED2 (the next round of price controls for DNOs) do not incentivise reinforcement if flexibility services could provide the requirements at a cheaper price[47]. This commitment should put flexibility on an even priority with traditional reinforcement and accelerate its progress towards business as usual. With WPD working to stimulate a flexibility market with Flexible Power, energy storage is already leaving the pure research area. The data from this trial is likely to prove extremely valuable in validating the viability of storage for the flexibility services DNOs require.

Batteries remain the ‘gold standard’ of grid assets, able to perform practically any role from fast response to increasing load, reactive power support or additional generation. There is significant competition remaining from older technologies, such as diesel and gas peaking plants, which can fulfil a single use case more cheaply but don’t have the dynamic flexibility that batteries offer. Batteries that operate for over 10 years should still be able to be performing useful roles for grid balancing etc at the end of their life, while other technologies may find themselves obsolete as their single revenue stream diminishes. However, there is not obvious clarity on what markets will exist for the next 10 years which makes it harder for developers to drive forward projects, as investors need a clear view of the revenues needed for meeting target IRR.

The overall market needs to help stimulate energy storage development, with longer contracts available, in order to get a critical mass of storage available to create a flexibility market. The installation of renewables was incentivised successfully as standard market forces were unable to create a viable investment case for these emerging technologies. But flexibility is being left and with revenues declining so precipitously energy storage isn’t competing, because the level of flexibility isn’t being rewarded.

This project has also proven the locational flexibility of storage. The battery was a containerised solution that was able to be delivered and removed by a contracted lift and moved on the back of an articulated lorry. This is an example of the potential future of storage, in areas with more rapidly changing requirements. The battery could be connected by a temporary connection for a few years, then removed and reinstalled at another location as network demands change. While this would create more expense than being able to leave the system at a single location, it demonstrates the ultimate flexibility of these assets. More about the removal of the battery can be found in the Battery Disposal Report noted in Section Further Information.

The final conclusion to draw is that storage is ready for the challenge of providing grid services. The fact that the test battery was able to fulfil nearly all of the requirements across several use cases, albeit with some initial teething problems, is a strong showing for the technology. It has been assumed that a future control system would have the minor software issues repaired, as test systems often have more loosely defined requirements than commercial systems and the majority of scenarios functioned perfectly, despite this being the first battery to implement several of these use cases.

[47] https://utilityweek.co.uk/all-six-dnos-sign-up-to-flexibility-pledge/

23. Further Information

The following documents relating to Solar Storage are available on the WPD innovation website;

- Solar Storage – Battery Disposal Report
- Solar Storage – Learning and Dissemination Webinar slides

Appendix A. Virtual Power Purchase Agreement

Please see the associated Excel spreadsheet

Appendix B. Base case calculation

Please see the associated Excel spreadsheet

Appendix C. RESolve Control Modes

The RESolve system operated using schedules that are programmed in advance by the user. When the selected time is reached, the schedule activates and operates the battery according to the control mode that has been chosen. The table below details the uses of the different control modes.

RESolve schedule summary table	
Control Mode	**Use**
Auto Export	Instructs the battery to export power for a set amount of time or for a set amount of kWh, at a specified kW.
Auto Import	Instructs the battery to import power for a set amount of time or for a set amount of kWh, at a specified kW.
Auto SoC	Charges or discharges the battery as appropriate to achieve the target SoC. The rate of charge or discharge is specified by the user in kW.
Peak Shaving	Charges or discharges the battery depending on whether solar generation is above or below pre-set kW values. This uses a PID control based on the readings from the solar PQM. During the tests it was nearly exclusively used in peak lopping (charging the battery) mode, rather than solar firming (discharging the battery) mode.
Ramp Rate Control	Using the input from the solar PQM, this mode monitors the rate of change of generation and instructs the battery to charge or discharge in opposition to that change. The threshold rate of change is set by the user, in kW/s.
Safe Mode	This is effectively a 'dummy' schedule, which can be set to make sure the battery doesn't do anything for a period of time.
Shutdown	This sends the shutdown signal to the BYD systems, turning off the battery and shutting down the PCS (power conversion system).
StartUp	This sends the start-up signal to the BYD system, reconnecting the PCSs and calculating the SoC. It takes a couple of minutes for the system to be ready.
Voltage Control	The only reactive power control mode. This requires a target voltage (kV) and max/min reactive power settings (kVAr). It also required calculating a control slope to say what percentage deviation from the target voltage would be sufficient to reach maximum reactive power response.
Combination Method	This allowed two or three control modes to be active at once. It prioritised reactive power over active power setpoints, and was set to sum the active power setpoints if more than one was active.

Appendix D. Commissioning Test Learning Summary

Commissioning Tests.

The battery was seen to perform all the required tests that included.

- Visual inspections
- Importing and exporting at different levels of power and power factor
- A set of round trip efficiency tests
- Tests for the control modes relating to the different use cases

The “COP5” meter and power quality meter were connected at the same point and shared CTs. There was a high degree of correlation between the meter advances, however the PQM gave results for energy delivered and received to 0.01 of a kWh whereas the COP5 meter display gave a resolution of 10kWh. Having proven that the readings were consistent, this enabled the PQM values to be used in testing in place of those from the COP5 meter which were specified in the test schedule. As the PQM values were displayed and recorded via the Resolve software, this enabled tests to be run overnight with confidence, rather than needing to inspect the Cop5 meter in real time.

Round trip efficiency tests

Round trip efficiency values are all comfortably higher than the battery specification requirements though there is some variability in the results where these were carried out more than once. Tests were repeated where the value was suspect (100%) or the test was interrupted, and differences between the two results suggest that a better approach would be to carry out the tests multiple times and take an average. This was not possible in the time allocated, however, it should be possible to calculate the efficiency values while testing the various use cases over the next year.

Cooling power used during the test would be expected to vary according to the temperature of the batteries at the start of the test, reflecting the operations prior to staring the efficiency test. Similarly the time of day and weather would also affect the cooling requirements. To a lesser extent, testing by human operation on site or remote control would also impact the values due to differences in lighting, heating and humidity control.

Use case tests

Ramp rate control was seen to operate on three occasions. Each time it showed a very clear and sensitive response with the battery output mirroring the variation in PV output.

The event triggered battery operation was seen to cut out if the original signal was removed and to cut back in if the trigger was reapplied. However where the signal was scheduled to run for a particular duration after triggering this would typically be shortened by approx. 30

seconds due to the time for the schedule to load, and the duration would be calculated from the first trigger and not be reset for subsequent triggers within the period of that duration. Additional triggers after the initial duration was complete would reset the timer. This is not expected to cause any problems, but would be useful knowledge for scheduling.

Two peak lopping modes were seen. In the first mode the combined output of the battery and PV system are compared to two thresholds. If the output is above the higher threshold, the battery discharges and the rate at which it discharges is proportional to the difference between the actual value and the upper threshold.

To prevent the battery from filling up from constant peak lopping, the battery is discharged if the level falls below the lower threshold. Again the speed of discharge relates to the distance below the threshold. This acts to fill in the troughs and to further smooth the output in addition to the peak lopping.

This is more complex than the simplified illustration of peak lopping below that implies a single limit that is not breached.

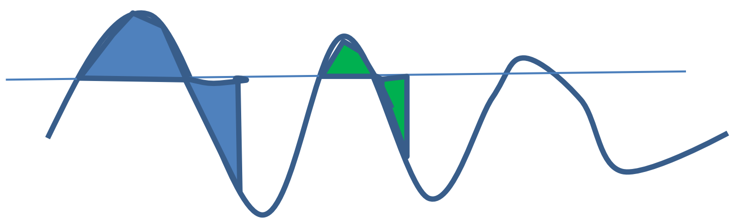

Most notably, the combined output will exceed the upper threshold which should not be considered as a hard stop. It is likely that the settings to achieve compliance under a particular limit can be determined by modelling and experience such that thresholds can be set appropriately.

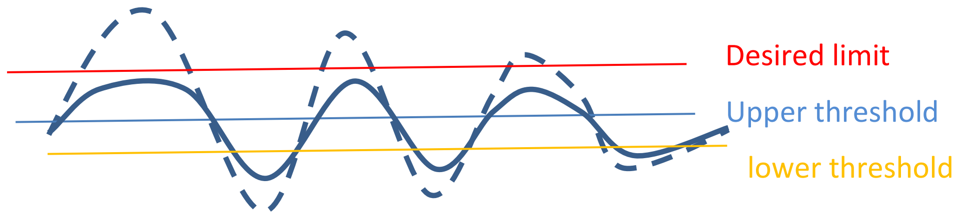

The other peak lopping mode triggers charging or discharging at full power once the upper or lower thresholds have been hit. This created a very "noisy" signal with rapid fluctuations between the limits. While this had the advantage that the desired limit could be set nearer the upper threshold, because the degree of overshoot would be reduced by switching to full power instantly, the likely impact on the battery and the network from the rapid changes between full charge and discharge would suggest this is not going to be a practical control mode on its own.

A combination approach using a proportional approach over one threshold with a breach of a higher threshold causing discharge at full capacity might be workable and should be explored during the project.

For the Voltage management tests the tap change operations at Millfield were logged so that their impact could be included when interpreting the results. The variation in voltage was seen to be considerable and highly dependent on the PV output which was very sensitive to cloud cover. These variations in background voltage were seen to be larger than the impact of the battery suggesting that this means of voltage control may only have limited effect in daylight hours. Again it is a little early to draw conclusions and the testing during the year will give a better indication.

Standby / Auxiliary power consumption

Once charged to 50% the battery SOC was seen to fall to 48% over three hours. The energy used to recharge the battery to 50% was used as a proxy for the energy losses due to charge dissipation within the battery, auxiliary loads within the battery and losses on the incoming transformer, switchgear and cables. The energy consumption with the battery left on SOC management was not measured which would have involved the battery topping itself up to 50% as required during the 3 hour period. This would be a useful comparison to inform battery scheduling to minimise energy losses.

BYD battery management system overrides

During the tests the BYD battery management system (BMS) superseded the Resolve control software.

One instance appeared to be due to the BYD estimates of SOC being very different between the four strings. The system interrupted the planned schedule to charge the battery at full power until it was charged to 100%. The SOC for the strings that were lightly loaded can be seen to recalibrate at full charge, suddenly jumping from approx. 80% to 100%. After this the battery discharged at full power before resuming the Resolve control mode.

SOC estimation has been an issue with other WPD battery projects. Both Sola Bristol and Falcon batteries had scheduled events to recalibrate SOC by charging and discharging. The learning here is that to avoid unwanted interruptions to the schedule we must be able to understand the rate at which SOC estimation drifts off and therefore when a recalibration event would be expected. Ultimately a predictive model that could determine the optimum frequency of recalibration events would be useful as it would ensure that calibration was not carried out too frequently but would prevent planned schedules being interrupted.

Another instance related to the battery temperature. As the batteries were warmer than would be expected, the BYD battery management system reduced the rate at which batteries could be charged or discharged. On investigation, the air conditioning units appeared to be triggering for only short periods of time. The setpoint for the air conditioning was reduced from 20 degrees C to 18 degrees C which resolved the issue.

Lastly, the default operating range of the battery is between 10% and 90% SOC. A test that required charging to 100% and discharging to 0% could only be operated between 98% and 4% as even with the BYD protection override facility the system still was operating in failsafe mode.

Override on SOC – imbalance between strings

During one of the initial tests, the operation to charge to a specific SOC resulted in a small override of the target and a subsequent discharge to achieve the required SOC. The override was relatively small and corrected within a five-minute period. No further overrides or undershoots were experienced and it is believed to have related to an imbalance between the SOC of the strings.

This was potentially related to the Emergency stop tests which would have shut down the battery to ensure safety and could have introduced imbalance between the strings. The overshoot was seen soon after the emergency stop tests and was not seen after the battery had recalibrated the SOC of the strings.