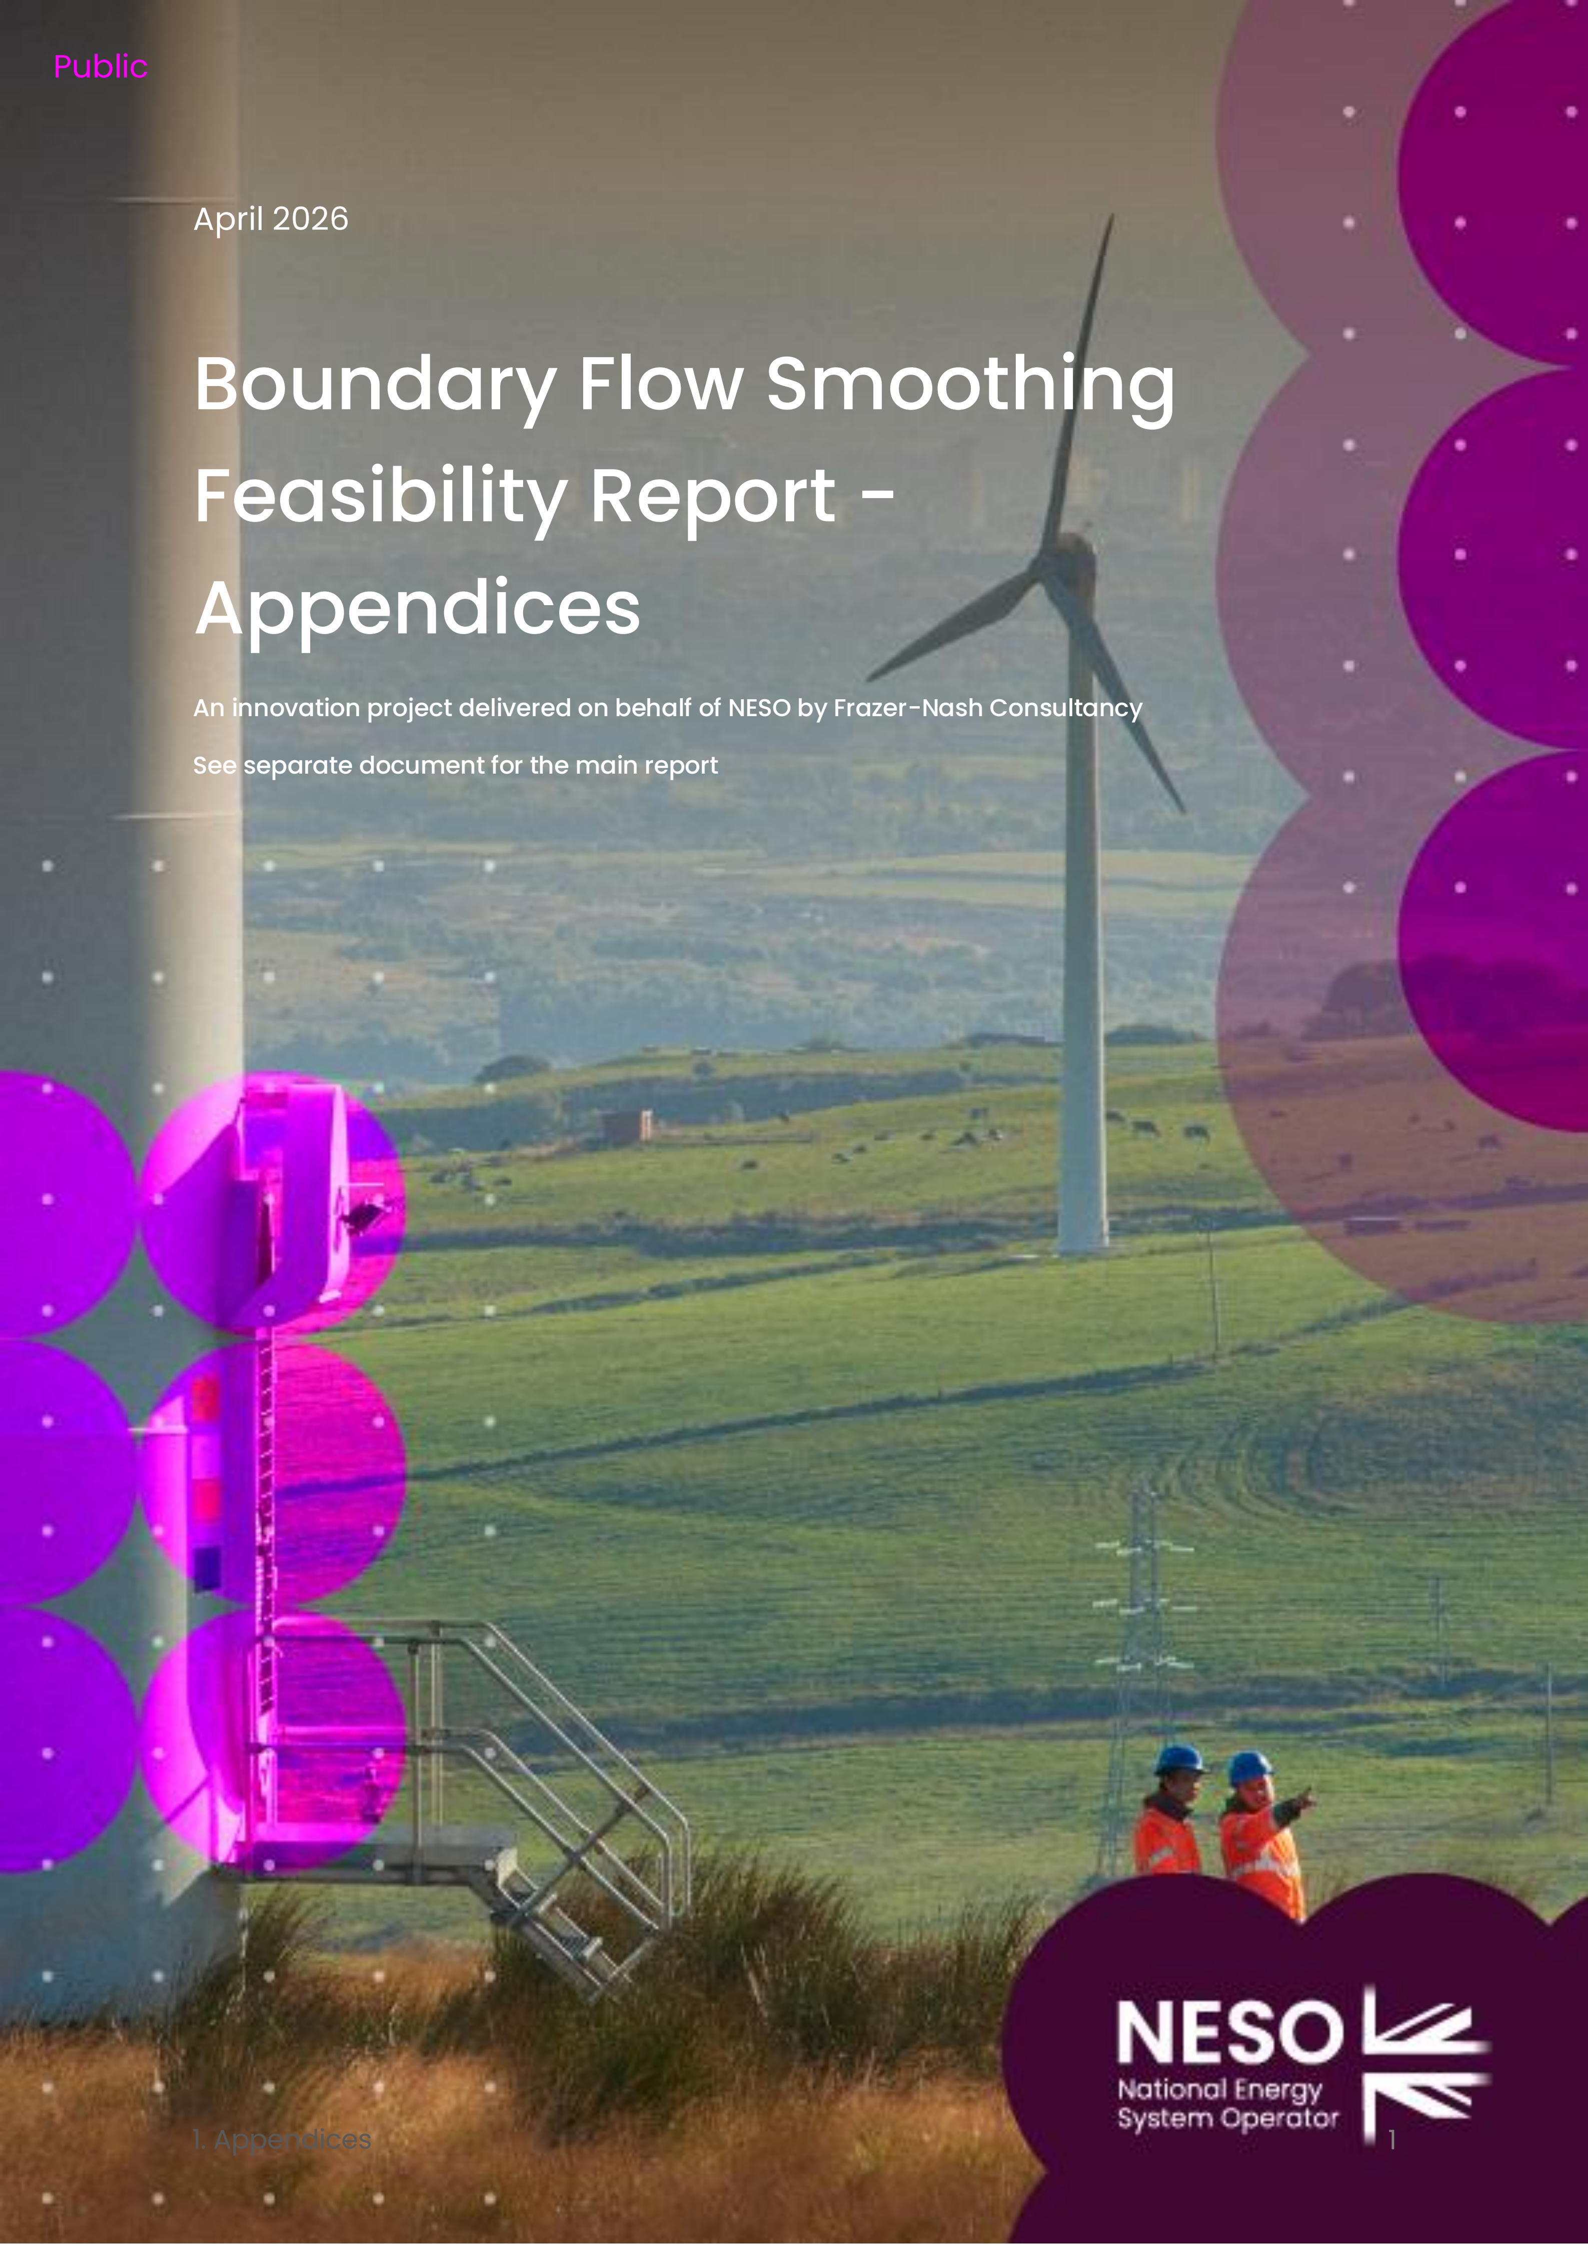

April 2026

Boundary Flow Smoothing Feasibility Report - Appendices

An innovation project delivered on behalf of NESO by Frazer-Nash Consultancy

See separate document for the main report

Contents

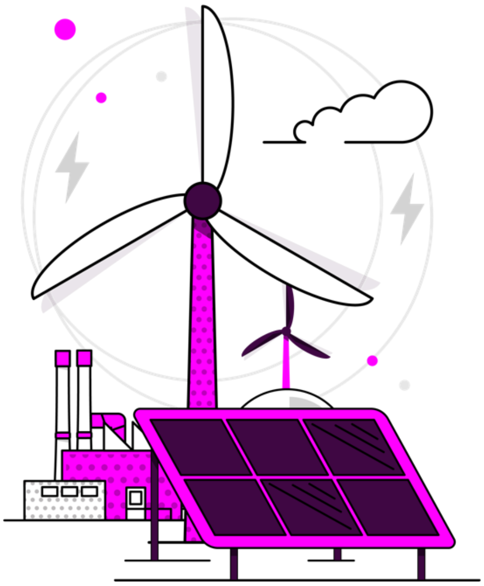

1. Appendices

Appendix A – Current Operational Practices

Transmission boundary capacity on the GB electricity network is often insufficient to effectively transport generation to demand areas and maintain system stability. High variability in boundary flows, often driven by renewable generation and fluctuating demand, can lead to frequent breaches of the boundary limits. Such breaches necessitate constraint actions, which are expensive and carbon-intensive, as constrained renewable generation is often replaced by fossil-fuel-based alternatives.

Engagement with the Electricity National Control Centre (the Control Room) established that boundary flows are ordinarily operated with no official safety margin in place to act as a buffer against the peaks of volatile flow. Instead, they aim to control flows to operate up to the constraint limit for the boundary. However, when stability constraints exist, a 100 MW margin will be implemented.

Analysis of historical data shows that boundary flows frequently exceed constraint limits.

Constraint limits are calculated by the planning team on a day-ahead basis using a network model based in DIgSILENT PowerFactory. Following this, outage plans and constraint limits are passed to the Control Room. The strategy team in the Control Room reviews and reassesses these limits based on recent changes, as well as the latest demand and generation forecasts, before issuing the Transmission System Operating Plan (TSOP), which is transferred to the transmission team four hours prior to real time. The limits are reviewed and updated in real time according to Power Network Analysis (PNA) results, which are updated every 10 minutes. The PNA is a network model which uses real-time (supervisory control and data acquisition(SCADA)) data. The updated thermal and voltage limits from PNA are used to inform the necessary Balancing Mechanism (BM) actions to adhere to these limits.

Despite these measures, short-term exceedances still occur due to manual interventions or competing priorities such as frequency control.

Appendix B – Literature Review

A thorough investigation into examples of boundary flow smoothing concepts or similar has been undertaken, which has included reviewing academic papers and industry /operator publications on projects.

From the literature review, no published examples were found within power systems globally that replicate exactly the concept being proposed as part of this project. However, Virtual Transmission Line (VTL) concepts are similar to the service being investigated. VTL can help to utilise more transmission line capacity and offer a short-term solution for constraints whilst new transmission lines are built.

Virtual Transmission Lines Overview

The key VTL concept utilises battery energy storage systems (BESS) at either end of a transmission line to use “Storage as Transmission”. As the transmission line reaches its limit, the BESS closest to the generation source charges while the BESS closest to the load at the other end of the VTL supplies the required additional load. This reduces the amount of renewable power constrained [1]. The concept is shown in Figure 1.

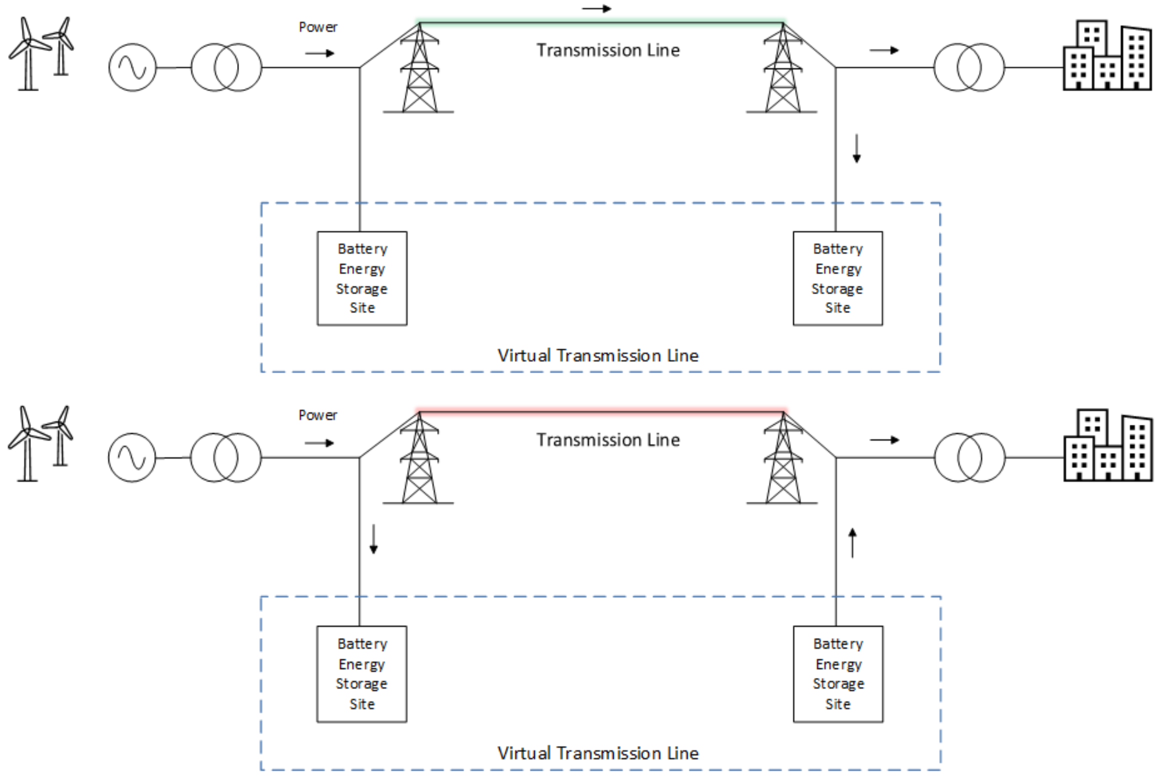

Figure 1: Virtual transmission line overview.

VTL and Storage as Transmission have generated considerable interest in the form of research, pilot projects and some operational examples [2] [3] [4] [5] [6] [7]. The rest of this section provides an overview of the most relevant project examples found.

Key Insights from Literature

The research has found a number of pilot projects that are deployed or implemented across the world, with project locations shown in Figure *2*.

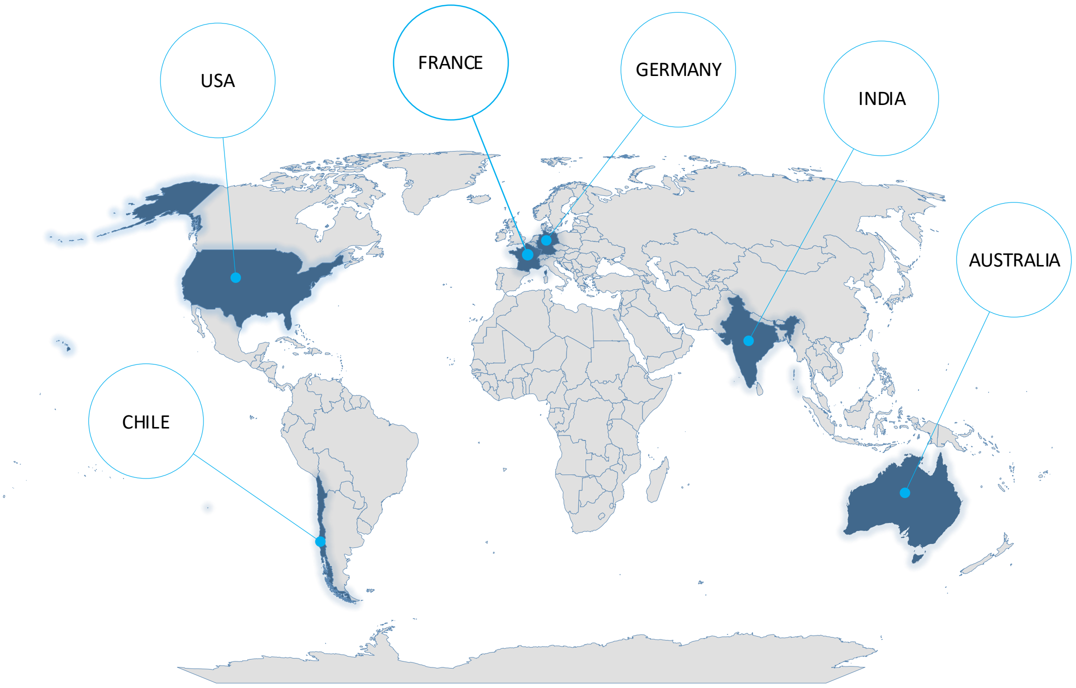

Figure 2: Relevant project examples locations.

Germany: Netzbooster

In Germany, the transmission system operators are implementing the "GridBooster" portfolio, which includes 250 MWh of energy storage capacity to ensure grid stability and lower network costs by providing backup transmission capacity [8].

This storage will enable more efficient operation of key constrained transmission lines running from north to south, reducing the need for redispatch of generation and bringing estimated annual cost savings of approximately 130 million euros.

The system aims to mimic N-1 redundancy so that both existing lines may be used fully. In a fault event, the Netzbooster battery will discharge. In a constraint event, it will store excess energy. The GridBooster initiative should save money for consumers and improve grid efficiency along congested transmission corridors [9] [10] [11] [12]. A visualisation of the project is shown in Figure 3.

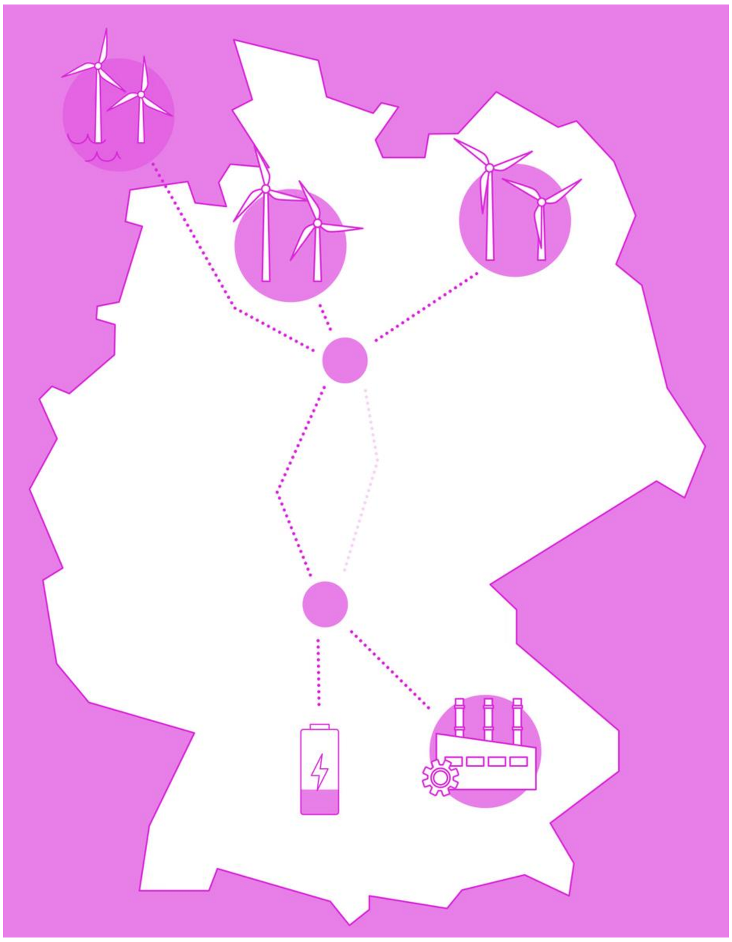

Figure 3: Netzbooster project visualisation.

Australia

In Australia, storage as transmission is being implemented and considered to enhance grid stability and increase transmission capacity [13] [14].

A 30 MW, 30 MWh Fluence battery system is deployed at the Ballarat Terminal Station in Victoria, providing grid stability services and flexible capacity at a critical transmission junction. Additionally, battery-based storage is being considered alongside traditional infrastructure to add capacity on key interstate transmission lines (interconnectors) that connect state electricity networks within the National Electricity Market.

These storage systems aim to reduce congestion, mitigate constraints, prevent outages, and enable better import/export of electricity between states during peak demand periods. Australia has over 1 GW of potential for deploying storage to bridge transmission capacity gaps and improve network reliability.

One of the key challenges in implementation is that existing investment frameworks for transmission networks are not designed to support the rapid deployment of storage as transmission assets.

France: RINGO

The French utility company RTE has implemented a pilot VTL project. Its purpose is to increase grid integration of renewable energy and to optimise power flow on the network.

The project comprises three geographically dispersed BESS with coordinated operation and creates a virtual link between regions with differing renewable output profiles. Energy may be stored at one location and released at another, enabling power to be either withdrawn or released without exceeding transfer limits of transmission lines. To comply with regulatory requirements that transmission companies may not buy or sell energy, the process of charging and discharging occurs simultaneously [15].

The BESS are 10 to 12 MW capacity and 2-hour duration assets controlled by an algorithm that adapts to the weather [16].

USA

In the U.S.A., California's Pacific Gas and Electric selected a 10 MW energy storage project to be part of a portfolio of solutions for the transmission network, the first of its kind in the country.

Chile

In Chile, a proposal for two new 200 MW energy storage sites has been submitted to the Chilean regulator.

If approved, the two new VTL projects will provide N-1 capacity to relieve congestion on the system. In particular this will assist the transmission of 700 MW of renewable generation coming online on the constrained side of the transmission corridor [16].

Literature Review Summary

An extensive literature review has concluded that, while there are numerous international examples of the use of storage to support transmission, **there are no published examples of transmission line flow smoothing being implemented to reduce constraints.** The differences in the aims of these international examples mean that there are no obvious lessons to be learnt from VTL projects that can be applied to this project.

The literature review has highlighted the following:

- VTL systems have been successfully implemented, with more planned, in various electricity transmission networks internationally.
- Many of these projects also deploy the assets during fault or outage scenarios whereas the boundary flow smoothing project is intended for general operational use.
- Storage as Transmission could be used in GB to boost transmission line capacity. However, it might bring complications in ownership, given the separation between NESO, as the system operator, and the transmission owners.

- The timeline for building Storage as Transmission assets may not be quick enough to bring benefits before planned transmission capacity upgrades. The boundary flow smoothing concept is planned as a short-term solution initially.

These projects demonstrate the feasibility of using storage to enhance transmission capacity but differ from the proposed smoothing concept, which targets operational volatility rather than outage or fault scenarios.

Appendix C – Data Overview

In delivering the project, the following datasets shown in Table 1 were used. The datasets have allowed us to analyse boundary power flows and see how these relate to BM actions in terms of energy and cost.

Table 1: Data sources and descriptions

Data received	**Description**
Bid Offer Acceptance (BOA) data	This is half-hourly Balancing Mechanism (BM) bids and offers data, filtered for constraint tagged actions. It provides information on whether constraints were import or export; the reason for the constraint, e.g. thermal or stability; replacement energy offers; the network location of the Balancing Mechanism Units (BMUs) ; the energy volume and price information.
Boundary flow meter data	This contains the boundary flow values in MW every second.
Constraint limit data	This provides the minute-by-minute constraint limit in MW for each boundary. Note: the Control Room reassesses the limit every 10 minutes, so generally the changes in constraint limits happen on a less granular basis.
SCOTEX (B6) intertrip constraint limits	SCOTEX (B6) uses an intertrip scheme to manage constraints on that boundary. When the intertrip scheme is active, the constraint limits are recorded in a separate dataset. The constraint limit dataset flags these active intertrip periods by recording a limit of 10,000 MW.
Network node file	This file identifies the network node that BMUs are connected to, and which nodes are feeding power across boundaries.
Palantir data	The Palantir data is more granular BM data. The volume and price information is provided on a minute-by-minute basis. It includes all bids and offers rather than just being filtered for the constraint actions.

Adjusted Flow Construction

The historical metered flow includes the effects of Control Room decisions. Therefore, an 'adjusted flow' has been calculated which attempts to recreate a counterfactual flow that has these Control Room decisions removed. The adjusted flow is calculated by unwinding the historical bids and offers from the metered boundary flow, as shown in Figure 4. This approach aligns with an existing NESO approach.

Metered Flow + Bids - Offers = Adjusted Flow

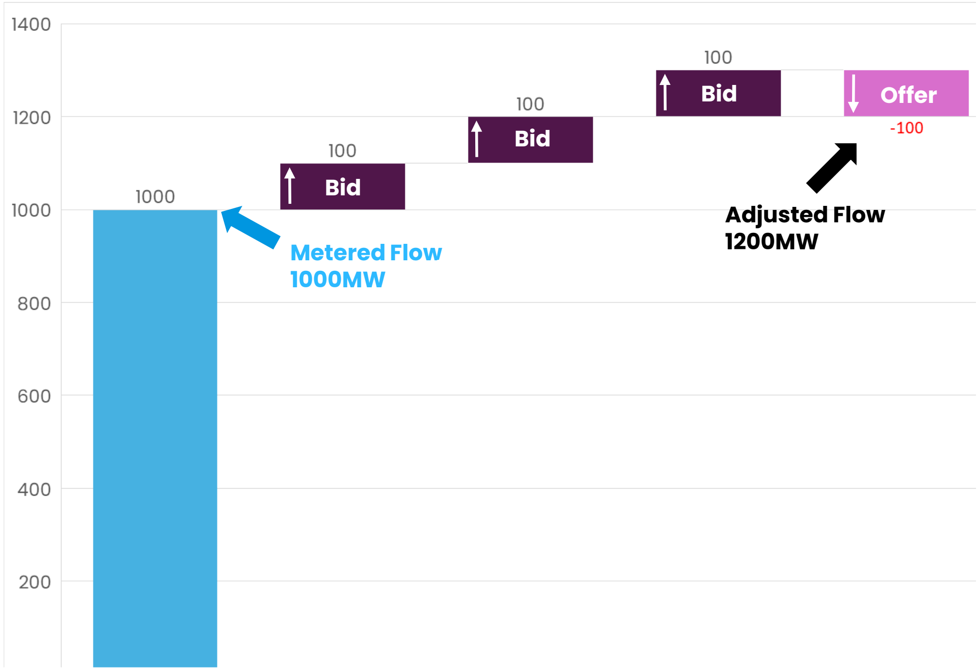

Figure 4: Adjusted flow construction.

Constrained periods are identified when the adjusted flow exceeds the constraint limit, signalling that BM actions would be needed.

Appendix D – Understanding the Scale of Boundary Flows and Constraints

Second-by-second boundary power flow and constraint data for the period 01/01/2024 to 31/03/2025 have been analysed. This initial analysis provided scale and context to the boundary constraint problem.

As described in the previous section, constrained periods are identified as being those when the adjusted flow is above the boundary constraint limit.

Analysis has investigated:

- How much the flow exceeds limits during constrained periods.
- Where metered flows are below the constraint limits, the additional energy that could be transferred through the boundary has been calculated (spare transfer capacity). This provides an upper bound for the volume of additional energy that could be transferred through the boundaries.

NESO identified that there are occasional quality issues with the constraint limits data. This results in an over-estimation of the number and severity of constraint exceedances in the analysis.

The metered boundary flow data and adjusted flows are shown in the following sections alongside plots of the data. The data from just constrained periods is also provided. All values are rounded to 3 significant figures. The statistics have been calculated from the flow data, constraint limits and BOA cost data. The different bars in the figures denote the following:

- The blue bar shows the total observed net flow of energy.
- The yellow bar denotes the total available energy transfer capacity, if the boundary flow had operated continuously at its limit.
- The green bar gives the energy transfer that occurred in exceedance of the boundary flow limit.
- The red bar indicates the ‘unused’ capacity – i.e. the difference between the blue and yellow bars.
- The brown bar denotes the total volume of energy turned down via the Balancing Mechanism and tagged to the boundary.

The analysis of the data demonstrates the large amounts of energy needing to be turned down and the differences between the constraint limits and flows.

The six boundaries form a nested set as described in the Boundaries of Interest section [17]. BOAs are tagged to the boundary closest to the constrained asset, which is not necessarily the boundary where the actual constraint was present. This means that the recorded turn-down volumes cannot be confidently attributed to constraints on the reported boundaries but could be the result of a constraint on any of the boundaries further south.

Metered Flows

Constrained Periods

Table 2: Metered flow data for constrained periods during analysis period (01/01/24 to 31/03/25)

Boundary	Metered Flow (MWh)	Limit (MWh)	Difference (Flow – Limit) (MWh)	Energy Transferred in Exceedance of the Flow Limit (MWh)	Action Cost (£)	Constrained Energy[1] (MWh)	Time Above Limit (days)	Percentage of Time Above Limit for Constrained Period
SSEN-S	4,680,000	5,130,000	-457,000	12,600	83,100,000	2,400,000	10.90	10.80%
SSE-SP2	7,530,000	8,100,000	-571,000	73,600	68,300,000	2,830,000	40.00	24.80%
NKILGRMO	3,740,000	4,650,000	-916,000	1,040	5,320,000	686,000	0.81	0.01%
SCOTEX	8,290,000	8,640,000	-344,000	411,000	26,600,000	562,000	16.60	21.50%
HARSPNBLY	11,100,000	11,400,000	-285,000	309,000	16,600,000	467,000	26.70	20.90%
SSHARN3	12,900,000	13,500,000	-666,000	18,400	19,300,000	344,000	6.90	9.90%

---

[1] Calculated by summing the total energy of constraint-tagged BOAs

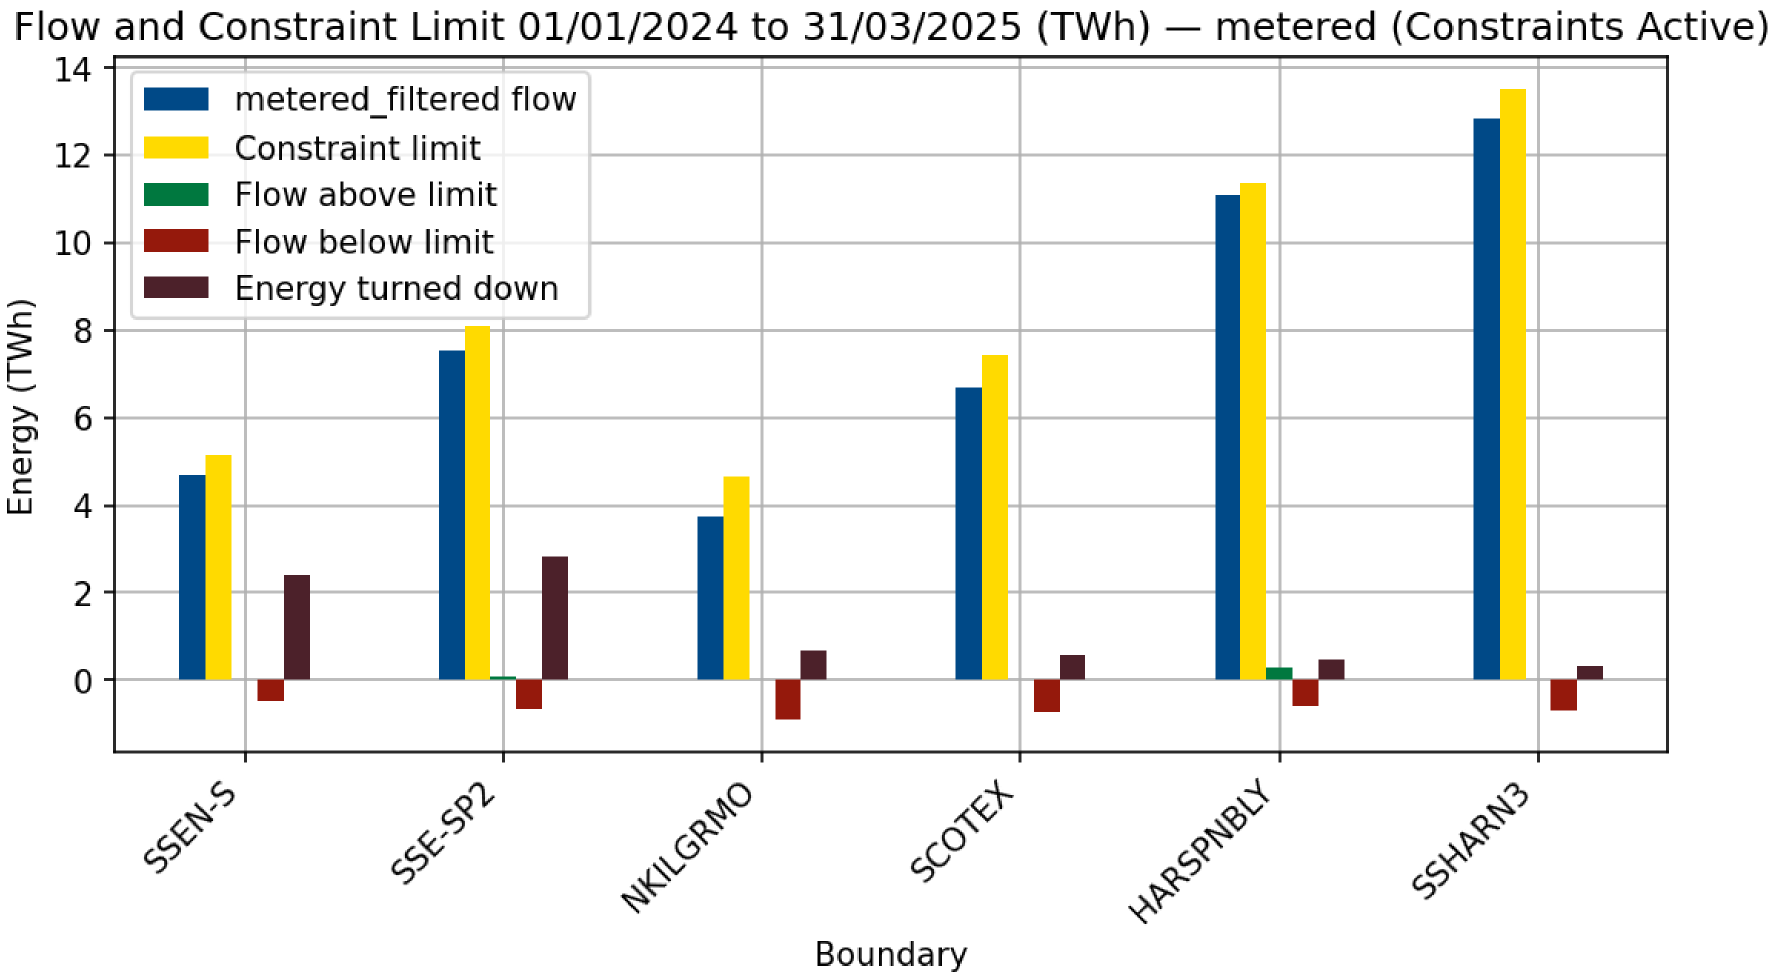

Figure 5: Metered flow data for constrained periods during analysis period (01/01/24 to 31/03/25).

Full Period

Table 3: Metered flow data for full analysis period (01/01/24 to 31/03/25)

Boundary	Metered Flow (MWh)	Limit (MWh)	Difference (Flow – Limit) (MWh)	Energy Transferred in Exceedance of the Flow Limit (MWh)	Action Cost (£)	Constrained Energy (MWh)	Time Above Limit (days)	Percentage of Time Above Limit for Constrained Period
SSEN-S	14,000,000	23,500,000	-9,500,000	12,600	96,800,000	2,810,000	10.90	3%
SSE-SP2	15,500,000	23,300,000	-7,820,000	73,600	71,800,000	3,060,000	40.00	11%
NKILGRMO	12,000,000	26,400,000	-14,400,000	1,050	4,620,000[2]	758,000	0.84	0%
SCOTEX	21,300,000	38,200,000	-16,900,000	411,000	37,900,000	866,000	16.60	5%
HARSPNBLY	31,100,000	51,600,000	-20,500,000	309,000	17,300,000	494,000	26.70	7%
SSHARN3	49,200,000	70,700,000	-21,500,000	18,400	20,400,000	370,000	6.90	2%

---

[2] This action cost total for the full period is lower than the constrained period due to the presence of negative BOAs reducing the total action cost.

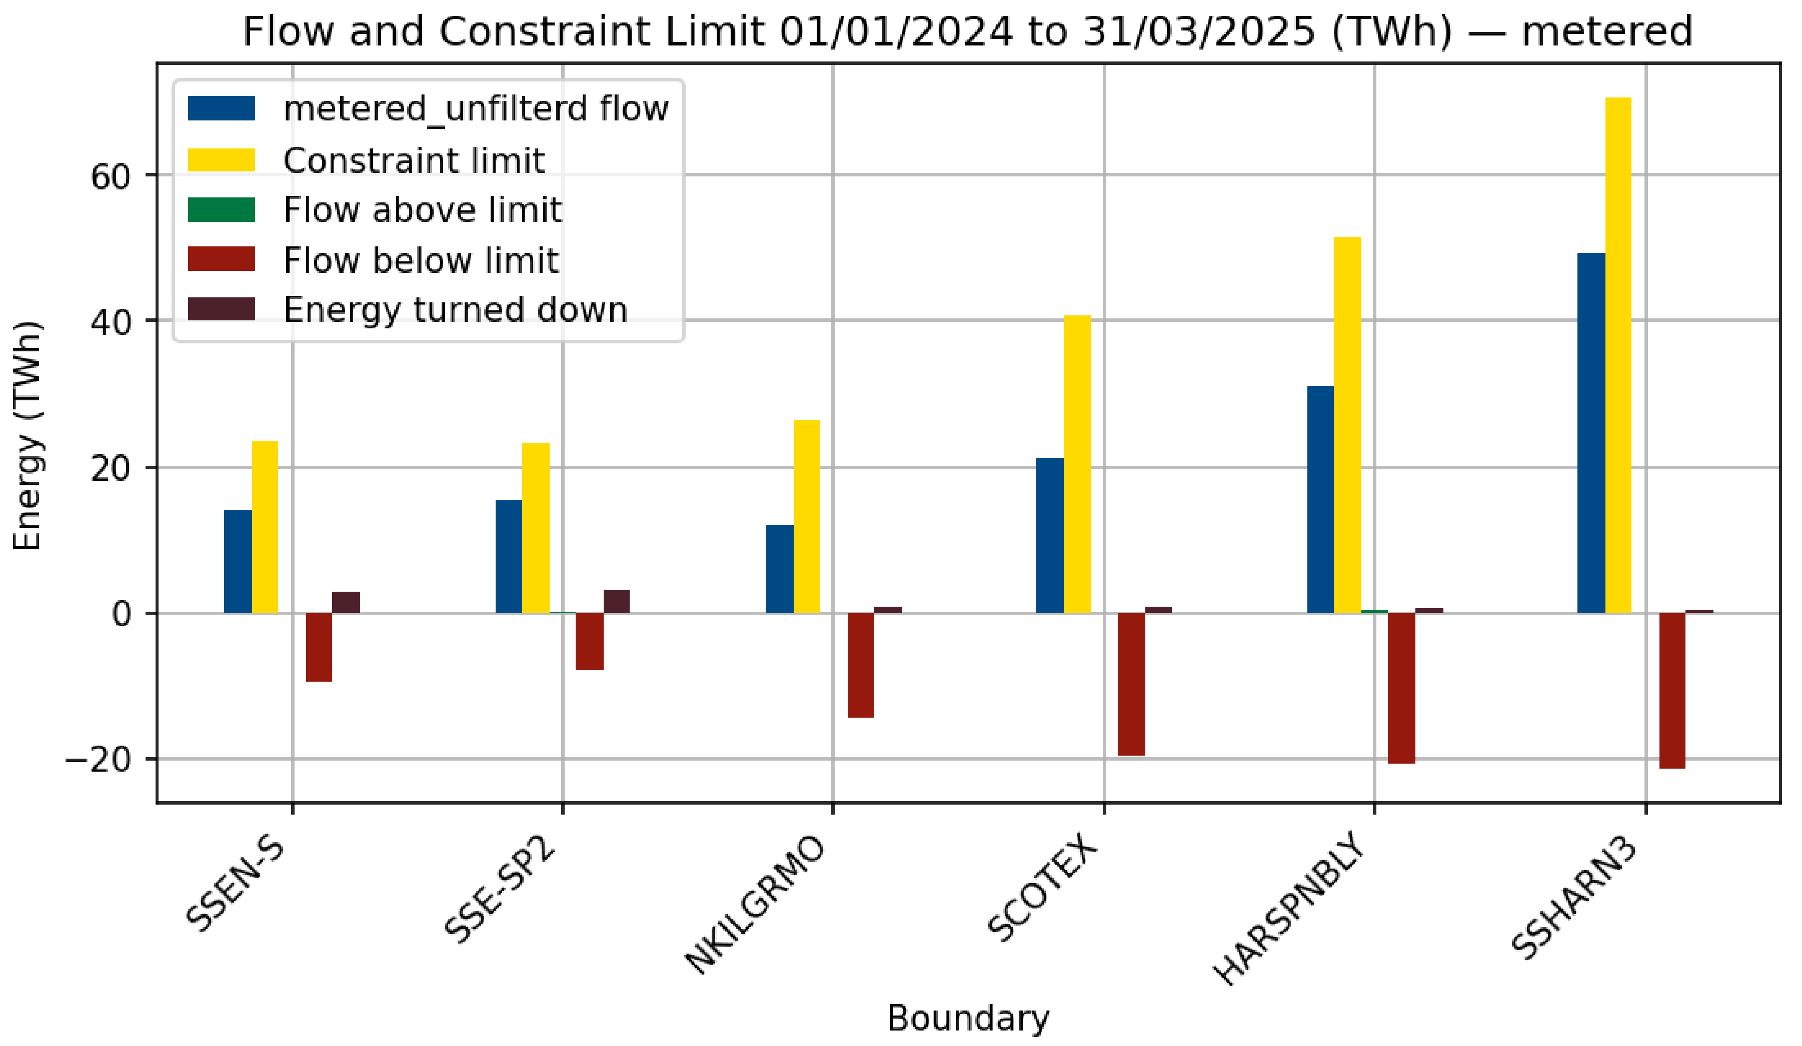

Figure 6: Metered flow data for full analysis period (01/01/24 to 31/03/25).

Adjusted Flows

Constrained Periods

Table 4: Adjusted flow data for constrained periods during analysis period (01/01/24 to 31/03/25)

Boundary	Adjusted Flow (MWh)	Limit (MWh)	Difference (Flow – Limit) (MWh)	Energy Transferred in Exceedance of the Flow Limit	Time Above Limit (days)	Percentage of Time Above Limit for Constrained Periods
SSEN-S	5,990,000	5,130,000	856,000	14%	101	22.1%
SSE-SP2	9,970,000	8,100,000	1,870,000	19%	161	35.4%
NKILGRMO	5,540,000	4,650,000	885,000	16%	85.7	18.8%
SCOTEX	10,200,000	8,640,000	1,530,000	15%	77.3	17.0%
HARSPNBLY	13,000,000	11,400,000	1,610,000	12%	86.3	19.0%
SSHARN3	14,600,000	13,500,000	1,050,000	7%	70	15.4%

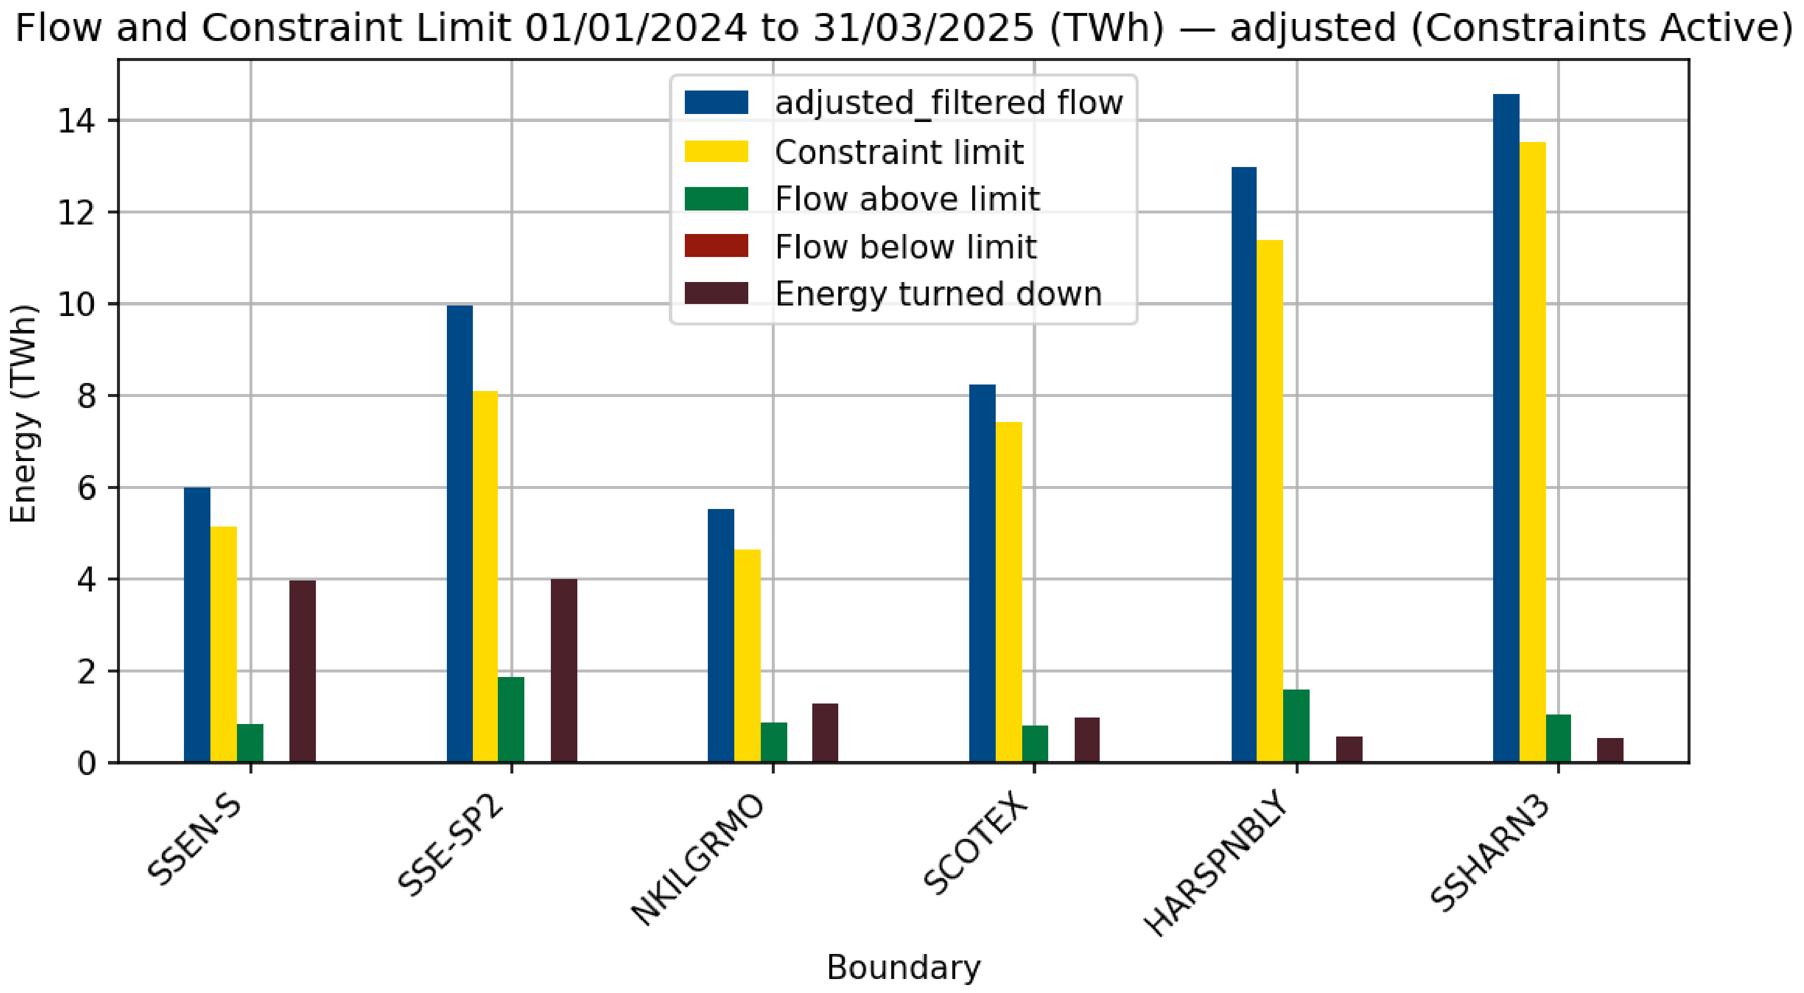

Figure 7: Adjusted flow data for constrained periods during analysis period (01/01/24 to 31/03/25)

Full Period

Table 5: Adjusted flow data for full analysis period (01/01/24 to 31/03/25)

Boundary	Adjusted Flow (MWh)	Limit (MWh)	Difference (Flow – Limit) (MWh)	Energy Transferred in Exceedance of the Flow Limit (MWh)	Time Above Limit (days)
SSEN-S	15,900,000	23,500,000	-7,600,000	856,000	101
SSE-SP2	18,500,000	23,300,000	-4,800,000	1,870,000	161
NKILGRMO	15,000,000	26,400,000	-11,400,000	885,000	85.7
SCOTEX	24,800,000	38,200,000	-13,400,000	1,530,000	77.3
HARSPNBLY	34,600,000	51,600,000	-17,000,000	1,610,000	86.3
SSHARN3	52,700,000	70,700,000	-18,000,000	1,050,000	70

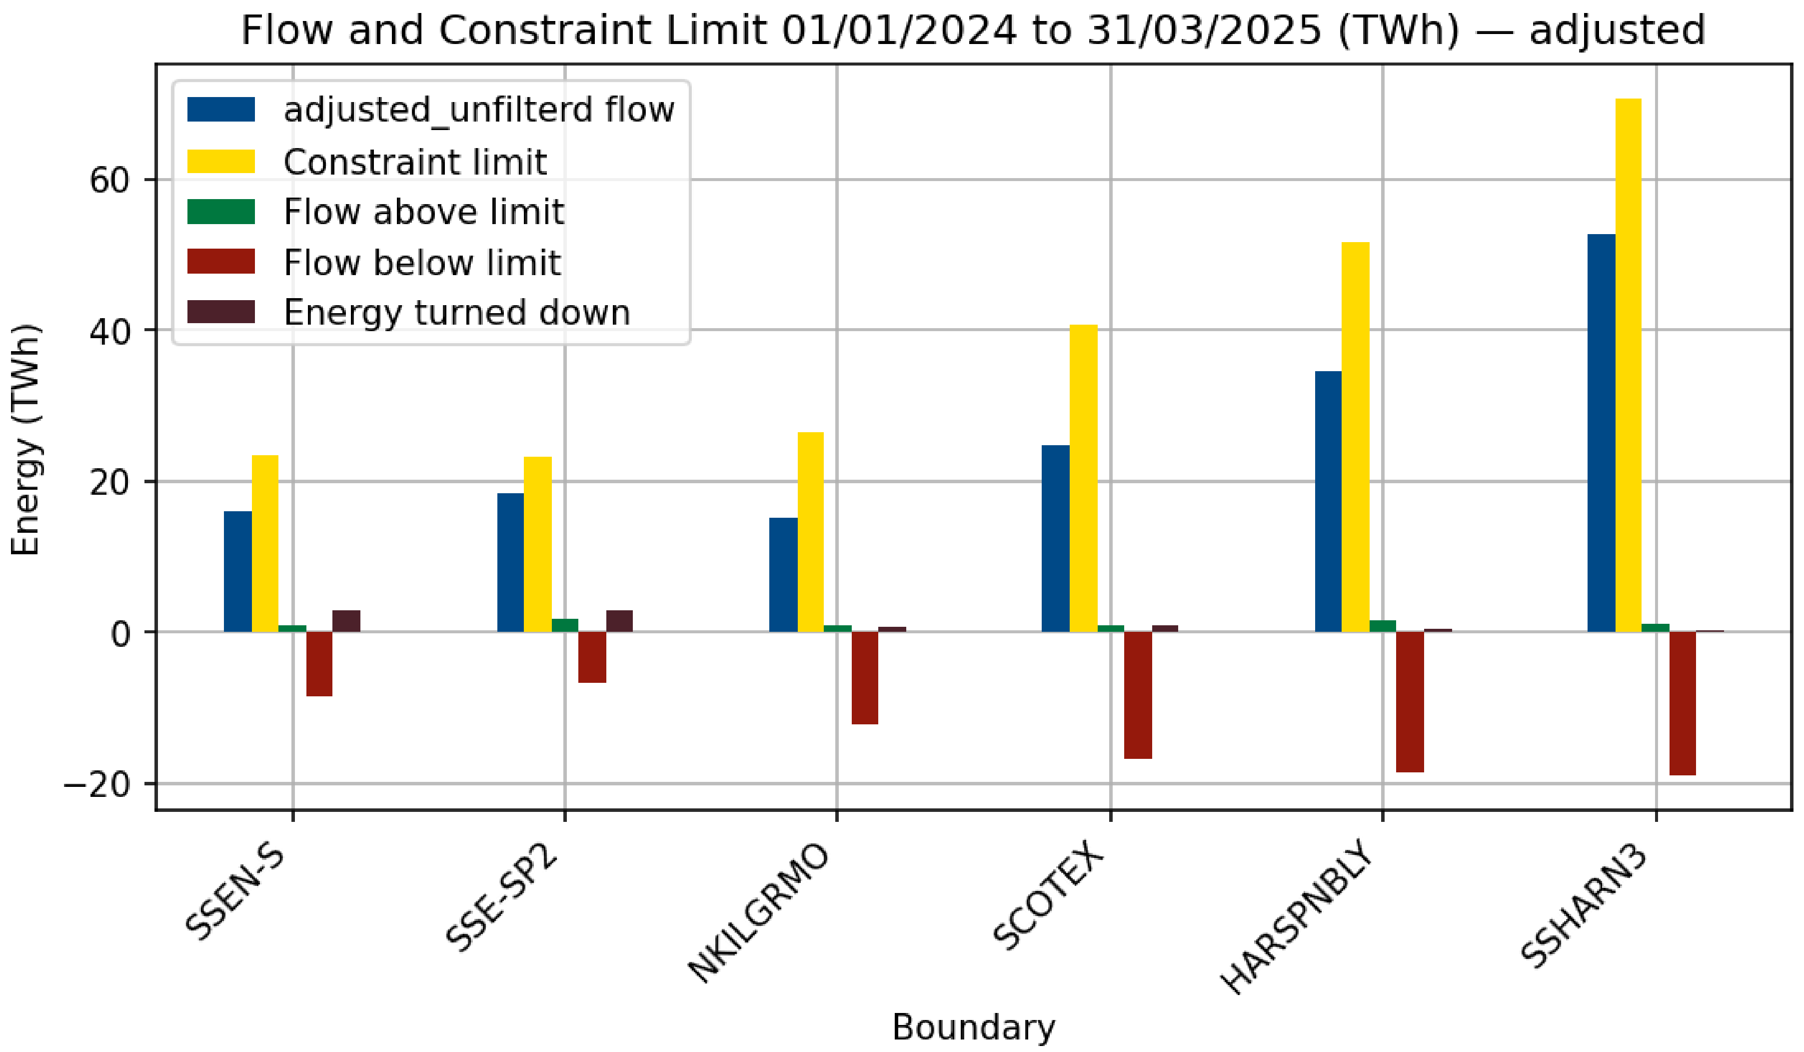

Figure 8: Adjusted flow data for full analysis period (01/01/24 to 31/03/25).

Comparison of Metered to Adjusted Flows

Table 6 is a comparison between the adjusted and metered flows for constrained periods (when the adjusted flow is above the limit). Adjusted flows are much higher than metered flows across all boundaries, showing the scale of BM activity and turn-down. SSE-SP2 shows the largest absolute gap, with adjusted flows exceeding metered flows by 2.44 million MWh. NKILGRMO has the largest relative difference, with adjusted flows 32% higher than metered flows.

Table 6: Comparison of metered to adjusted flow data for constrained periods during analysis period (01/01/24 to 31/03/25)

Boundary	Metered Flow (MWh)	Adjusted Flow (MWh)	Limit (MWh)	Energy Transferred in Exceedance of the Flow Limit – Metered (MWh)	Energy Transferred in Exceedance of the Flow Limit - Adjusted (MWh)	Time Above Limit - Metered (days)	Time Above Limit – Adjusted (days)
SSEN-S	4,680,000	5,990,000	5,130,000	12,600	856,000	10.90	101.00
SSE-SP2	7,530,000	9,970,000	8,090,000	73,600	1,870,000	40.00	161.00
NKILGRMO	3,740,000	5,540,000	4,650,000	1,040	885,000	0.81	85.70
SCOTEX	8,290,000	10,200,000	7,420,000	411,000	1,530,000	16.60	77.30
HARSPNBLY	11,100,000	13,000,000	11,400,000	309,000	1,610,000	26.70	86.30
SSHARN3	12,900,000	14,600,000	13,500,000	18,400	1,050,000	6.90	70.00

Appendix E – Distribution of Mean Flow

The plots in Figure 9 and Figure 10 illustrate the breakdown of flow values for each boundary into constrained and unconstrained flow period components. In each case, the plot on the left presents the distribution of unconstrained flows in pink alongside the distribution of all flows in blue. The plots on the right display the distribution of flows when the flow is constrained. This constrained flow corresponds to the difference between the blue and pink lines shown in the left-hand plots.

The six boundaries can be grouped into three categories based on the shapes of their flow distributions. The three most northerly boundaries exhibit a distribution with a primary peak in constrained flow at approximately 1800 MW and a smaller secondary peak at approximately 2400 MW. In contrast, the three southerly boundaries display a different pattern, characterised by a less clearly defined peak. This peak increases from 5000 MW at SCOTEX to 7300 MW at SSHARN-3.

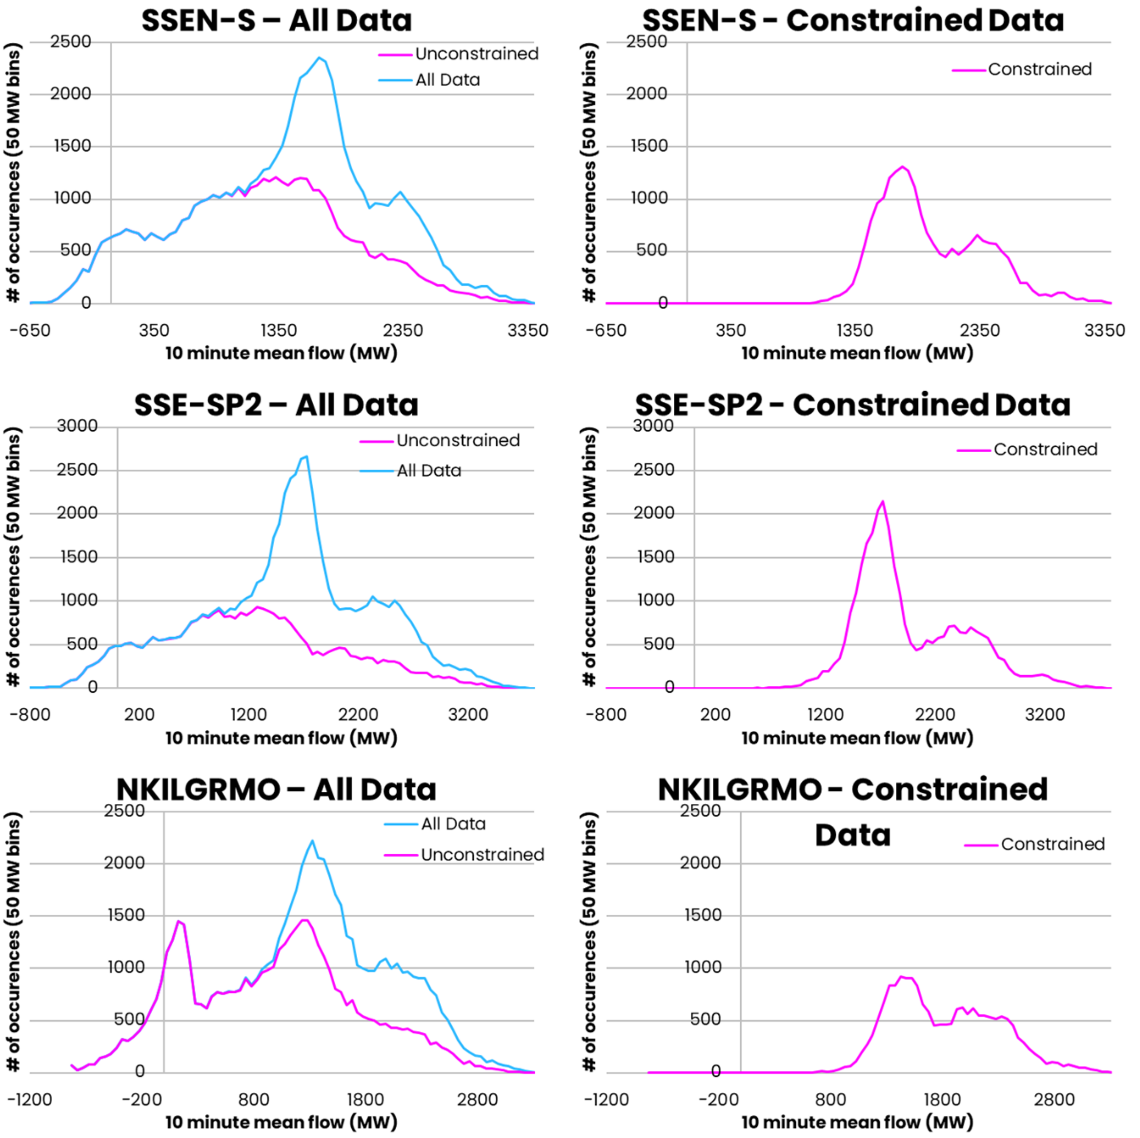

Figure 9: Distribution of mean flow for the three most northernly boundaries.

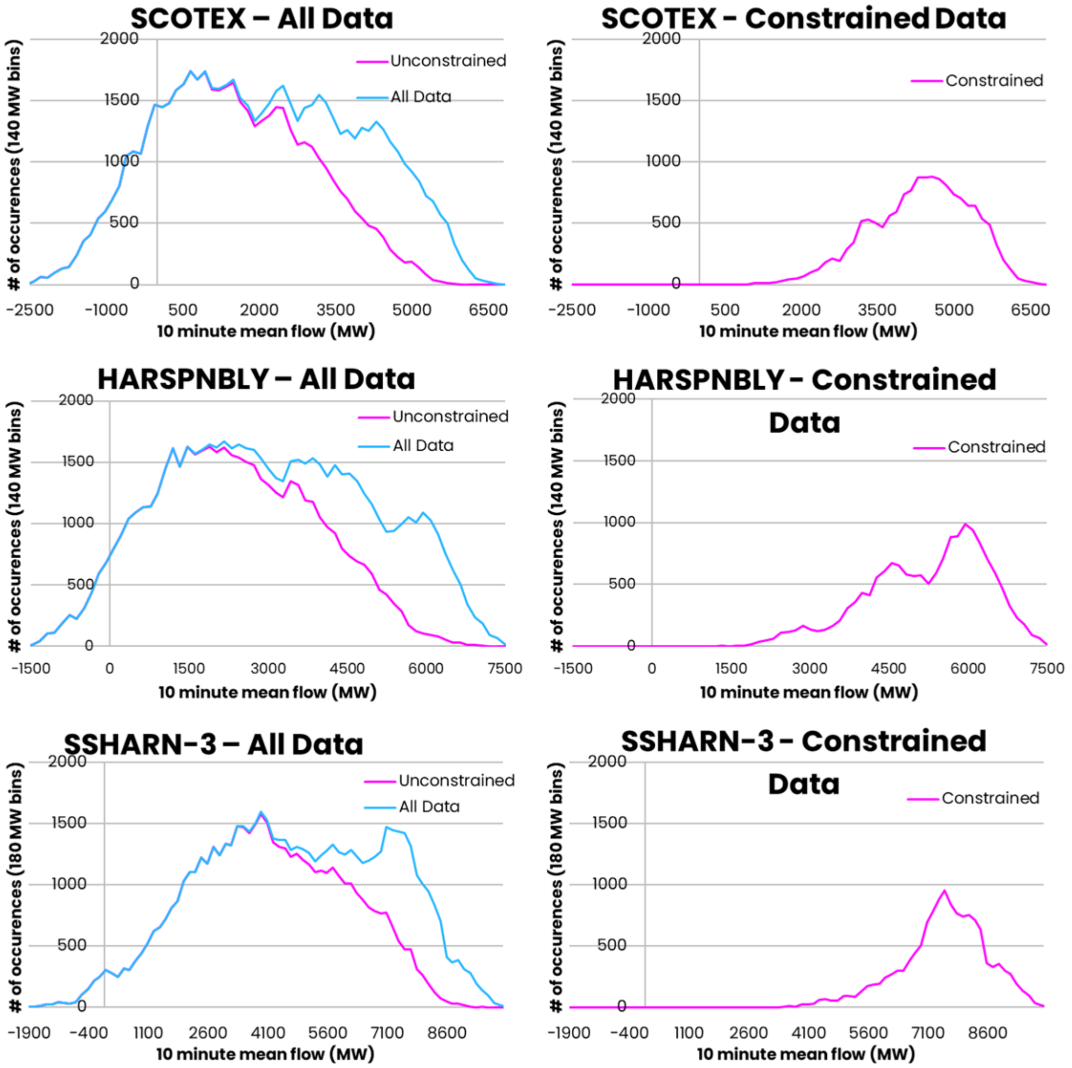

Figure 10: Distribution of mean flow for the three most southernly boundaries.

Appendix F – Variability and Volatility Analysis

As part of the project, boundary flow variability and volatility was explored. The variability of a time series describes the spread of the data points and their distribution about the mean value. It is typically measured by range and variance. Low range and/or variance means low variability.

Volatility is the variability of the time series over a defined time interval, for example, the variance or range of boundary flow values over a day, over an hour or over a period that typifies the operational realities within the Control Room.

The Control Room calculates boundary constraint limits every 10 minutes. A volatility assessment over a period of between 5 and 15 minutes would reflect operational timescales for the Control Room to respond to fluctuations in the altered flow.

A number of statistics were investigated to describe boundary flow time series. Including metrics to quantify the variability and volatility of the flow as well as analysis of the typical frequencies of the fluctuations in flow. These included:

- Range and variance
- Spectral techniques (Fourier analysis)
- Realised volatility

After investigation these have not been used to assess the smoothing effect and instead the statistical approach in the main body of the report [17] was chosen as a better approach. However, some of these did feed into the smoothing assessment metrics that were used. A description of the metrics is included in Appendix G – Smoothing Algorithms.

Range and Standard Deviation

Two simple statistics that can be efficiently calculated and understood are the range and standard deviation of flows across different time windows. Both of these statistics increase as the variability increases.

To calculate the variance statistics, the flow data is grouped into half-hour periods, and the range and standard deviation of flows within each period is calculated. By plotting the cumulative distribution of these values per half-hour period, the distribution of range and standard deviation across extended time frames becomes apparent.

Range

- The range is the difference in the minimum and maximum power flow over a given period.

- A high range within a small time period suggests that the power flow has high variability.
- Over longer time periods the range becomes less useful as a statistic as it offers no insight into the distribution of values.

Standard Deviation

- The standard deviation gives a measure of how far from the mean the values within a period are.
- A higher standard deviation indicates a wider range of values.
- The standard deviation is less sensitive to outliers than more simplistic metrics such as the range.

These metrics are straightforward and easy to interpret. They are effective for assessing overall variability but less sensitive to short-term fluctuations.

Range and Standard Deviation Distribution

The figures below show the distributions of the data for the period 01/01/2024 to 31/03/2025. In Figure 11, the y-axis is the range of the flow for the given boundary in a half-hour window. The x-axis represents the percentage of half-hour periods over the period in which the range was less than the corresponding x-axis value. For example, if you take the SCOTEX boundary and split it into thirty-minute windows, 80% of those windows will have a flow range of 500 MW or less. Almost 100% will have a range of 1000 MW or less. The shape of the curve indicates how consistent the variability is over the year. The same interpretation applies to the standard deviation plot in Figure 12.

The boundaries analysed fall into two groups based on their distributions for the range and standard deviation plots. HARSPNBLY, SCOTEX, and SSHARN3 constitute one group that is characterised by higher average variance statistics, while NKILGRMO, SSEN-S, and SSE-SP2 form a second group with comparatively lower variance statistics.

Applying a smoothing algorithm is expected to reduce the variance statistics for each boundary, resulting in a general downward shift of the distribution curves and a lower maximum standard deviation (the value at 100%).

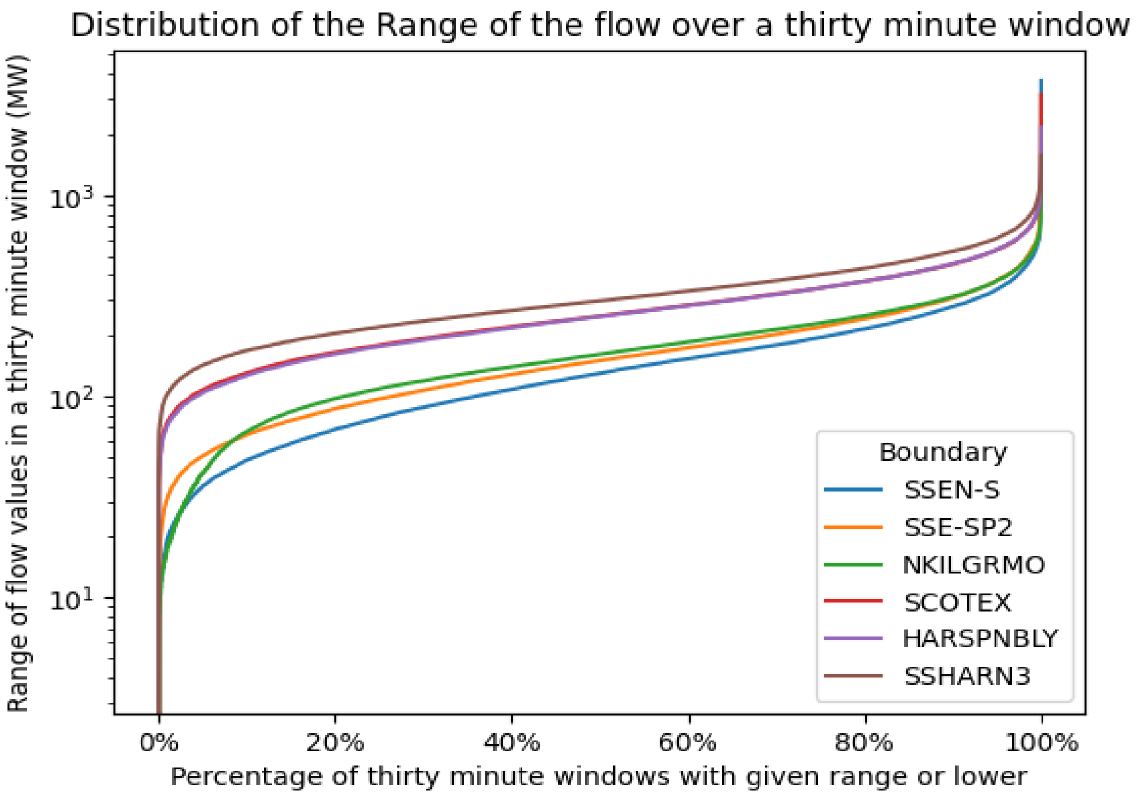

Figure 11: Distribution of the range of boundary flows in thirty-minute windows.

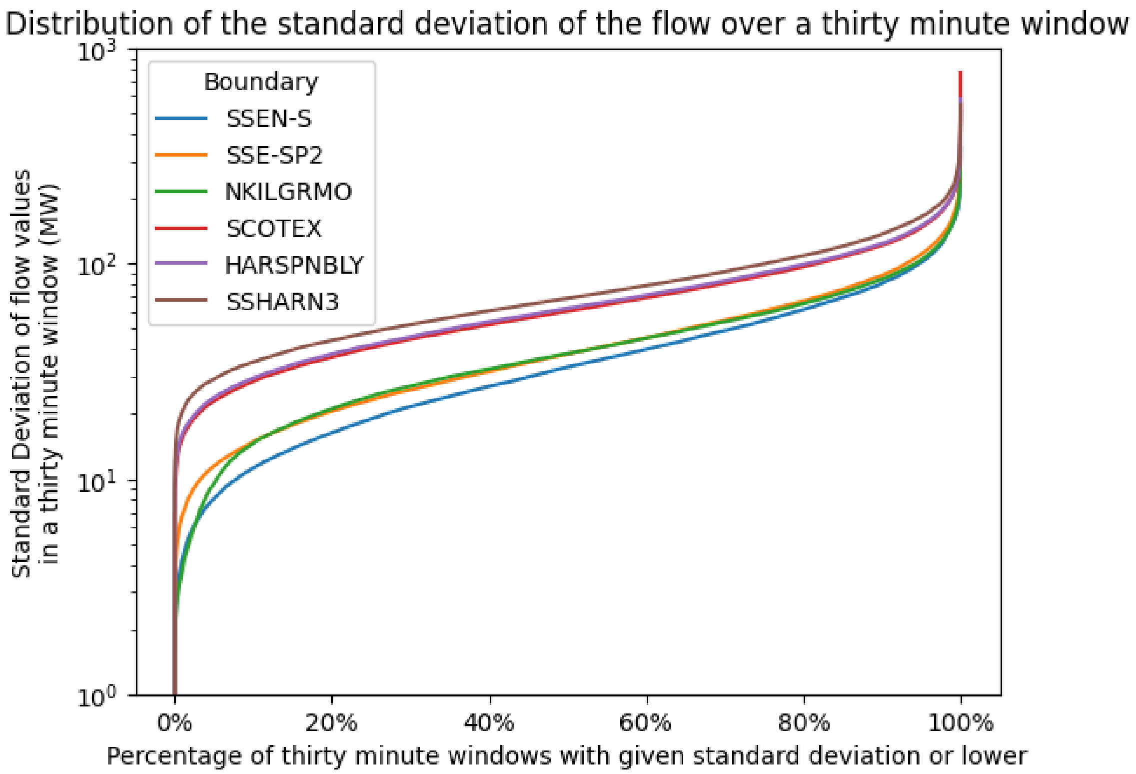

Figure 12: Distribution of the standard deviation of boundary flows in thirty-minute windows.

Fourier Analysis

Fourier analysis can be used to break down a signal into its component frequencies. With the power flow signal there will be high frequency components representing the faster fluctuations in power flow and lower frequency components, which will contribute to the longer-term trends in the signal.

When the power flow signals were broken down using the Fourier transform there were no dominant frequencies that contributed to the signal variability more than others. The presence of such signals would have allowed for further analysis into why those frequencies were present and the potential of damping them to provide a smoother power flow. Without dominant frequencies this targeted damping is not possible.

Instead, it may be possible to damp just the high frequencies to provide a smoother, less volatile flow, as shown in Figure 13. On the right are three signals to which varying levels of damping have been applied. The top chart shows all high frequencies being damped, second chart shows some high frequencies being damped and the bottom chart shows minimal damping. The smoothness increases as the higher frequency components are damped.

Preliminary results indicate that high-frequency damping could improve flow smoothness; however, this approach was not explored further due to prioritising the main analysis.

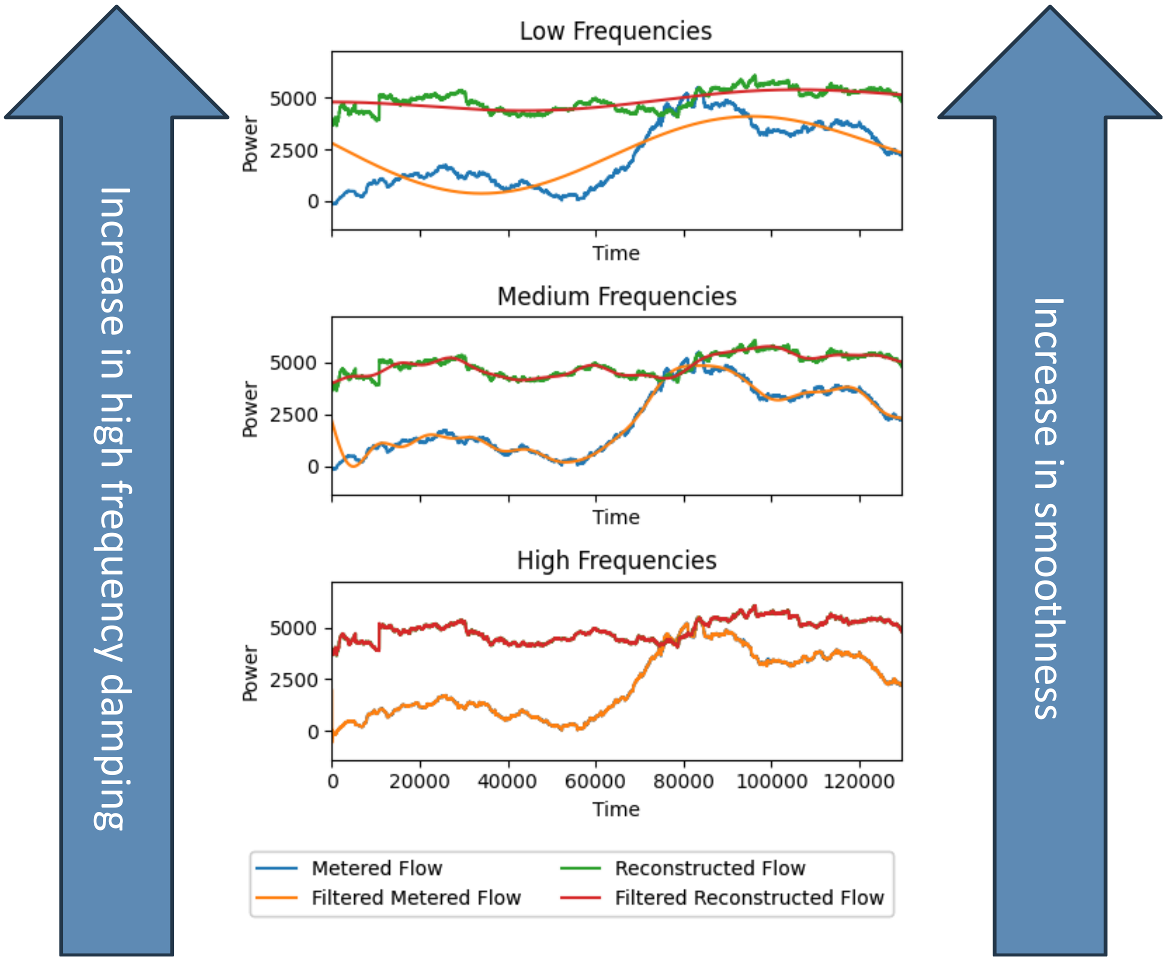

Figure 13: Fourier analysis approach and results.

Realised Volatility

The realised volatility metric can be used to assess how volatile the power flow is. It uses the historic power flows to assess how much the power flow across a boundary is changing. This metric takes the difference in power flows on a second-by-second basis and combines these differences over a rolling 30-minute period to calculate a single metric. As shown in the equation below, N will be the number of samples in the period (1800 in this case), the level of volatility to assess can be tuned by varying this.

$$V = \sqrt{\frac{1}{n}\sum_{t=1}^{n}\left(\ln\frac{P_t}{P_{t-1}}\right)^2}$$

This metric increases with more volatile flow. Unlike the range and standard deviation, which look at rolling windows and assess the values within, the realised volatility accounts for how the flow values change within the window.

Figure 14 shows that the realised volatility is on average higher for the SCOTEX than the SSEN-S boundary. This is to be expected given the power flows across the two boundaries. However, the mean values of the realised volatility are very similar for the two flows shown, which makes comparison difficult. This closeness in values along with the relative complexity when compared to other metrics has led to the decision to not use this metric in subsequent analyses.

This metric captures intra-window fluctuations but is complex and offers limited differentiation across boundaries. It was not selected for further application due to its low comparative value.

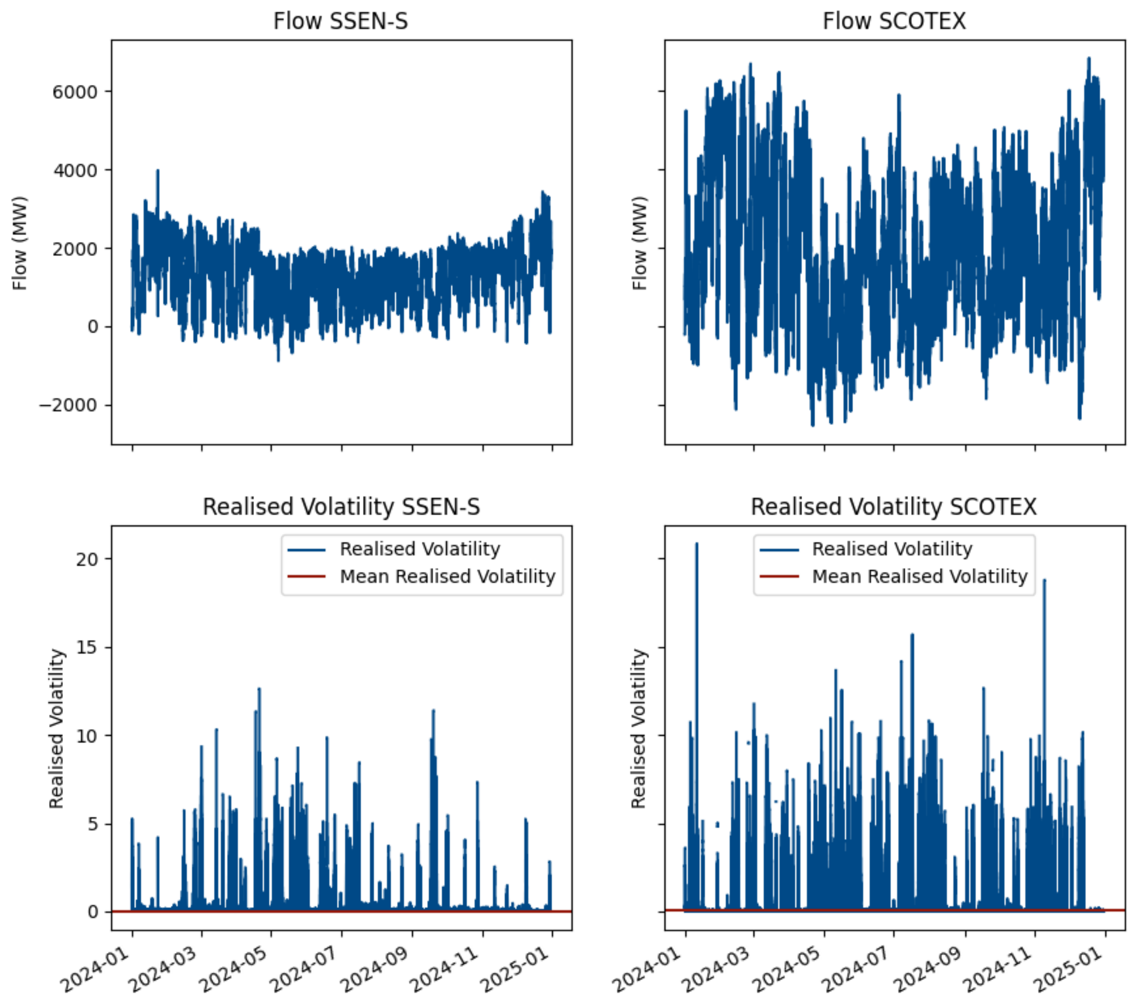

Figure 14: Realised volatility of boundary flows.

Appendix G – Smoothing Algorithms

Smoothing Algorithm Exploration Overview

This section describes the development and assessment of boundary-flow smoothing algorithms and evaluates their impact on boundary conditions used in the modelling workflow. The work involved identifying suitable smoothing approaches, defining a robust testing methodology, and assessing algorithm performance against a set of agreed smoothing assessment metrics. The resulting analysis provides a quantitative view of how smoothing affects boundary flows.

The work is structured as follows:

1. **Literature review:** Identification of candidate smoothing algorithms and rationale for their inclusion in the initial screening.
2. **Assessment methodology:** Description of the testing framework used to compare algorithm performance.
3. **Algorithm evaluation:** Results from the first stage of testing, based on a single boundary.
4. **Algorithm optimisation:** Optimisation of the preferred algorithm for applying to boundary flows.

Full List of Smoothing Approaches

Established approaches were reviewed for boundary flow conditioning and this identified a number of options that cover both frequency filtering and amplitude limitation.

- **Frequency-focused** methods reduce variability (volatility/variance) by attenuating high frequency components of a signal. Applying this approach will produce a smoother flow, enabling more accurate prediction.
- **Amplitude-focused** methods act to limit the magnitude of abrupt changes or reduce the duration of fluctuations that are above the boundary flow limit.

Selection criteria were: demonstrated use in the literature, parameter interpretability, computational efficiency, and relevance to network operation. Table 7 highlights all algorithms that were initially considered following the literature review.

Table 7: Smoothing algorithms found through literature review

Algorithm	**Core Idea**	**How the algorithm may help**	**Primary Focus**	**Reference**
Simple Moving Average (SMA)	The SMA replaces each data point with the average of the most recent window of measurements. By pooling several samples into one value, short and rapid fluctuations are removed and the signal becomes steadier. It quickly removes sudden power fluctuations that can go above the boundary limit.	Quickly removes volatile behaviour that tends to cause brief boundary breaches.	**Frequency**	[18] [19]
Exponential Moving Average (EWMA)	The EWMA updates the flow values by blending the latest measurement with the previous smoothed value using a fixed weight. Recent data are given more influence than older data, so the filter responds faster than an SMA with similar smoothing.	Reduces fast fluctuations while keeping the smoothing delay small, which helps maintain boundary compliance without over-constraining.	**Frequency**	[18] [19]
Savitzky–Golay (SG)	The SG filter fits a simple polynomial to the most recent block of data and uses that fitted curve to estimate the current value. Instead of merely averaging, it preserves the local shape of the signal. Therefore, ramps and peaks that reflect real changes are kept, while random noise is reduced.	Smooths measurement noise but retains genuine ramps, which avoids unnecessary constraint when the underlying flow is truly changing. This helps it track up to the boundary more faithfully while still damping the jitter that causes brief exceedances.	**Frequency**	[20]
Gaussian Kernel Smoothing	The signal is convolved with a bell-shaped (Gaussian) set of weights so that nearby samples contribute most and distant samples contribute very little.	Gives a steady signal with little variance.	**Frequency**	[21] [22]

Algorithm	**Core Idea**	**How the algorithm may help**	**Primary Focus**	**Reference**
FFT Low-Pass Filter	FFT (Fast Fourier Transform) converts the time series into frequency components. Apply a low-pass filter (keep components below a chosen cutoff, suppress higher ones) then inverse-transform to get a smoother time signal.	Effectively removes the noise that causes short breaches. Needs careful setup to avoid overshoot when turned back into a time signal.	**Frequency**	[23]
Ramp-Rate Control (Slew Limiter)	Caps the permitted change between consecutive values, enforcing a maximum ΔP (change in power) per time step regardless of measurement volatility.	Puts a limit on how much the power flow can change. This directly helps to stay within the boundary constraint limits, especially during sudden changes.	**Amplitude**	[19]
Holt Double-Exponential	It keeps track of two things at the same time: the current value and the direction of travel (up or down). If the flow starts trending up, the smoothed line moves up with it; if the trend fades, it eases off. This means it reacts to real ramps while still cutting random noise.	When the flow rises or falls for a while (e.g., a long ramp), it follows that change sooner than a simple one-step smoother. This is more effective at staying within the boundary constraint limits with less delay and avoids over-constraining during sustained changes.	**Frequency**	

Algorithm	**Core Idea**	**How the algorithm may help**	**Primary Focus**	**Reference**
LOWESS / LOESS	LOWESS performs a local regression around each point, weighting nearby samples more and down-weighting outliers. Therefore, the amount of smoothing adapts to the current behaviour of the series.	Adapts to volatile versus calm periods, allowing the power to run close to the limit with fewer small breaches.	**Frequency**	
Rolling-Mean Peak Shaving	Computes a trailing rolling mean over a chosen window and uses that value as a maximum value. Whenever the instantaneous flow exceeds the dynamic maximum, the flow is clipped back toward the rolling mean (with storage absorbing the excess). The dynamic maximum updates continuously as new data enter the window, so it follows slow trends while suppressing short spikes.	It limits large excursions without needing a fixed external limit, reducing amplitude-driven breaches. Because the cap lags the signal, it naturally provides headroom during volatile periods and prevents brief spikes from crossing boundaries.	**Amplitude**	[24] [25]
Rolling-Percentile Peak Shaving	The controller computes a trailing percentile (e.g., 90th or 95th) over a chosen window and uses that as the dynamic maximum. Values above the percentile are clipped, while typical values below the dynamic maximum pass through unaltered. This makes the cap robust to outliers and better targeted at true peaks.	It trims only the most extreme peaks, reducing boundary exceedances with less overall constraint than a rolling mean. Tuning the percentile directly controls how aggressive the shaving is, balancing compliance versus energy moved.	**Amplitude**	[24] [25]

Algorithm	**Core Idea**	**How the algorithm may help**	**Primary Focus**	**Reference**
Limit-based Peak Shaving (100 MWh capacity)	A binary charge/discharge rule relative to the fixed boundary using a finite-capacity device (usable energy ≈ 100 MWh). When measured flow is above the limit, the controller charges at the available power (until state of charge (SOC)/full or limit met) to pull net flow down. When flow is below the limit, it discharges at the available power (until SOC/empty or limit met) to lift net flow up.	Allows understanding of boundary compliance with realistic device constraints.	**Amplitude**	[24] [25]
Limit-based Peak Shaving (no capacity limit)	Same binary charge/discharge rule relative to the boundary, but modelled with effectively unlimited energy capacity.	Serves as an upper-bound benchmark for boundary limit compliance. Shows best-case shaving if energy capacity were not a constraint. Useful for capacity sizing and cost–benefit studies.	**Amplitude**	

Qualitative Assessment

Before moving to quantitative testing, a qualitative screening process was undertaken to rule out methods unlikely to add value. Each algorithm was assessed for:

1. How many parameters need tuning and how sensitive results are to them.
2. Redundancy with other methods that already cover the same behaviour.
3. Whether the extra complexity would deliver meaningful metric improvements.

Algorithms that did not pass the screening were not progressed. Table 8 summarises which algorithms were not passed through to the next stage and why this was the case.

Table 8: Smoothing algorithms discontinued through qualitative screening

Algorithm	Core Idea Summary	Main Reference	Reason for not taking forward	Further explanation
Savitzky–Golay	Local polynomial fit over a sliding window.	[4]	Too many parameters or parameters that are difficult to tune.	Requires several choices (window size, polynomial order) and known to produce distortions at the start and end of the data, which adds special rules and delay. The extra complexity does not bring clear benefits for the objectives.
Ramp-Rate Control	Enforce a maximum power change per time step.	[2]	Too many parameters or parameters that are difficult to tune.	Choosing one safe ramp rate for all conditions is challenging and unlikely to be beneficial.
Holt Double-Exponential	Level + trend exponential smoothing.	—	Redundancy with other methods.	Similar to EWMA but requires more tuning of parameters.
Rolling-Mean Peak Shaving	Clip to the mean of a past window.	[8] [9]	The added complexity would not deliver meaningful metric improvements.	The true limit is already known. Using a rolling mean as the cap targets the wrong value and causes avoidable charging/discharging (wasted energy) without better limit compliance.
Rolling-Percentile Peak Shaving	Clip to a high (or low) percentile.	[8] [9]	The added complexity would not deliver meaningful metric improvements.	Same issue as the rolling mean: it caps against a moving statistic, not the known limit. This leads to unnecessary actions and wasted energy without improving performance against the actual limit.

Candidate Smoothing Approaches Evaluated

The down-selected set of smoothing approaches covers both frequency filtering and amplitude limitation.

Frequency Filtering Smoothing Approaches

Simple Moving Average (SMA): Computes the mean over a trailing time window. It reduces short-term fluctuations but introduces a delay of about half the window length. A longer window smooths more but increases delay. SMA is a common baseline in PV smoothing studies with storage [18] [19]. Principal parameter: window length.

Exponentially Weighted Moving Average (EWMA): Updates the smoothed value by blending the latest measurement with the previous smoothed value. For a given level of smoothing it usually gives less delay than a simple moving average and runs efficiently in real time. EWMA has been tested on the same photovoltaic (PV) data and ramp rate limits as simple moving average and low pass filters, allowing direct comparison of delay, smoothness, and storage use [18]. When short-term forecasts are added to the controller, asset cycling is lower than with fixed filter settings alone [19]. Principal parameter: smoothing factor α or an equivalent time constant.

LOWESS or LOESS: A local regression smoother that fits a simple curve to nearby data and moves along the series. It follows gradual changes and aims to preserve the overall trend while reducing noise. A robust option can reduce the influence of outliers. It is more computationally demanding than SMA or EWMA and needs a sensible span setting; behaviour near the start and end of the series can be sensitive to that choice. Principal parameter: span or fraction of data used in each local fit.

Gaussian kernel smoothing: A moving average that uses bell-shaped weights, so nearby points count more and distant points count less. It produces a very smooth result with few ripples. In the literature, Gaussian smoothing is applied to renewable generation with explicit power constraints and is shown to achieve target smoothness with smaller required storage than simple moving averages [20]. The algorithm is also used as a benchmark alongside moving average, low pass, and Savitzky-Golay filters in PV and BESS studies [21]. Principal parameter: kernel standard deviation or full width at half maximum.

Low-pass filter: First and second order low-pass filters are widely used to shape the bandwidth of power fluctuations and to generate storage set points. They are implemented and compared in PV smoothing studies and also appear within hybrid storage control strategies for wind generation smoothing [23]. Principal parameters: time constant for first order and natural frequency with damping ratio for second order.

Amplitude Constraining Smoothing Approaches

Limit-based peak shaving with 100 MWh maximum: When the flow is above the constraint limit, the asset charges and when below it the asset discharges, while enforcing an energy capacity of 100 MWh. This directly targets exceedance reduction and turns smoothing into an operational dispatch under an energy budget. In the literature, Rahimi sets fixed limits and demonstrates a BESS controller with explicit capacity and state of charge management for peak reduction [24]. Groß et al. apply a similar concept in industry but set the limit using probabilistic quantile forecasts to balance missed peaks against cycling and cost [25]. Principal parameters: energy capacity.

Limit-based peak shaving without an explicit maximum: The same control action is applied without an energy cap, which removes the saturation constraint and provides a theoretical reference. Principal parameter: limit profile.

Assessment Methodology

The seven algorithms identified in the preceding section were taken forward for assessment. The following steps describe the assessment methodology (each step corresponds to the flowchart numbering shown in Figure 15):

1. **Constructing the adjusted flow**
   Using metered data together with bid and offer data from Palantir, an adjusted flow is constructed by unwinding bids and offers. This provides a clearer picture of the flow before any balancing mechanism actions were taken. The detailed methodology for this step is described in the Adjusted Flow Construction section.
2. **Identify constrained periods**
   A constrained period has been defined as any period in which the adjusted flow exceeds the boundary limit. To merge short gaps and reflect settlement behaviour, a 30-minute buffer on either side is added when joining adjacent periods.
3. **Smoothing metered flow**
   The smoothing algorithms are applied to the metered flow, but only within the constrained periods identified in step 2. This assumes the flow can be altered to perfectly match the algorithms output. Due to a time lag in assets acquiring actions there would be a discrepancy between the algorithm's output and the actual smoothed flow[3].
4. **Metric evaluation**
   The smoothing assessment metrics are used to determine the impact of smoothing by comparing the smoothed flow against the metered flow.

---

[3] The algorithms have been considered under the assumption that they can instantly affect the actual flow to the smoothed flow value. In network deployment this would not be possible due to time lags in sending signals to assets and then those assets responding. It is likely that assets would need a data feed of the metered flow and the smoothed flow point, they would then scale their output to ensure the metered flow moved towards the smoothed flow value.

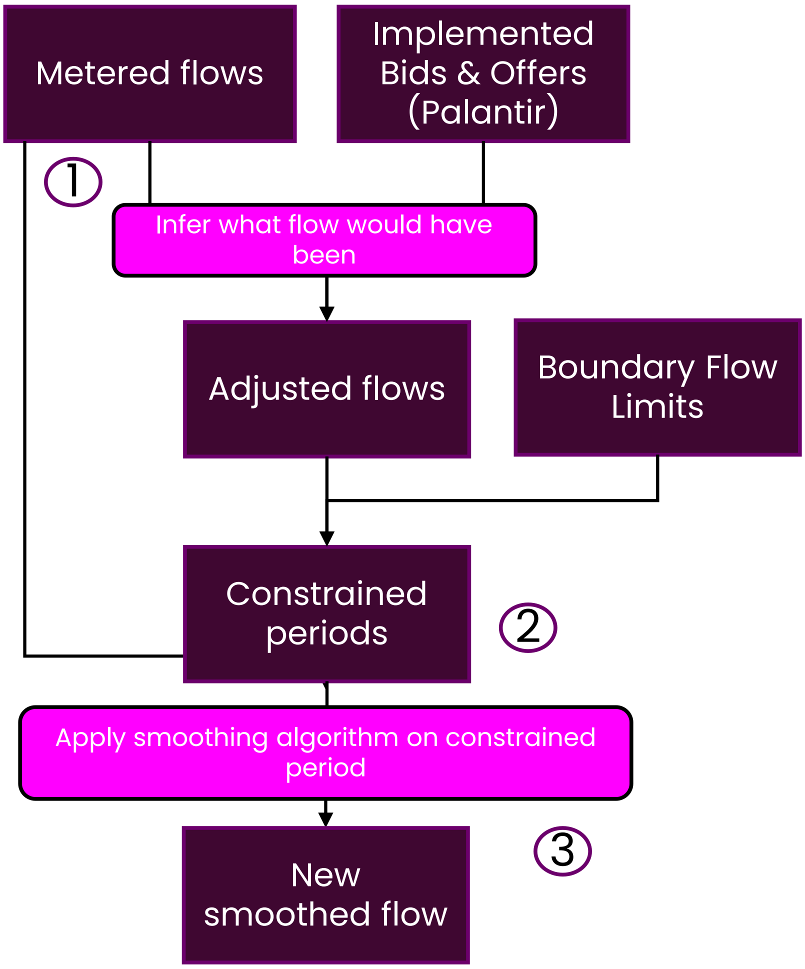

Figure 15: Assessment methodology workflow.

Smoothing Assessment Metrics

Three sets of smoothing assessment metrics have been developed that assess the success of smoothing from different perspectives. They are used to compare the success of the smoothing algorithms and to decide which should progress further based on the results. The three sets cover:

1. Statistical Metrics
2. Constraint Limit Exceedance Metrics
3. Asset Metrics

The three sets are used to assess the impact on smoothing individually and then, to conclude, all the assessment metrics are evaluated together to understand a complete picture on the success and impact of smoothing.

Statistical Metrics

First, the statistical metrics are used to understand how smooth the flow is from a mathematical perspective. Table 9 summarises the statistical metrics used.

Table 9: Statistical Smoothing Assessment Metrics

Metric	Variability or Volatility	Description	Reason to measure
90th Percentile Variation	Variability	The 90th percentile of power within each time window. It is the power level exceeded for the top 10% of the duration of a time window. This metric measures the mean variation in the 90th percentile power between consecutive windows.	A drop in this metric indicates reduced amplitude fluctuations.
Auto Correlation	Volatility	Measures how closely power values at one time point relate to previous values, indicating stability or volatility over time.	Assesses how predictable or persistent flow patterns are, helping evaluate the effectiveness of smoothing.
Smoothness Index (Adams & Van Deventer)	Volatility	A mathematical measure of smoothness defined by curvature.	Provides an objective measure of smoothness, helping assess if smoothness is desired when combined with other metrics.
Standard Deviation of Centred Moving Average Residuals	Both	The residual differences between observed values and the centred moving average.	Compares smoothed algorithms to a ‘smooth’ flow that has been created using future values.

Constraint Limit Exceedance Metrics

Secondly, Control Room feedback showed that it was necessary to characterise exceedances and evaluate whether smoothing had reduced their number and extent. A well-performing smoothing algorithm should also help reduce the number, scale and magnitude of exceedances. Table 10 summarises the constraint limit exceedance metrics used.

Table 10: Constraint Limit Exceedance Assessment Metrics

Metric	Description	Reason to measure
Constraint Limit Exceedance Number	The number of 10-minute periods where a constraint limit breach occurs.	Captures whether constraint limit exceedances have been reduced.
Constraint Limit Exceedance Energy	The average energy transferred above the constraint limit during a 10-minute period where a breach has occurred.	Captures the change in magnitude of constraint limit exceedances.
Constraint Limit Exceedance Time	The average time spent above the constraint limit during a 10-minute period where a breach has occurred.	Captures whether the time spent above the constraint limit has changed.

Asset Metrics

Thirdly, smoothing must be assessed in terms of asset requirements, as the service must be commercially viable. Any extra energy or power needed to smooth flows must be justified by the benefits gained through reduced constraint actions.

Table 11: Asset Smoothing Assessment Metrics

Metric	**Description**	**Reason to measure**
Peak Energy	The absolute peak of import or export energy (MWh) required for the assessed period.	Captures the overall energy exchanged, providing insight into the amount of energy needed to provide the smoothing.
Cumulative Energy	Considering both import and export power, what is the largest cumulative energy (MWh) required to undertake smoothing during the assessed period.	Offers a clearer indication of the type and size of asset that may be required.
Peak Power	The absolute peak MW import or export value for the assessed period.	Captures the power requirements needed to provide smoothing as an indication for the asset requirements.

Algorithm Assessment

Candidate smoothing algorithms were assessed using the sets of metrics defined above. The assessment took the following approach:

1. The effect of smoothing was tested on a single boundary. SSE-SP2 was selected as the primary test case, as this is the most constrained boundary from the adjusted flow data: 161 days of time above the limit (35.4% of the period 01/01/24 to 31/03/25).
2. The algorithms were checked to ensure they ran quickly enough to operate in real time, and they were assessed against the chosen performance metrics.
3. The close results meant that a sensitivity analysis was needed. Not all metrics are necessarily worth an equal importance weighting for NESO, so a repeated randomised assignment of sensitivities was undertaken to ensure robustness of results when choosing the preferred algorithm.
4. Conclusions were drawn from the previous steps on the relative strengths and weaknesses of the algorithms to choose a preferred algorithm.

Initial Algorithm Assessment Results

Table 12 shows the results for each of the assessment metrics. Each column is colour-coded with the best result in green and the worst result in red.

There are three groups of columns in the table, corresponding to the three different classes of metric: smoothing metrics, exceedance metrics and asset metrics.

Table 12: Algorithm assessment results

Algorithm	**Smoothness Index**	**90th Percentile Variation (MW)**	**Std Dev of Centred Moving Ave Residuals**	**Auto Correlation**	**Constraint Limit Exceedance Number**	**Constraint Limit Exceedance Energy (MWh)**	**Constraint Limit Exceedance Time (min)**	**Peak Energy (MWh)**	**Cumulative Energy (MWh)**	**Peak Power (MW)**
Simple Moving Average	67,972,336	1992.46	17.98	0.99	8960.00	331.95	10.00	24,596.77	35.62	2040.61
Exponential Weighted Moving Average	63,373,792	1978.29	9.96	1.00	6986.00	328.45	10.00	132,344.96	294.94	1929.39
Frequency Filter Smoothing	59,915,716	1992.32	17.69	0.99	8705.00	331.60	10.00	32,025.28	42.53	1997.59
Gaussian Smoothing	64,190,152	1985.62	9.96	1.00	8072.00	330.99	10.00	38,608.74	70.96	1963.84
LOWESS Smoothing	66,365,084	1982.36	8.93	1.00	7532.00	330.10	10.00	57,715.88	499.69	1779.10
Peak Shaving With Limit (100MWh)	311,963,232	1988.35	19.51	0.99	2176.00	326.27	10.00	44,785.81	100.00	1413.48
Peak Shaving With No Limit	299,729,312	1984.06	18.57	0.99	0.00	0.00	0.00	84,088.39	2445.93	1413.48
Unsmoothed Flow	329,283,072	1994.72	20.71	0.99	9451.00	331.86	10.00	0.00	0.00	0.00

Normalised Algorithm Assessment Results

To allow for easy comparison of results and to combine the assessment metrics into a single score, each column was normalised between 0 and 1. A single score was produced for each algorithm by calculating a weighted sum of the metrics. A lower number indicates a better score. Table 13 shows the normalised assessment metrics for each algorithm. Combining these metrics with a uniform weighting produces the scores in the final column.

Table 13: Normalised algorithm assessment results

Algorithm	Smoothness Index	90th Percentile Variation (MW)	Std Dev of Centred Moving Ave Residuals	Auto Correlation	Constraint Limit Exceedance Number	Constraint Limit Exceedance Energy (MWh)	Constraint Limit Exceedance Time (min)	Peak Energy (MWh)	Cumulative Energy (MWh)	Peak Power (MW)	Total
Simple Moving Average	0.03	0.86	0.77	1.00	0.95	1.00	1.00	0.19	0.01	1.00	0.68
Exponential Weighted Moving Average	0.01	0.00	0.09	0.00	0.74	0.99	1.00	1.00	0.12	0.95	0.49
Frequency Filter Smoothing	0.00	0.85	0.74	1.00	0.92	1.00	1.00	0.24	0.02	0.98	0.68
Gaussian Smoothing	0.02	0.45	0.09	0.00	0.85	1.00	1.00	0.29	0.03	0.96	0.47
LOWESS Smoothing	0.02	0.25	0.00	0.00	0.80	0.99	1.00	0.44	0.20	0.87	0.46
Peak Shaving With Limit (100MWh)	0.94	0.61	0.90	1.00	0.23	0.98	1.00	0.34	0.04	0.69	0.67
Peak Shaving With No Limit	0.89	0.35	0.82	1.00	0.00	0.00	0.00	0.64	1.00	0.69	0.54
Unsmoothed Flow	1.00	1.00	1.00	1.00	1.00	1.00	1.00	0.00	0.00	0.00	0.70

Three algorithms perform similarly well: the Exponentially Weighted Moving Average, Gaussian Smoothing and LOWESS Smoothing. As an initial assumption, all metrics have been weighted equally as they are all deemed to represent important aspects of smoothing. Due to these results, it is difficult to say which of these algorithms has performed best overall at this stage, as a change in weightings may affect the final score. Therefore, the next step taken was to assess which algorithm performs the best with a sensitivity analysis conducted on the weightings of the assessment metrics.

Sensitivity Analysis Results

The sensitivity analysis involved taking the normalised assessment metrics and assigning each one a random weighting between 0 and 1. The total score is then calculated by summing these weighted metrics. This is done 100,000 times with different weights each time and a distribution of the total scores calculated. This allows a comparison of the various algorithms as the weightings are varied. By looking at the distribution, it is possible to see how algorithms perform with various sets of weightings and the algorithm performance can be compared without any prior knowledge of what the relative importance weightings should be.

Each of the assessment metric sets had an individual sensitivity analysis performed on them before all ten were assessed together. The following sections cover the sensitivity analyses on:

1. Statistical metrics
2. Constraint limit exceedance metrics
3. Asset metrics
4. All metrics

Statistical Smoothing Metrics Sensitivity Analysis

When considering the assessment metrics that relate to the statistical efficacy of the smoothing algorithms, EWMA is clearly the best algorithm. This is shown by the tight distribution of the scores on the box plot in Figure 16. This shows that, in almost all cases, the EWMA has the best total score, leading to the conclusion that the weightings used are largely irrelevant as the EWMA will score the best across most prioritisation options. There are some extreme cases where the Gaussian or LOWESS smoothing will outperform but the likelihood of these weightings aligning with the desired weightings are low. The assessment metrics considered for this analysis were:

1. 90th percentile variation
2. Autocorrelation
3. Smoothness index (Adams & Van Deventer)
4. Standard deviation of centred moving average residuals

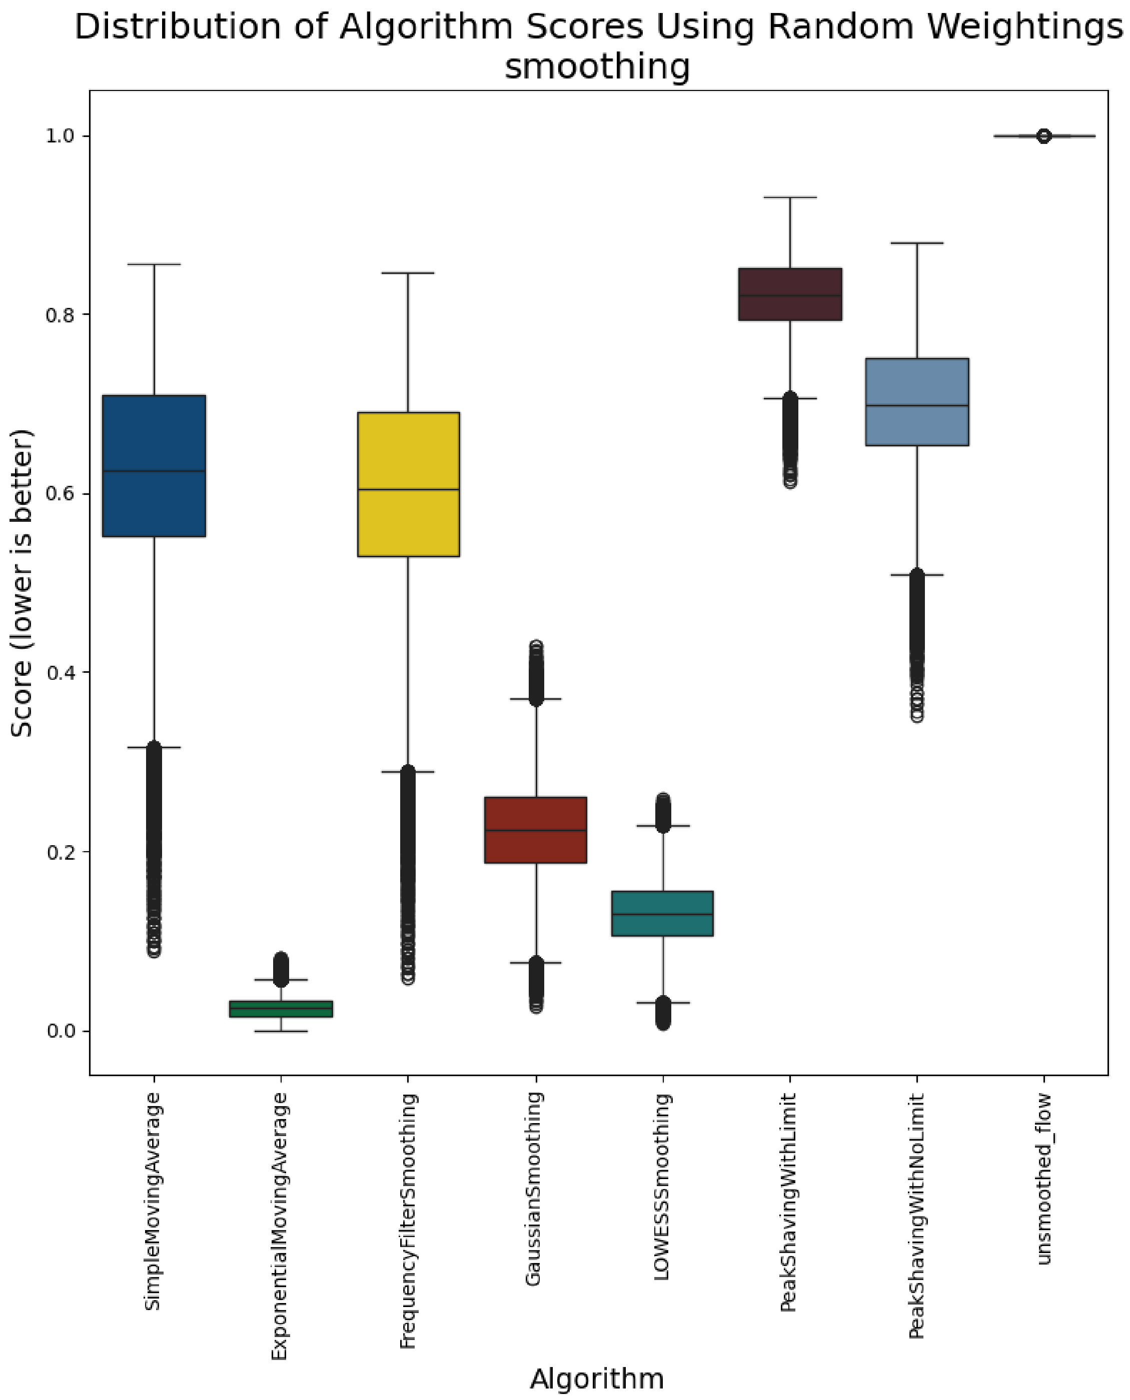

Figure 16: Statistical smoothing metrics sensitivity analysis box plot.

Constraint Limit Exceedance Smoothing Metrics Sensitivity Analysis

When considering the metrics that relate to reducing constraint exceedance, the weightings applied to sum the metric scores have a strong influence on which algorithm ranks the best. This can be seen in the box plot in Figure 17, where the whiskers overlap for most of the algorithms on the plot. This suggests that the weight should be manually chosen, depending on which metric is prioritised for minimisation.

The peak shaving algorithms perform very well, due to the goal of these algorithms being to minimise exceedances rather than smoothing. Peak shaving with no limit will remove all exceedances and will therefore achieve the perfect score of 0.0 for this group of metrics. The assessment metrics considered for this analysis were:

1. Constraint limit exceedance number
2. Constraint limit exceedance energy
3. Constraint limit exceedance time

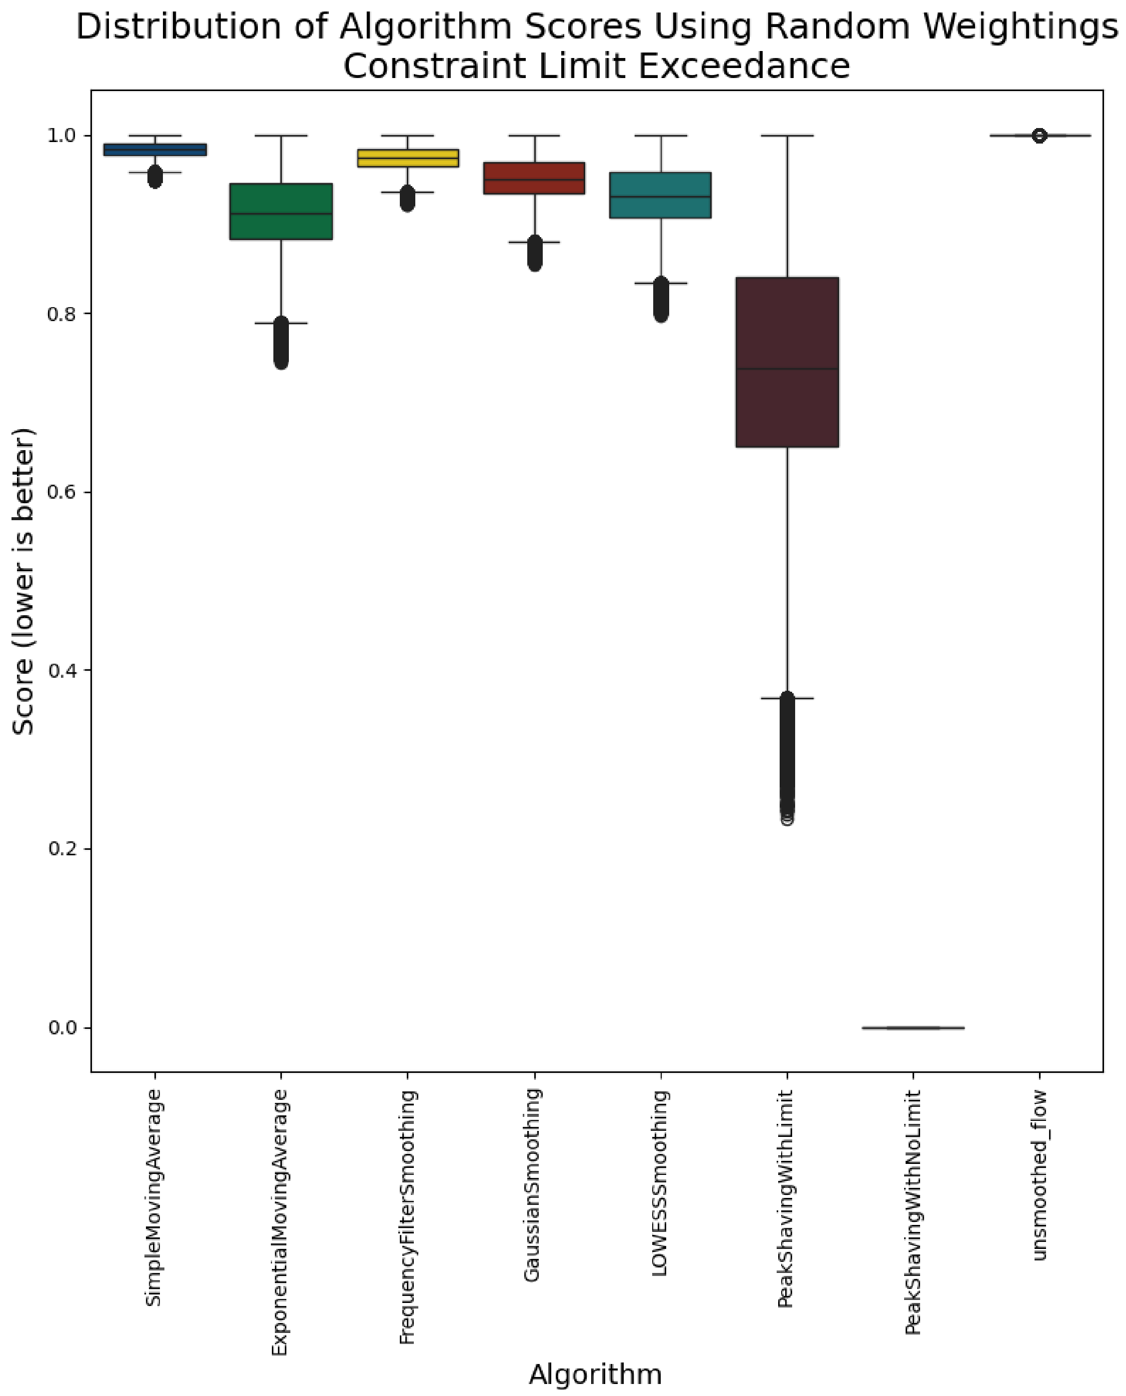

Figure 17: Constraint limit exceedance smoothing metrics sensitivity analysis box plot.

Asset Smoothing Metrics Sensitivity Analysis

When looking at the demands of the smoothing algorithm on the assets providing that smoothing, most of the algorithms overlap in their scores as shown in Figure 18. EWMA does score worse than some other algorithms, but the width of the distributions makes it hard to choose a clear winner.

The peak shaving algorithms are interesting here as the ‘peak shaving with limit’ is limited in the power that can be used for smoothing, which will restrict the demands on the smoothing asset, giving it a better score. The peak shaving with no limit scores poorly as all the energy that flows over the limit must be absorbed, giving it high scores for the asset metrics.

As would be expected, the unsmoothed flow scores perfectly against these assessment criteria as the smoothing asset is not used. The assessment metrics considered for this analysis were:

1. Peak energy
2. Cumulative energy
3. Peak power

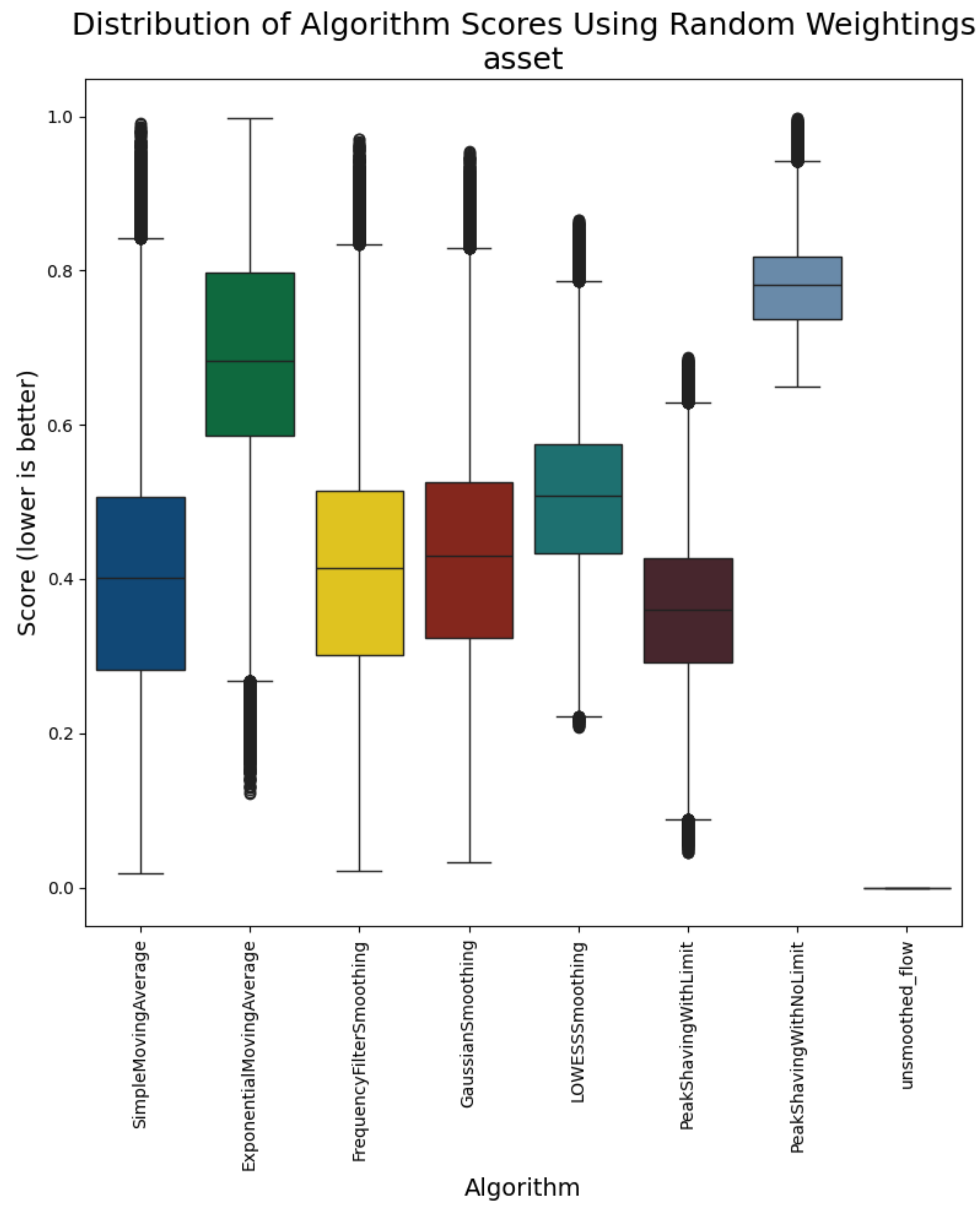

Figure 18: Asset smoothing metrics sensitivity analysis box plot

All Smoothing Metrics Sensitivity Analysis

Combining all ten metrics together allows for a full sensitivity analysis to be performed. The results of this have a wide distribution, suggesting that the relative importance weighting between the asset usage and the smoothing metrics will be important. The EWMA smoothing algorithm performs the best on average with Gaussian smoothing and LOWESS smoothing scoring similarly. These three algorithms were also close in the results of the normalised assessment. The sensitivity analysis is shown in Figure 19.

Based on these results, EWMA will be used for the subsequent analyses. If the importance of the assessment metrics differed from being equally important then the preferred algorithm may differ. For the purposes of this project, the feasibility of smoothing is being investigated and this is not necessarily dependent on the smoothing algorithm. A deployed smoothing approach would need to be optimised for operational service usage based on more detailed network modelling. This can be explored at a later date dependent on the results of this project.

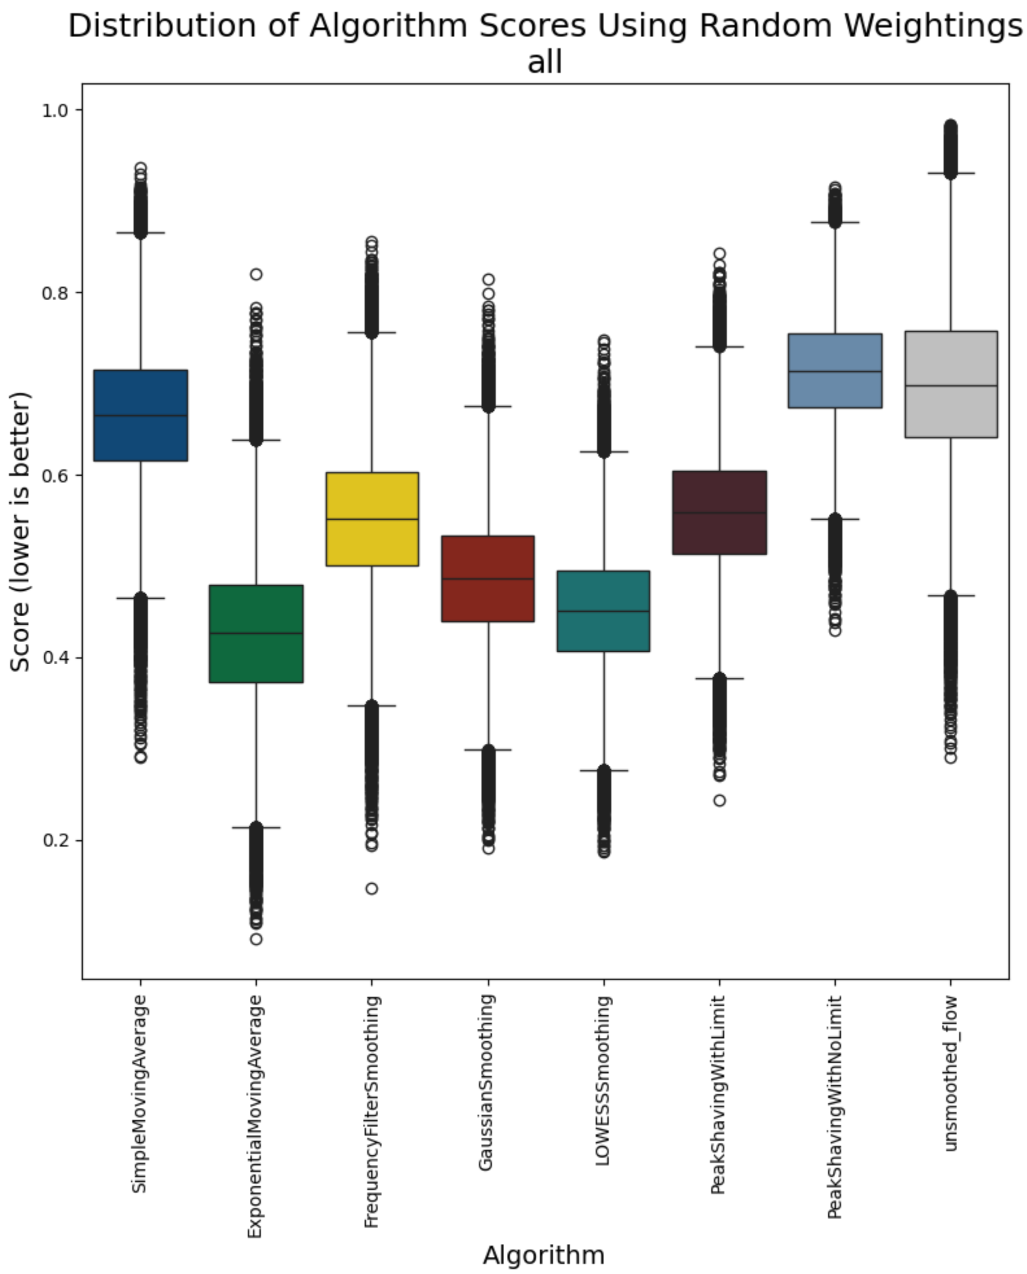

Figure 19: Sensitivity analysis of all metrics.

Appendix H – Algorithm Optimisation Results

The EWMA has a tuning parameter α that determines how quickly the smoothed flow reacts to changes in the flow. The tuning parameter is used to reduce the impact of metered flow data from further in the past by down-weighting older data points. The formulation for EWMA is:

$$s_0 = x_0$$

$$s_t = \alpha x_t + (1 - \alpha) s_{t-1}$$

Where x_0 is the first flow value, x_t is the flow value at time t, and s_t is the smoothed flow at time t.

This takes a weighted sum of the historical data and the most recent data. The smaller the value of the tuning parameter, the more weight is given to historic data, giving a smoother flow.

For our initial assessments, a tuning parameter of 0.0016 was used as it looked to visually smooth the data to a sufficient level. To ensure the most appropriate tuning parameter was selected, EWMA was run using multiple parameter values. Figure 20 shows an example of the effect of varying the smoothing factor on the resulting flow.

By comparing the assessment metrics from each of these runs it is possible to see how changing the tuning parameter affects how the smoothing algorithm performs. A sensitivity analysis of the effect of the tuning parameter on the score was performed and the results are shown in this section.

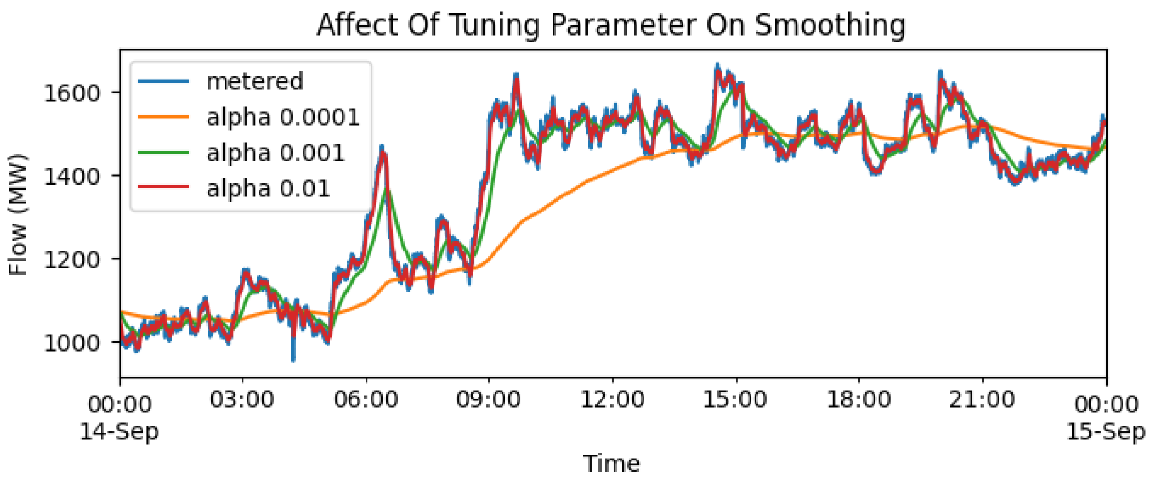

Figure 20: Tuning parameter effect on smoothing.

Statistical Smoothing Metrics Sensitivity Analysis

A sensitivity analysis shows that the score is not strongly affected by the weightings. In general, the higher the tuning parameter the worse the statistical metrics score. This is due to a higher tuning algorithm giving less weighting to historical data causing the smoothed flow to more closely resemble the unsmoothed flow. The assessment metrics considered for this analysis were:

1. 90th percentile variation
2. Autocorrelation
3. Smoothness index (Adams & Van Deventer)
4. Standard deviation of centred moving average residuals

The analysis of all the metrics is shown in Figure 21. The individual metric analyses are shown in Figure 22 . They combine to show the curve in Figure 21.

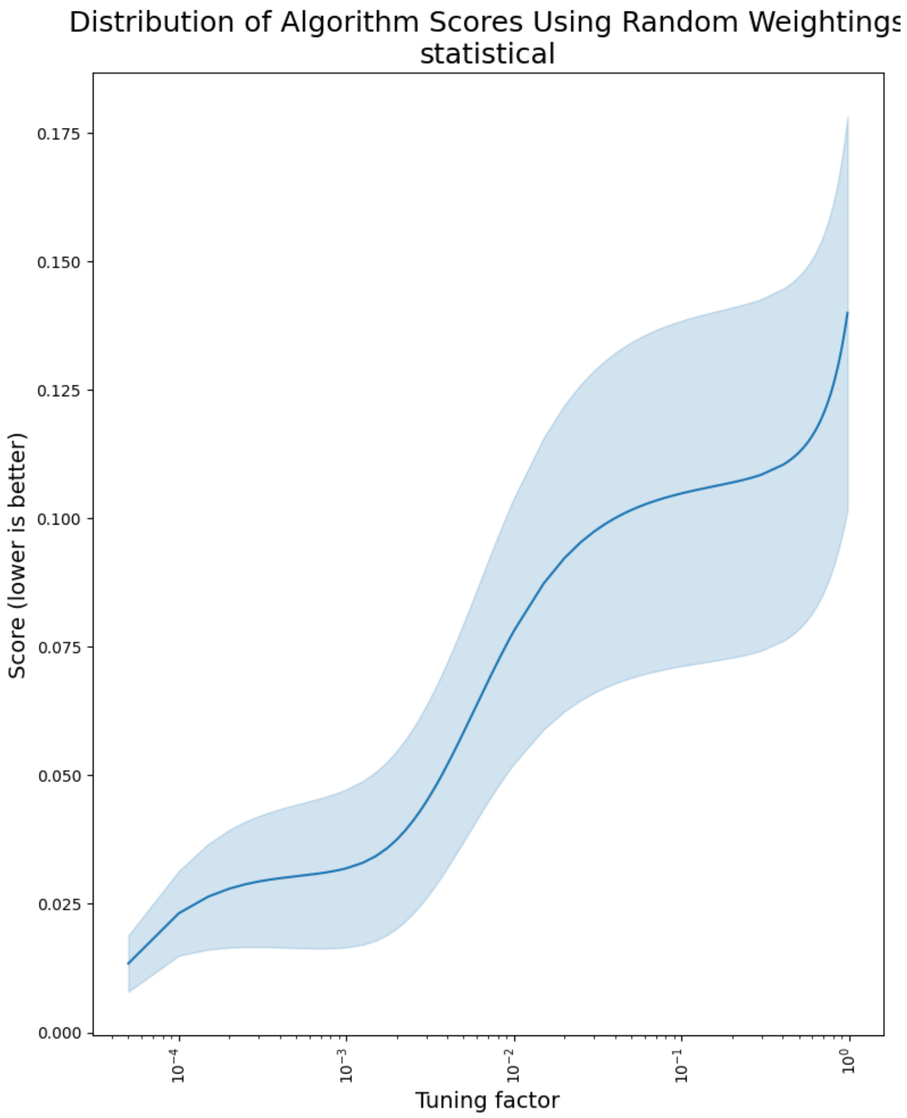

Figure 21: Statistical smoothing metrics sensitivity analysis.

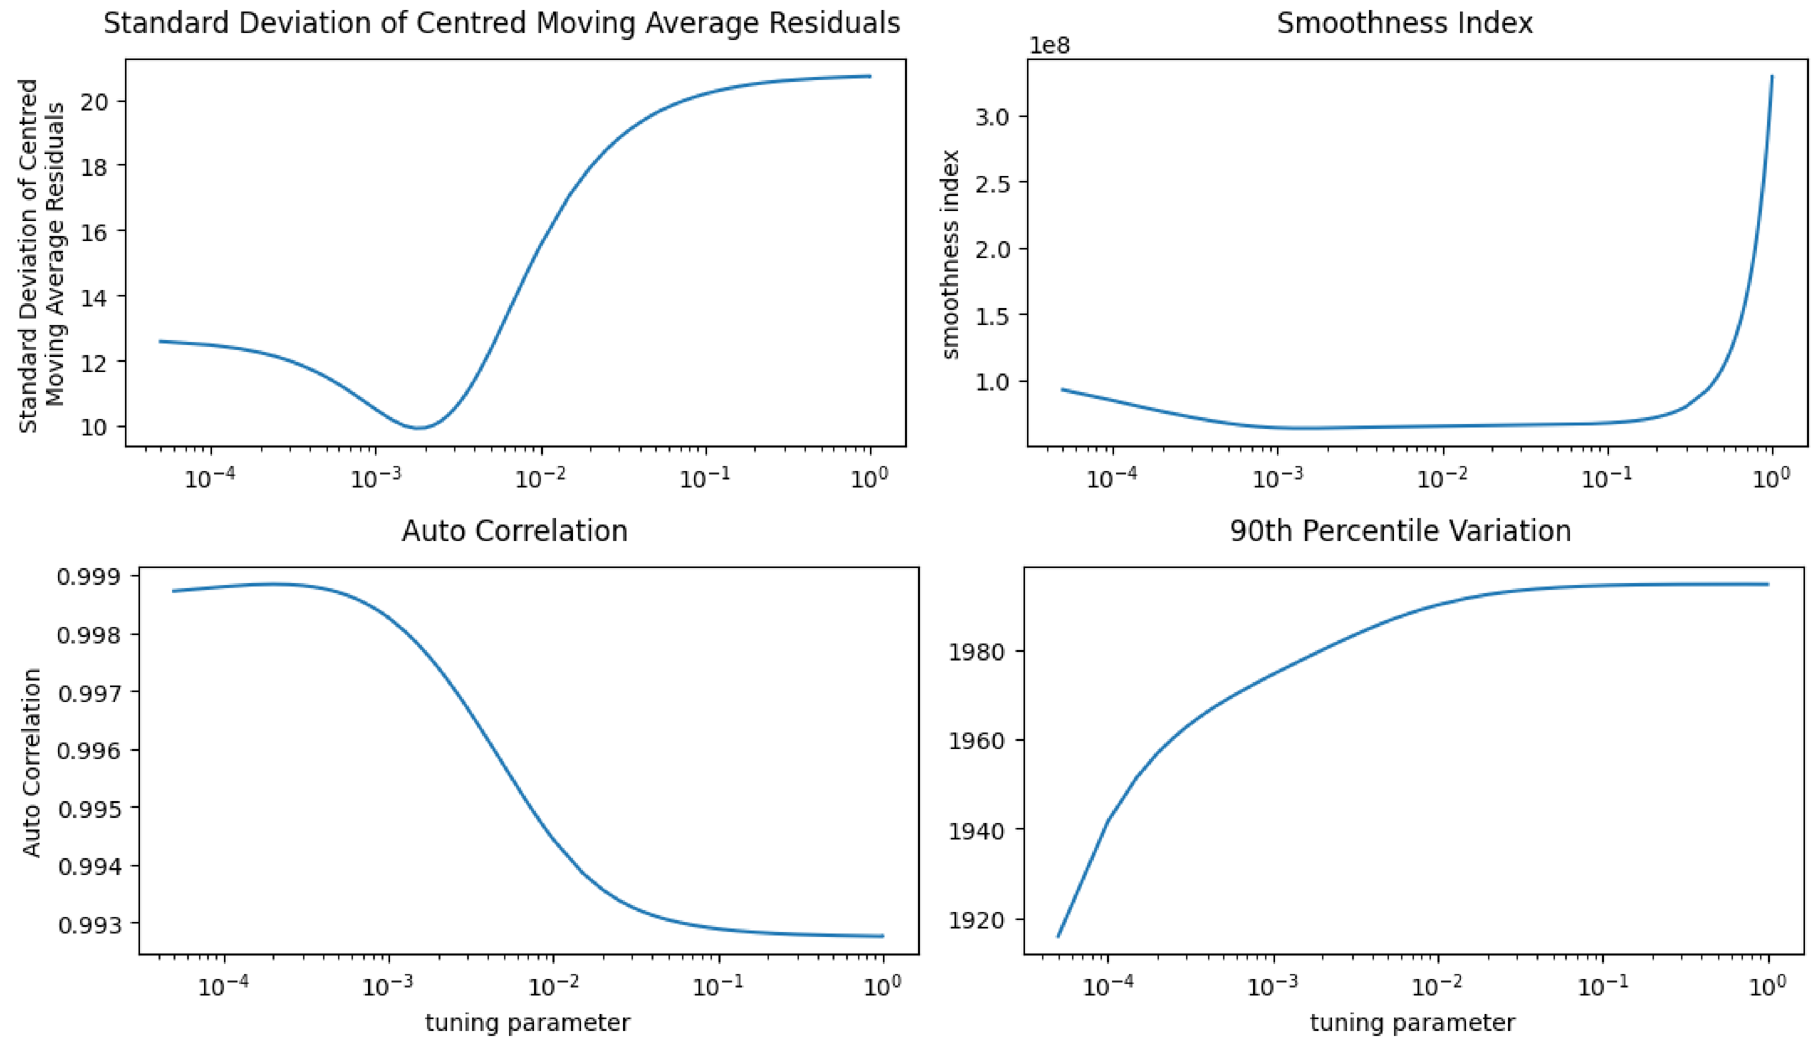

Figure 22: Statistical smoothing individual metrics sensitivity analysis.

Constraint Limit Exceedance Metrics Sensitivity Analysis

As the tuning parameter increases, the number and average size of limit exceedances worsens. This is due to the smoothed flow being more volatile. Smaller tuning parameters lead to low scores for all three metrics, indicating little sensitivity to changes in the weighting.

When combined with the previous result, this suggests that a smaller tuning parameter is almost always better. However, as seen later, there is a trade-off between smoothing efficacy and the technical requirements for the smoothing asset.

The assessment metrics considered for this analysis were:

1. Constraint limit exceedance number
2. Constraint limit exceedance energy
3. Constraint limit exceedance time

The analysis of all the metrics is shown in Figure 23. The individual metric analyses are shown in Figure 24. These are combined for the curve in Figure 23.

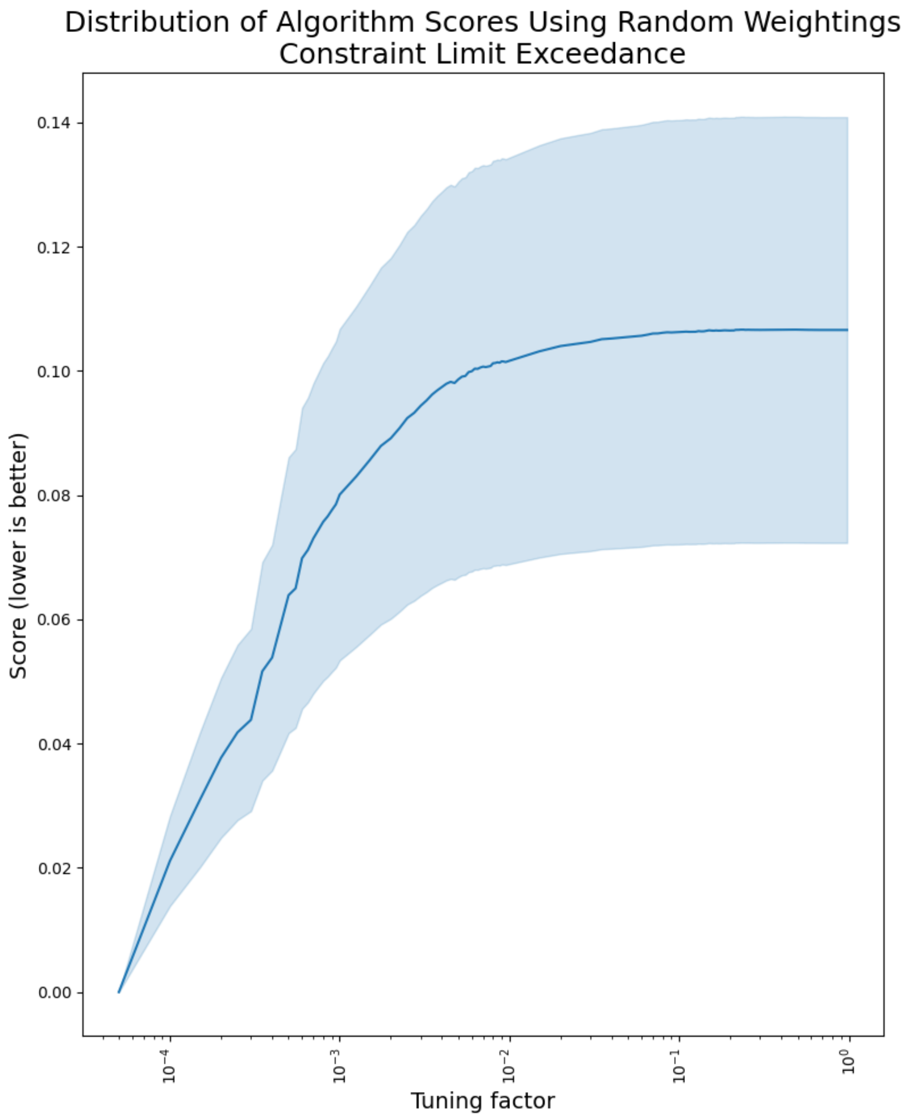

Figure 23: Constraint limit exceedance metrics sensitivity analysis.

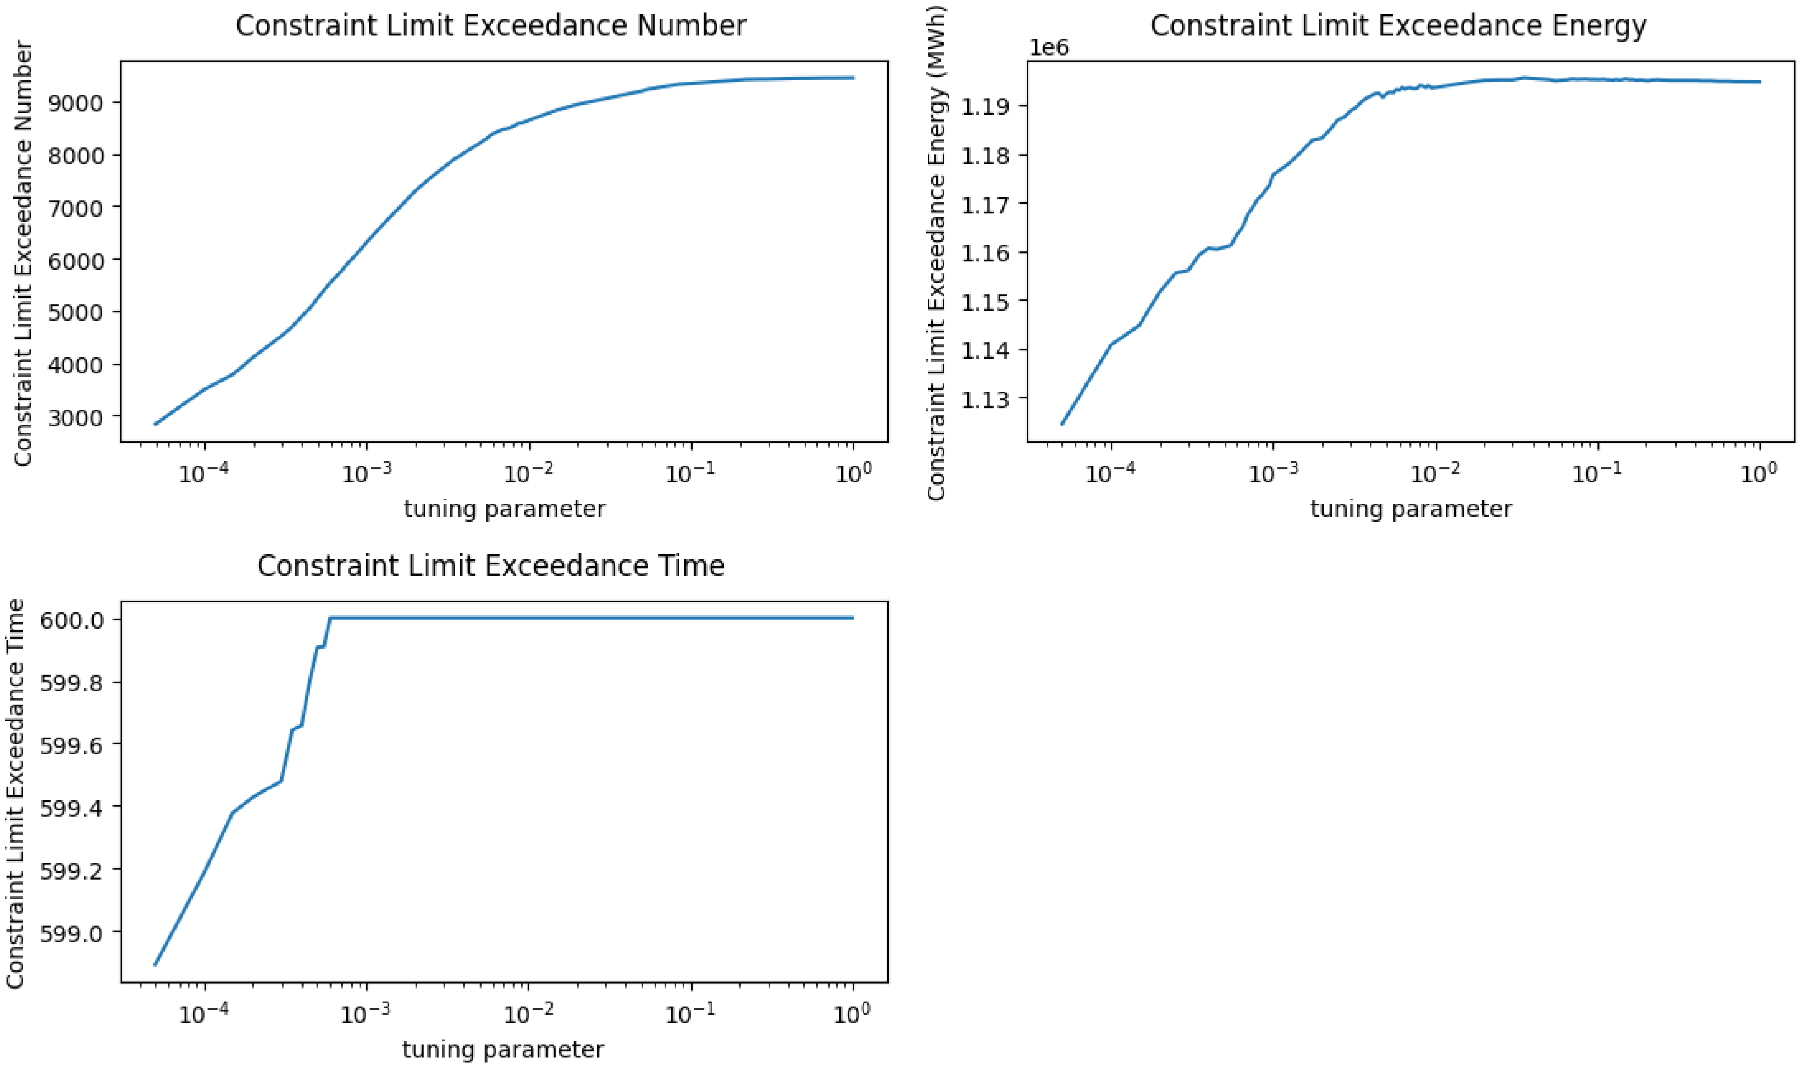

Figure 24: Constraint limit exceedance individual metrics sensitivity analysis.

Asset Smoothing Metrics - Sensitivity Analysis

The asset power and energy requirements will be important for understanding the impact of smoothing. Figure 25 shows a 24-hour sample of flow that was smoothed using the EWMA and an alpha value of 0.0016. The requirements of the assets used for smoothing can be seen in the lower two plots of Figure 25.

The middle plot shows the power flow into and out of the asset that would be needed to provide the smoothing. Rapid changes in power are needed to achieve the smoothing, although the magnitude is well-bounded over the course of 24 hours.

The bottom plot shows the cumulative energy that would be needed to achieve the smoothing. As the power demand flips from positive to negative, the asset switches between import and export. Here the asset imports more than it exports. However, this excess energy could be exported when the flow is not constrained.

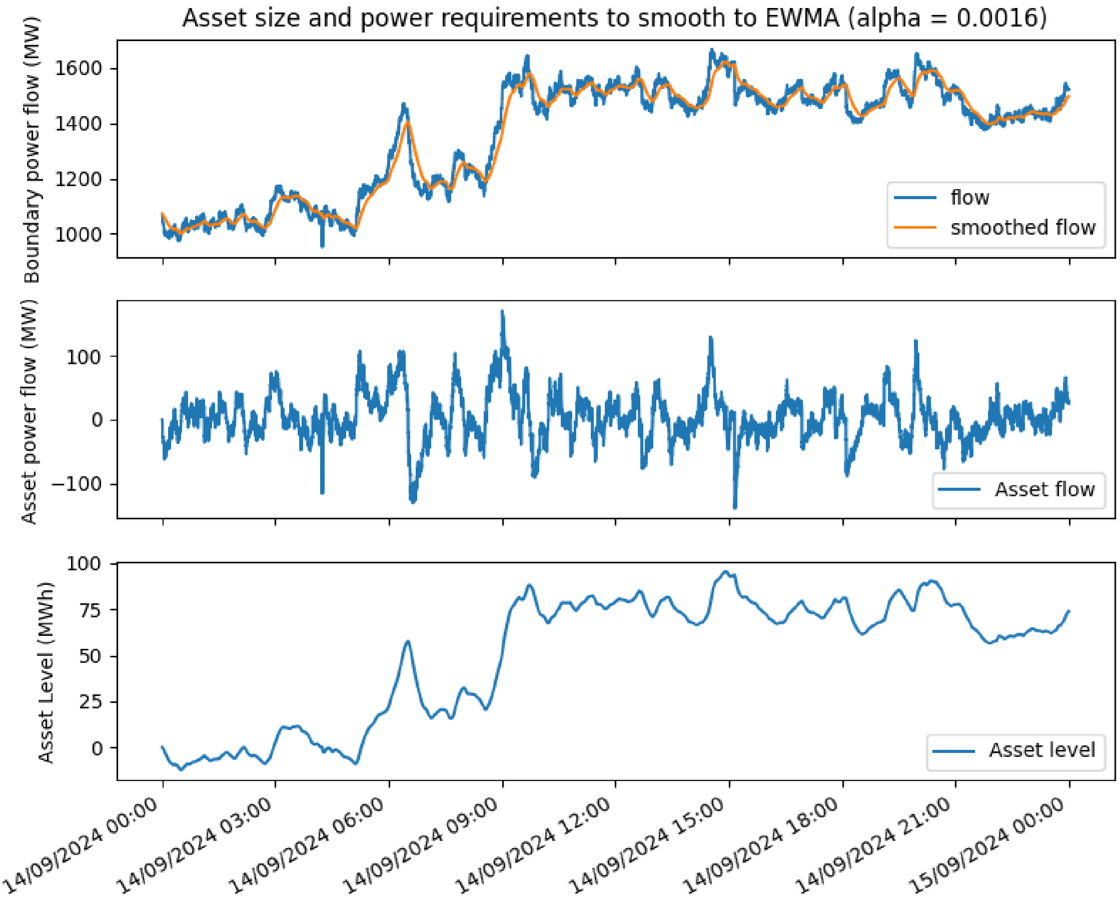

Figure 25: Asset smoothing requirements.

When a sensitivity analysis is performed with the asset metrics, as the tuning parameter increases the EWMA scores better on the asset metrics. This is to be expected as with larger tuning parameters the EWMA weights the past data less. At the extreme, a tuning parameter of 1 provides no smoothing and therefore has no asset requirements.

For values less than 0.001 the score gets significantly worse as the asset is utilised more intensively.

The assessment metrics considered for this analysis were:

1. Peak energy
2. Cumulative energy
3. Peak power

The analysis of all the metrics is shown in Figure 26. The individual metric analyses are shown in Figure 27. These are combined to give the curve in Figure 26.

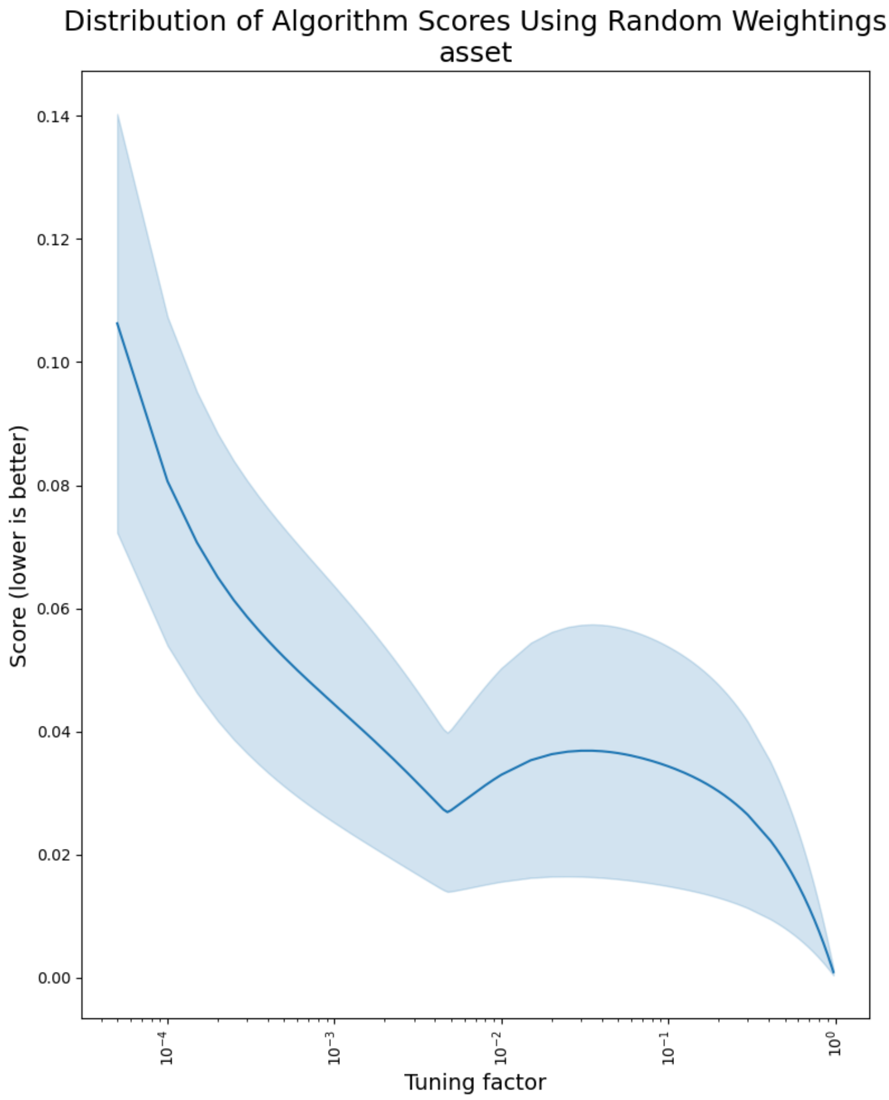

Figure 26: Asset smoothing metrics sensitivity analysis.

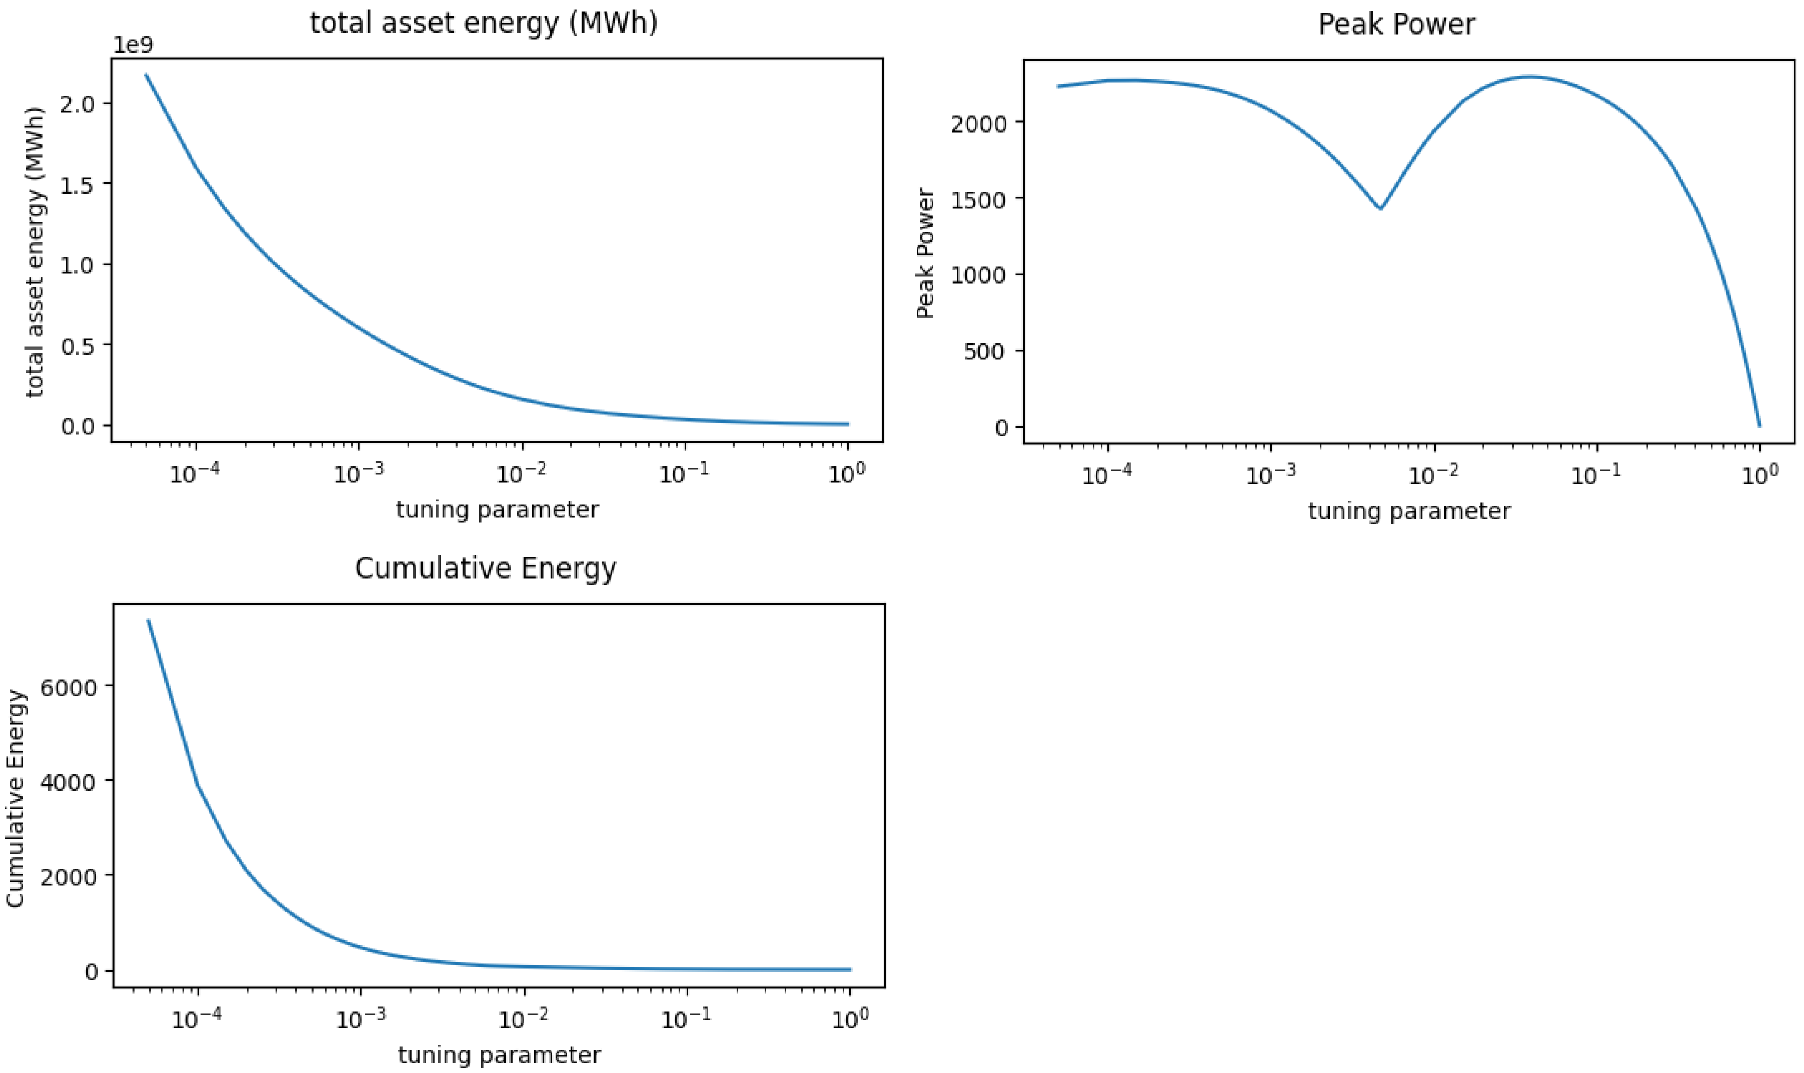

Figure 27: Asset smoothing individual metrics sensitivity analysis.

Full Sensitivity Analysis with Asset Smoothing Metrics

When combining the 10 previously discussed metrics and performing a final sensitivity analysis, the score follows an interesting shape as shown in Figure 28. For very low values of the tuning parameter the score is high. This is due to the low alpha values requiring an asset to have large power and energy storage capacities in order to deliver smoothing. At high values of the tuning parameter, the score is also high due to the EWMA delivering poor smoothing performance. For values of around 0.001, the algorithm performs the best looking at all the metrics. For values in this area the requirements of the asset and the smoothing performance appear to be most optimal.

The spread of the results does show there is a sensitivity to the weightings. Between 0.0001 and 0.01 there is a lot of overlap, so any targeted prioritisation on weightings will determine where in this band the tuning parameter should be set.

Table 14 shows the assessment metrics for varying values of the tuning parameter. The anticorrelation between smoothing performance and asset usage can be seen in the final three columns compared to the earlier columns. If one type of metric were to be prioritised, then this may change what tuning parameter is optimal. Currently a parameter has been chosen that balances the different types of metrics as they are all important.

Table 14: Assessment metric results under different tuning parameters

alpha	Smoothness Index	90th Percentile Variation (MW)	Std Dev of Centred Moving Ave Residuals	Auto Correlation	Constraint Limit Exceedance Number	Constraint Limit Exceedance Energy (MWh)	Constraint Limit Exceedance Time (min)	Peak Energy (MWh)	Cumulative Energy (MWh)	Peak Power (MW)
0.0001	84469504.00	1941.56	12.48	1.00	3497.00	316.86	9.99	444790.34	3884.11	2268.33
0.0010	63740992.00	1974.52	10.51	1.00	6311.00	326.56	10.00	166448.61	463.02	2070.58
0.0100	65061336.00	1990.10	15.52	0.99	8640.00	331.54	10.00	42843.95	59.85	1935.17
0.1000	67354520.00	1994.44	20.17	0.99	9343.00	331.99	10.00	8006.48	8.82	2171.31
1.0000	329283072.00	1994.72	20.71	0.99	9451.00	331.86	10.00	0.00	0.00	0.00

Figure 28 shows the sensitivity weightings distribution.

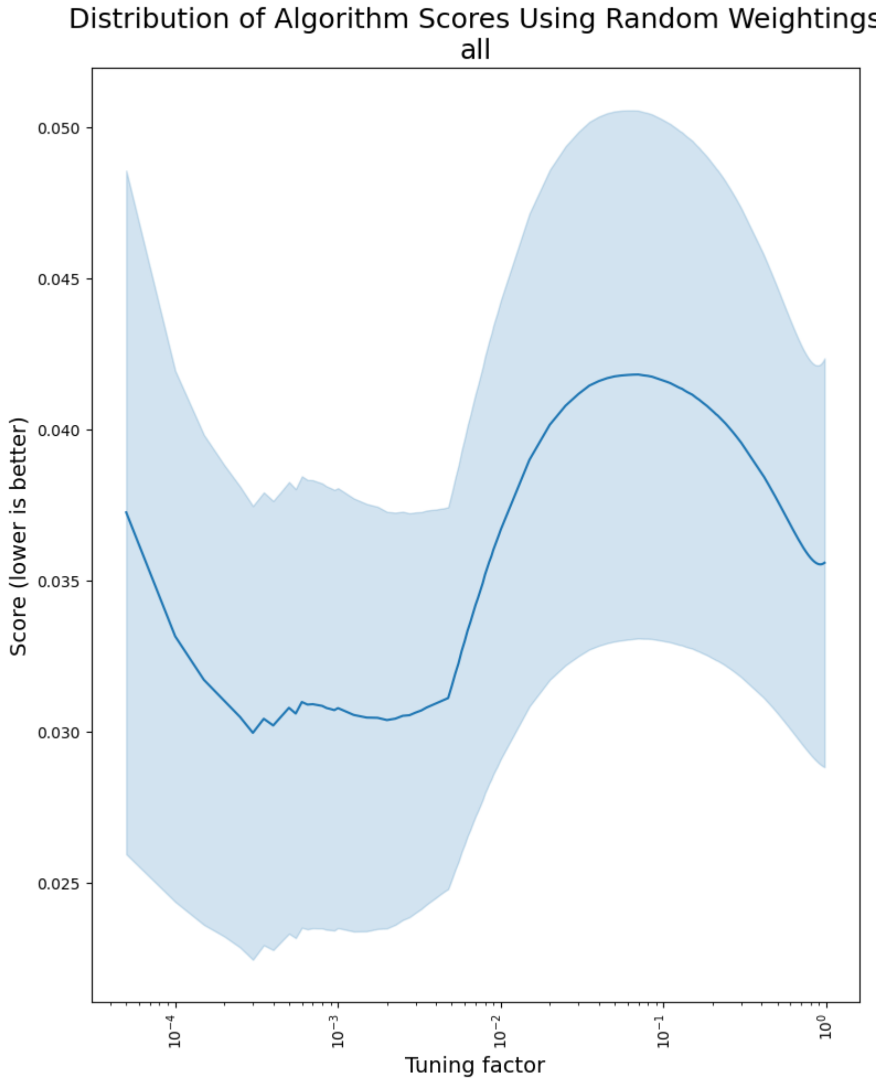

Figure 28: Full sensitivity analysis with all ten metrics.

Algorithm Optimisation and Conclusions

Seven algorithms were evaluated using multiple assessment metrics. These metrics served as inputs to a sensitivity analysis, which revealed that certain algorithms excelled in specific areas. For example, peak shaving algorithms performed poorly in producing a smooth flow but were effective at reducing the frequency and magnitude of constraint limit exceedances.

While no algorithm was the best for all assessment metrics, the exponentially weighted moving average (EWMA) algorithm was deemed the optimal overall choice when considering a combination of all assessment metrics. This was chosen as the preferred smoothing approach for further analysis.

EWMA's behaviour is governed by a tuning parameter (alpha) that affects how the algorithm weights historical data. The optimisation of the tuning parameter using a sensitivity analysis has resulted in a value of 0.001 being chosen for use with smoothing. This value balances the performance of the smoothing algorithm and the sizing requirements of asset that could deliver the service. This alpha value can be tweaked to ensure the asset requirements are technically feasible while keeping the smoothing efficacy relatively high.

It is worth noting that there was significant sensitivity to the relative importance weighting of the assessment metrics and, should NESO in the future wish to reprioritise certain metrics, the preferred smoothing algorithm settings may be different.

Overall, across the different smoothing algorithm options and the optimised EWMA set, the metrics suggest that the use of the optimised smoothing algorithms would:

- Decrease flow volatility by approximately 80% (using the smoothing index) compared to unsmoothed flow.
- Reduce the number of constraint exceedances by between 15-35%.
- Achieve the above benefits for a required asset capability of between 100 and 500 MWh and peak power of approximately 2 GW.

Appendix I – Stakeholder Engagement Feedback

Engagement Approach

Twelve industry stakeholders expressed interest in the project, and all but one participated in an interview. The stakeholders covered a range of technology types including:

- BESS developers, operators and optimisers.
- Aggregators/optimisers of flexibility technologies including BESS, district heat, industrial demand, EV chargers, hydro stations.
- Flexible data centres (cryptocurrency mining).
- Wind farm operator (including co-located assets).

The majority of stakeholders that participated were lithium-ion BESS operators or optimisers. The stakeholders were open with their responses and these have informed the technology characterisation that was performed as part of the project.

A briefing pack was issued in advance of holding one-to-one meetings to gather feedback and insights from the stakeholders. Written feedback was also encouraged via Microsoft Forms in advance of the meetings. The questions shared in the briefing pack were asked in order to understand:

- Technical characteristics of assets (in general, and of specific assets that could provide a service).
- Operational characteristics.
- Wider considerations.

The feedback from these interviews and online form responses is summarised in this section. The feedback themes have been summarised under:

- Operation.
- Data and Dispatch.
- Location.
- Service Design.

Operation Feedback

- A boundary flow service will be a change in behaviour for assets so there needs to be certainty on the requirements and sufficient incentive to change behaviour.

- Ramp rate limits on individual BESS (e.g. BMU assets are limited to a ramp rate of 50 MW per minute when changes are larger than 300 MW from Grid Code) may mean that the service needs to be provided by a larger number of assets to meet power requirements of the service.
- If a high energy throughput service is needed, this will increase degradation of energy storage assets. A higher service price may be needed to reflect the high usage and impact on asset health.
- Energy-limited assets will need to be able to perform state of charge (SoC) management i.e., be able to export if the service has required a lot of import.
- Uncertainty over asset SoC for the following service procurement period (they do not know what they will be required to dispatch) limits what capacity assets will commit. Providing certainty on final SoC would provide more confidence for assets to participate.

Data and Dispatch Feedback

- The dispatch and operation would be different to current services, for example frequency response where assets can measure frequency themselves and dispatch accordingly.
- Feedback suggested that the BOA control point could be utilised to send instructions as quickly as every few seconds, as opposed to the minute timescales currently utilised for BOA instructions.
- For service operation, automated dispatch was the clear preference with no manual dispatch given the likely response time requirements.
- This service would likely require a boundary meter feed and a constraint limit feed to assets for them to operate according to service parameters autonomously and without manual instructions.
- There will need to be some mechanism for assets to send signals to NESO on their SoC while participating.

Location Feedback

- There are grid scale assets located in the regions of interest for the boundaries of interest (100’s MW, including some Long Duration Energy Storage (LDES) BESS).
- There are a large number of planned and under construction projects some of which are very large projects (100+ MW).
- Due to the locational nature of the service, there is likely to be insufficient grid-scale assets currently connected to have a liquid market so it is important to allow as many Flexibility Service Providers (FSP) as possible to participate.
- Further detailed network modelling will be needed to understand how the location of an asset on the grid and its connection voltage will influence its smoothing capability.

- Further analysis will be needed for distribution-connected assets to consider the effect of local constraints.
- Many distribution assets in Scotland will have active network management (ANM) connections which will also require further analysis to determine whether this will interfere with being able to provide the service (i.e. not able to achieve required service export because of constraints).

Service Design Feedback

- Aligning the procurement of this service with other services, in particular with the Enduring Auction Capability (EAC), will allow assets to optimise the procurement of different services and give them more certainty on being able to provide the service.
- There could be an opportunity to optimise this service with frequency response if that service becomes procured locationally e.g. procure less frequency response in Scotland if smoothing is active.
- The ability to stack this service with other services will allow assets to provide it at a lower cost.
- Initial service prices could be higher than expected during enduring operation to encourage early participation (as was the case for other services).
- If it is a low utilisation service then availability payments would probably be required for participation. Relatively high utilisation would be required for some demand response to be economical.
- Long-term contracts would be beneficial to encourage investment in some large-scale assets.
- Multi-day procurement periods may help mitigate issues around SoC management.

Appendix J – Technology Options

This section outlines flexibility technology options and includes any insights gained from research and engagement.

Battery Energy Storage Systems (BESS)

A BESS is an electrochemical asset that stores electrical energy in rechargeable battery cells for later use, enabling grid stabilisation, load shifting, frequency regulation, and integration of intermittent renewable energy sources.

- All BESS are designed to be capable of providing Dynamic Containment (DC) (this is the fastest frequency response service and requires response within 0.5 s).
- 2-hour assets are the most typical assets of the stakeholders engaged, although many developers/operators are now looking to expand duration to 3-4 hours. Some are considering participating in LDES schemes of >8 hours, with the right incentives.
- There is also a trend towards larger power capacity BESS, with new developments typically between 200–600 MW (but up to ~1 GW).
- BESS assets are generally providing a full suite of ancillary and balancing services to NESO (frequency response, reserve, BM, wholesale, stability and reactive power).

Flexible Data Centres

Flexible Data Centres contribute to grid balancing by dynamically modulating their power consumption in response to grid signals, leveraging demand-side response capabilities and load shifting to absorb excess generation or reduce load during peak demand periods, thereby enhancing grid resilience and supporting renewable integration.

- Data centres can provide demand turn-up and turn-down, adjusting their processing capabilities, as necessary.
- They are primarily a turn-up asset as their operating model is cryptocurrency mining and relies on using power when it is cheap. Turn-down relies on them already operating at a set level which has to be incentivised.
- The design of the data centres would be tuned to the expected utilisation level (type of computers/servers) and power level can be scaled.
- They offer potentially limitless 'storage' capacity as they are an energy consumer.
- A key challenge is balancing remuneration with loss of revenue resulting from participation in grid activities.

Hydro

Hydropower supports energy grid operability by providing rapid-response, dispatchable generation through reservoir and pumped-storage systems, enabling frequency regulation, load balancing, and reserve capacity to accommodate fluctuations in demand and intermittent renewable output.

- Hydro is a well-developed flexibility asset that can import/export energy, although they do this over longer time periods than BESS due to slower ramp rates.
- Smaller hydro sites typically have quicker ramp rates than larger generation sites.
- Energy storage capacities can be significantly larger than BESS (orders of magnitude) as they are limited only by the size of the upper and lower reservoirs.
- Geographically constrained to areas of high elevation delta, often in Scotland but away from system boundaries where power constraints are apparent.

Aggregated EV

Aggregated Electric Vehicles (EVs), coordinated via smart charging and vehicle-to-grid (V2G) technologies, function as a distributed energy resource that can provide grid services such as peak shaving, and load balancing by dynamically adjusting charging patterns and exporting stored energy during periods of grid stress.

- Aggregated EV charging coordinates multiple electric vehicles to charge intelligently, often using smart algorithms and grid signals to optimise charge periods.
- Due to increasing adoption of EVs across the UK, aggregated EV charging could be deployed to areas of high constraint with superior locational advantage over other 'static' technologies.
- EV chargers have not implemented the type of charging required for smoothing, i.e. charging from ancillary service signals, they currently just operate by charging due to simple price signals. New control algorithms and signals would be needed to enable this fast-acting charging capability.
- The times when a smoothing service might be required could be outside of typical overnight charging schedules, so the higher energy costs for charging would need to be compensated.

Industrial Demand & District Heating

Industrial demand and district heating systems contribute to energy grid balancing by enabling flexible load management and thermal energy storage, allowing for strategic load shifting, peak demand reduction, and integration of surplus renewable electricity through power-to-heat conversion and coordinated demand-side response.

- District heating systems can utilise constrained periods by heating up communal hot water tanks, which are insulated and act as thermal stores.
- Engagement with stakeholders suggests these heating technologies require up to 10 minutes to respond to signals and may need further time (several minutes) to ramp up.
- Industrial customers can also participate in constraint management by increasing/decreasing production capacity, managing cold storage or utilising on-site storage (or equivalent) assets.
- Industrial load shedding can be done within 1 second.
- Industrial compressors can be ramped up within 3-4 minutes to provide demand turn-up.

Additional Technologies

On top of the key technologies identified in the preceding section, a comprehensive list of the technologies was considered. Alternative storage technologies that could also provide a smoothing function are listed below however at this point in time they have been triaged as unviable for smoothing:

Liquid Air Energy Storage (LAES)

Cryogenically cooled air is liquified and stored at low pressure during times of thermal constraint. When capacity allows, air is evaporated to drive a turbine, producing power.

- LAES is typically flexible in terms of deployment; can be developed near grid boundaries or energy sources.
- LAES is typically considered an LDES asset and could provide around 12 to 24 hours of storage. Given plant characteristics, it would likely be unidirectional smoothing during a charging phase.

Compressed Air Storage

Similar to LAES but without the phase change. Requires large reservoirs i.e. underground caverns to store air.

- Compressed air storage provides very long duration storage, but is limited to geologically/seismically inactive regions.

Flow Batteries

Energy is stored in liquid electrolytes that are pumped through a cell to generate electricity.

- Flow batteries are a scalable technology with traditionally longer lifespan than conventional chemical BESS, due to separation of the energy storage medium (electrolyte) from the electrochemical process.
- Can offer higher energy capacity than Li-ion BESS.

There are no identified projects or sites for the additional viable technologies (although there appear to be some projects based on LAES and flow batteries in development in Scotland) so focus will remain on the already identified technologies as assets are available and the technology has been tested.

Some storage technologies provide high power density, but lower energy density. These are better suited to short term stability or response services, and likely will not meet the duration requirements of a smoothing service. Additionally, there are demand assets that are unlikely to have the required ramp rates and flexibility to provide a smoothing service. These technologies are:

Supercapacitors

Energy is stored electrostatically between electrodes using materials that allow for a much higher capacitance than traditional capacitors.

- Supercapacitors provide high power density, rapid charge/discharge cycles, and long cycle life (all superior to BESS)
- Suffer from voltage leakage and are unsuitable for long-duration storage.

Flywheel Energy Storage Systems (FESS)

Flywheels are rotating masses that store energy as rotational kinetic energy.

- Flywheels provide high power density and fast response. They are 'spinning' assets so also contribute to frequency stabilisation.

- Mechanical friction causes energy losses and therefore is unsuitable for long-duration storage.

Power-to-X

Power-to-X captures hydrogen production facilities (i.e., electrolysers) that will use energy from the grid to produce hydrogen when grid signals indicate that a demand increase is required.

- Some electrolysers are more flexible than others, but likely on par with response times of industrial demand.
- They have been identified as unviable because they will struggle with faster ramp rates and are better suited for bulk constraints compared to a smoothing service.

Appendix K – Technology Options Suitability Analysis

Table 15 shows the technologies assessed for suitability against different smoothing service criteria. The suitability assessment identified technologies that should be progressed to archetype development, and these technologies will be evaluated against smoothing service requirements.

The criteria are:

- Technical Capability (power rating, energy capacity, power direction, duration of sustained action, response time, ramp rate).
- Technology Maturity ('ready to go', 'needs some development', 'not deployed at scale').
- Asset Availability ('pre-existing', 'close to commissioning', 'in development with connection', 'not yet planned').

The technologies have been given a Red Amber Green (RAG) rating:

- Green indicates that the technology performs well against the criteria.
- Amber indicates that the technology has potential to meet the criteria.
- Red indicates that the technology currently performs badly against the criteria.

Table 15: Technology suitability analysis

Technology	Participating Stakeholders[4]	Technical Capability	Technical Maturity	Asset Availability
BESS	10	🟩	🟩	🟩
Flexible Data Centres	2	🟩	🟩	🟧
Hydro	1	🟧	🟩	🟧
Aggregated EV charging	3	🟩	🟧	🟩
Industrial Demand & District Heating	1	🟧	🟩	🟧
LAES	0	🟧	🟧	🟥 (some relevant developments)
Compressed Air Storage	0	🟧	🟧	🟥
Flow Batteries	0	🟩	🟧	🟥 (some relevant developments)
Supercapacitors	0	🟥	🟩	🟥
FESS	0	🟥	🟩	🟥
Power-to-X	0	🟧	🟥	🟥

---

[4] In total 12 organisations participated in the stakeholder engagement with some organisations covering multiple technology types.

Appendix L – Technology Archetypes

Table 16 shows the technology archetypes developed [26]. These will be used to compare with service requirements to determine viable technologies.

Table 16: Technology archetypes

Asset	Power rating (typical)	Energy capacity (typical)	Duration of sustained action (typical)	Power direction	Response time	Ramp rate
BESS	50 to 600 MW (with a trend towards newer developments >200 MW, and can be over 1 GW)	50 to 1200 MWh	1 to 4 hrs (with trend of increasing duration)	Generation up/down Demand turn up/down	Sub-second (0.5 s baseline depends on communication and plant)	100% in <1 second
LDES BESS	> 100 MW	> 800 MWh	8 to 10 hrs (>8 hrs end of life capacity requirement part of LDES cap-and-floor scheme)	Generation up/down Demand turn up/down	Sub-second (0.5 s baseline depends on communication and plant)	100% in <1 second
Flexible Data Centres	20 to 100 MW (but scalable upwards)	Unlimited	Unlimited	Demand turn up/down	<2 seconds	100% in <5 seconds

Asset	**Power rating (typical)**	**Energy capacity (typical)**	**Duration of sustained action (typical)**	**Power direction**	**Response time**	**Ramp rate**
Small Hydro	<5 MW	Varies, subject to reservoir capacity	Varies, subject to reservoir capacity	Generation up/down	Seconds-minutes	100% in 2 minutes
Large Hydro	>5 MW	Varies, subject to reservoir capacity (up to 10 GWh, e.g. Dinorwig PSH)	Varies, subject to reservoir capacity	Generation up/down Demand turn up/down (if pumped)	Minutes	100% in 3-5 minutes
Industrial Demand	50 kW to >1 MW	-	Minutes-Hours	Demand turn up/down	Seconds-Minutes	100% in 3-8 minutes
Aggregated EV	50 kW to >1 MW	-	Minutes-Hours	Demand turn up/down	Seconds-Minutes	100% / 3-8 minutes
District Heat	50 kW to >1 MW	-	Minutes-Hours	Demand turn up/down	Minutes	100% in 10 minutes

Appendix M – Project Assumptions

The assumptions in Table 17 were used to simplify the problem and enable the feasibility of smoothing to be assessed.

Table 17: Assumptions used during the project

Category	Number	Assumption	Justification	Impact on results
Data	A01	Constraint tagged BOAs do not completely capture all constraint actions.	This is a well-known issue as tagging is done after the fact, so some actions are missed. The adjusted flows have been used to determine constrained periods instead, although this does remove all BOAs rather than purely constraint BOAs.	The pool of BOAs to select from when assessing the benefits via unwinding BOAs is reduced. The methodology to identify constrained periods will overestimate the constrained periods relative to the number of tagged BOAs.
	A02	Constraint BOAs are not always correctly tagged to the right boundary.	BOAs are tagged to the closest boundary rather than the boundary where they are addressing a constraint.	Costs are wrongly associated and some BOAs may not be available to choose through the nested boundaries analysis.
	A03	Inaccuracies in data are not sufficient to materially bias results.	The inaccuracies are not judged to be significant.	Any inaccuracies should not materially affect the project findings.
Statistical Analysis	A04	Historical boundary flow data can be used to gain insight into the likely statistical behaviour of future boundary flows.	The analysis has shown that historical, metered data can be used to understand the behaviour of the smoothing algorithms. It was found	The analysis is based on the statistical probability of how historical boundary flows have changed. If the distributions

Category	Number	Assumption	Justification	Impact on results
			that, statistically, flows behave differently depending on whether Control Room constraint actions were taken. At high power flows, there are sufficient numbers of both constrained and unconstrained periods to develop the statistics of how the mean flow evolves and how flow fluctuates upwards from this. These two factors can be used to infer how likely smoothing is to deliver additional, usable capacity on constrained lines.	change the analysis would need to be updated.
	A05	The metered flow was the most realistic flow to use for smoothing	The statistics of the adjusted flow are significantly different to those of the metered flow. Particularly, the adjusted flows have much lower instances of small differences between successive 10-minute periods and much higher instances of very large differences. This is not representative of either the constrained or unconstrained flows within metered data.	The analysis is more realistic when using metered flows.
	A06	Boundary flows may be considered in 10-minute averaged windows.	It matches operational timescales: Control Room updates flow limits every 10 minutes.	Basis of modelling.

Category	Number	Assumption	Justification	Impact on results
			Generator ramping aligns with 10-minute timescales. Thermal ratings are often based on 10-minute overload durations.	
	A07	First-difference distributions are an appropriate proxy for short-term forecast noise, i.e., macro trends (wind, weather, demand patterns) do not materially affect differences between adjacent 10-minute windows.	Macro trends typically impact flows on timescales greater than or equal to a settlement period and are unlikely to be the main cause of change between consecutive 10-minute periods.	The impacts of these is not zero and therefore the breadth of the noise-only distribution is likely to be narrower than that used in the project. This will lead to an over-estimation of the safety margin and hence an overestimation of the project benefits.
	A08	First-difference distributions based on the highest 70% of unconstrained flows.	This most closely matches the situations in which smoothing might be deployed and therefore reflects the volatility that impacts on Control Room decisions.	This choice of flow subset is the most representative of volatile and unconstrained flows.
	A09	In assessing the smoothing algorithms, the metrics capture entirely distinct characteristics (i.e. no overlapping and no double-counting).	A sensitivity analysis and normalisation of results was undertaken to ensure the algorithm chosen was the most appropriate.	The choice of smoothing algorithm is assumed to not have a material difference on results.
	A10	Exceedance probability is a proxy for the harm of exceedance.	This would be better assessed by considering the energy typically associated with an exceedance, but	Results comparing the energy transferred over the limit show very close agreement over all

Category	Number	Assumption	Justification	Impact on results
			this is more computationally challenging.	boundaries apart from SSEN-S. Overall the value across all boundaries is ~97% of that seen in the unsmoothed case. The discrepancy means that the smoothed case is very slightly underestimating the benefits from BOA reduction, as some of the benefit is still captured as a reduction in asset harm.
Modelling	A11	Smoothing being active is calculated independently for each boundary.	Smoothing is operational during constrained periods regardless of whether the benefits of smoothing are being calculated.	The analysis overestimates asset usage due to this assumption.
	A12	BOAs are assumed to operate for the whole settlement period to convert energy into power.	The analysis calculates additional power that can be transmitted through boundaries within 10-minute periods. BOAs are recorded as changes in energy over 30-minute settlement periods. The energy is converted to power assuming the energy was delivered constantly over the settlement period. An intra-settlement breakdown is not available so this assumption must be used.	The change in power will be slightly inaccurate for the revised flows.

Category	Number	Assumption	Justification	Impact on results
	A13	BOAs associated with the boundary where additional power is being allowed through are used first.	The analysis logic unwinds constraint bids associated with the boundary being assessed rather than the most expensive bid of any boundary north.	The benefits will be slightly underestimated if there are more expensive bids further north.
	A14	Network latencies and losses are not considered.	The smoothing occurs instantaneously and all the assets output effects the boundary flow. In reality this is unrealistic due to latencies and the effect of losses on the network. These simplifications were necessary for the analysis to determine the feasibility of the concept as encompassed by the scope of this project. Wider network effects would need to be considered as part of the scope of a follow-on project.	The smoothing impact on flows is overestimated.
	A15	Asset losses are not considered.	The analysis has been undertaken from a technology agnostic perspective so exact asset losses are difficult to implement at this stage. Should detailed feasibility studies be undertaken this should be considered in the future.	The smoothing impact on flows is overestimated.
	A16	Independently smoothing each boundary produces an adequate approximation without modelling	Keeping the modelling at a level of complexity that is commensurate with the stage in the project lifecycle (i.e. early feasibility).	Overestimates the amount of energy that smoothing assets need to move in order to deliver the algorithm requirements.

Category	Number	Assumption	Justification	Impact on results
		the propagation of smoothed flows to downstream boundaries.		
	A17	Additional flow is only useful if it can exit the nested set at SSHARN3.	Unwinding BOAs while SSHARN-3 is constrained means constraint actions would need to be retaken.	Appropriate for a network operation assumption.
	A18	BOA power is constant throughout the three 10-minute periods of a settlement period	Allows the inference of a power flow increase from the energy data that is contained in the BOA data.	Occasional, very small calculation discrepancies depending on whether data is considered per 10-minutes or per settlement period.
	A19	BOAs may be unwound without causing secondary system impacts either within the current settlement period or in future ones. Nor does it produce flow bottlenecks further south in the network.	Not possible to model these effects without including network modelling.	Unwinding of BOAs may be overestimated.
	A20	Smoothing would have no effect on other BM or wholesale market prices.	Allows the use of historical prices.	Basis of modelling.
	A21	All replacement energy offers south of SSHARN3 are viable to counteract the unwinding of a bid.	There is no way to link offers to bids from the data. Choice is made on the basis that all are viable and the most expensive would be unwound.	This overestimates the savings.
	A22	The ceiling price calculated does not include the following: the costs	These are unknown at this stage.	Potentially the costs will eliminate all savings, and the project Net

Category	**Number**	**Assumption**	**Justification**	**Impact on results**
		of developing the service and the market, the costs of running the market. These are assumed to be of-the-order-of single millions of GBP.		Present Value (NPV) will be close to or below zero.
	A23	Unwound BOAs add zero-volatility flow.	The volatility of the output from the unwound generation is not included. The actual output that would have been added is unknowable. The added flow is assumed to deliver the energy removed by the original BOA at a uniform power over a 10-minute period.	Underestimation of the volatility of the revised, BOA-unwound flow, leading to slight over-estimate of the benefits of smoothing.
Operation	A24	The Control Room does not formally apply a safety margin but this is implicitly present within decision making when flows are volatile meaning that operators take more conservative BOA actions when flow is volatile.	If operators consistently attempt to keep the flow at the limit, there would be higher incidence of exceedance. The fact that this is not observed suggests that short-term volatility plays a role in decision-making.	If this assumption is false, there can be no BOA-related benefit from smoothing.
	A25	Operator behaviour can be considered to be consistent and repeatable over time.	Without an assumption of consistency of process, the analysis is impossible.	Basis of modelling.
	A26	The Control Room would alter BOA choices if volatility were reduced.	Ties in with assumptions on volatility playing a role in decisions.	Basis of modelling.

Category	**Number**	**Assumption**	**Justification**	**Impact on results**
	A27	Operator behaviour may be inferred from historical exceedance occurrence.	This provides the link between exceedance and behaviour.	Basis of modelling.
	A28	Forecasts used by the Control Room are unbiased and accurate i.e., any discrepancy between the forecast and reality is due to irreducible noise.	There is no access to the forecasts, so it is not possible to judge the extent to which forecast accuracy contributes to exceedance.	Over-estimate of the impacts as some of the exceedances will be caused by forecast error.

2. Glossary

Abbreviation	*Description*
Adjusted Flow	Power flow with the bids and offers unwound from the **metered flow** to calculate a flow that precedes any actions taken by the Control Room
BESS	Battery Energy Storage Systems
BM	Balancing Mechanism
BMU	Balancing Mechanism Unit
BOA	Bid Offer Acceptance
CAES	Compressed Air Energy Storage
CBA	Cost Benefit Analysis
Constrained Period	Period where the **adjusted flow** was above the constraint limit i.e. Control Room actions had to be taken to manage constraints
Constraint Limit	The assigned power flow capacity of a transmission boundary. If the flow in absence of BOA actions is above the limit, then the boundary is referred to as constrained. Limits are updated on a 10-minute basis
ENCC	Electricity National Control Centre
EV	Electric Vehicle
EWMA	Exponentially Weighted Moving Average
Exceedances	When the boundary flow exceeds the constraint limit
FESS	Flywheel Energy Storage Systems
FSP	Flexibility Service Provider
GB	Great Britain
Implied Safety Margin	The implied safety margin of power flow below the constraint limit that predicts a certain level of limit exceedance based off the mean flow change distributions
LAES	Liquid Air Energy Storage
LDES	Long Duration Electricity Storage

Operational Margin Deployment Cases	The CBA deployment cases that assess the effect of different operational margin for when smoothing benefits are calculated
Metered Flow	Historical power flow data recorded by meters crossing boundaries
PNA	Power Network Analysis
Power-to-X	Power-to-X captures hydrogen production facilities (i.e., electrolysers) that will use energy from the grid to produce hydrogen when grid signals indicate that a demand increase is required.
PV	Photovoltaic
Revised Flow	The **metered flow** across a boundary with bids unwound to allow additional power through without exceeding the **implied safety margin**
SMA	Simple Moving Average
Smoothed Flow	Power flow that has been smoothed. For the CBA results the **revised flow** has been smoothed
SoC	State of Charge
Unconstrained Period	Period where the **adjusted flow** was below the constraint limit i.e. no Control Room actions had to be taken
VTL	Virtual Transmission Lines: Using storage as a transmission asset to enhance transmission capacity

3. References

[1] Fluence, “Redrawing the Network Map: Energy Storage as Virtual Transmission,” [Online]. Available: https://info.fluenceenergy.com/hubfs/Collateral/Storage%20as%20Transmission%20White%20Paper.pdf.

[2] U. S. D. o. Energy, “Enabling Principles for Dual Participation by Energy Storage as a Transmission and Market Asset,” [Online]. Available: https://www.pnnl.gov/main/publications/external/technical_reports/PNNL-32196.pdf.

[3] I. R. E. A. (IRENA), “VIRTUAL POWER LINES INNOVATION LANDSCAPE BRIEF,” [Online]. Available: https://www.irena.org/-/media/Files/IRENA/Agency/Publication/2020/Jul/IRENA_Virtual_power_lines_2020.pdf?la=en&hash=C58043124D596D1CF75395066817C38B55AC1983.

[4] C. M. I. S. Andrea Biancardi, “Battery electricity storage as both a complement and substitute for cross-border interconnection,,” *Energy Policy,* vol. 189, no. https://www.sciencedirect.com/science/article/pii/S030142152400154X, 2024.

[5] S. M. M. H. B. V. D. A. S. D. V. G. S. W. B. M. M. L. V. H. Giovanni De Carne, “The role of energy storage systems for a secure energy supply: A comprehensive review of system needs and technology solutions,” *Electric Power Systems Research,* vol. 236, no. https://www.sciencedirect.com/science/article/pii/S0378779624008496?ref=pdf_download&fr=RR-2&rr=99cec52dbaa16364, 2024.

[6] Fluence, “Building Virtual Transmission: Critical Elements of Energy Storage for Network Services,” [Online]. Available: https://info.fluenceenergy.com/hubfs/Building%20Virtual%20Transmission.pdf.

[7] I. G. a. E. B. P. Ilyushin, “Application Methods of Efficient Electricity Storage Systems to Improve Flexibility of Power Systems,” in *International Russian Automation Conference (RusAutoCon)*, 2024.

[8] Transnet BW, “Grid Booster,” [Online]. Available: https://www.transnetbw.de/en/company/portrait/innovations/grid-booster.

[9] Amprion, “Decentralised Grid Booster,” [Online]. Available: https://www.amprion.net/Grid-expansion/Our-Projects/Dezentraler-Netzbooster/.

[10] Re-Twin Energy, “Industry Update | Amprion GmbH’s New “Netzbooster” Tender: A Game-Changer for Grid Stability,” [Online]. Available: https://re-twin.energy/blog/news/industry-update-amprion-gmbhs-new-netzbooster-tender-a-game-changer-for-grid-stability.

[11] Fluence, “Grid Boosters as innovative solution to optimize power grids,” [Online]. Available: https://www.aes.com/sites/aes.com/files/2024-04/Consentec%20%26%20Fluence_Grid%20Boosters%20as%20Innovative%20Solution%20to%20Optimize%20Power%20Grids_Part%201.pdf.

[12] Fluence, “TenneT Grid Boosters Will Enhance Transmission Capacity and Help Reduce Costs for Customers in Germany,” [Online]. Available: https://info.fluenceenergy.com/hubfs/Project%20Spotlight_TenneT_Fluence_PS-013-01-EN.pdf.

[13] Renew Economy, “Why networks think battery storage may be smarter choice than more poles and wires,” [Online]. Available: https://reneweconomy.com.au/why-networks-think-battery-storage-may-be-smarter-choice-than-more-poles-and-wires-16241/.

[14] Renew Economy, “New transmission lines are controversial – batteries and virtual lines may ease the pain,” [Online]. Available: https://reneweconomy.com.au/new-transmission-lines-are-controversial-batteries-and-virtual-lines-may-ease-the-pain/.

[15] RTE France, “Electricity storage - the RINGO experiment,” [Online]. Available: https://www.rte-france.com/projets/stockage-electricite-ringo#Leprojet.

[16] M. &. P. J. &. Q. E. &. V. V. &. S. A. &. A. H. &. M. J. &. C. R. Agüero, “Virtual Transmission Solution Based on Battery Energy Storage Systems to Boost Transmission Capacity,” *Journal of Modern Power Systems*

and Clean Energy, vol. 12, no. https://www.researchgate.net/publication/379296115_Virtual_Transmission_Solution_Based_on_Battery_Energy_Storage_Systems_to_Boost_Transmission_Capacity, pp. 466-474, 2024.

[17] Frazer-Nash Consultancy, "Boundary Flow Smoothing Feasibility Report," FNC 029835-60019R, 2026.

[18] J. Martins, S. Spataru, D. Sera, D.-I. Stroe and A. Lashab, "Comparative Study of Ramp-Rate Control Algorithms for PV with Energy Storage Systems," *Energies,* vol. 12, p. 1342, 2019.

[19] A. González-Moreno, J. Marcos, I. de la Parra and L. Marroyo, "A PV Ramp-Rate Control Strategy to Extend Battery Lifespan Using Forecasting," *Applied Energy,* vol. 323, 2022.

[20] A. Atif and M. Khalid, "Savitzky–Golay Filtering for Solar Power Smoothing and Ramp Rate Reduction Based on Controlled Battery Energy Storage," *IEEE Access,* vol. 8, 2020.

[21] "Smoothing of Renewable Energy Generation Using Gaussian-Based Method with Power Constraints," *IEEE Access,* vol. 8, 2020.

[22] M. Mao, Y. Tang, J. Chen, F. Ma, Z. Li, H. Ma, H. Sun, C. Yin and H. Li, "An Overview of Solar Photovoltaic Power Smoothing Control Strategies Based on Energy Storage Technology," *Energies,* vol. 18, p. 909, 2025.

[23] X. Han, F. Chen, X. Cui, Y. Li and X. Li, "A Power Smoothing Control Strategy and Optimized Allocation of Battery Capacity Based on Hybrid Storage Energy Technology," *Energies,* vol. 5, 2012.

[24] M. Rahimi, "A Simple and Effective Approach for Peak-Load Shaving Using BESS," 2017.

[25] R. Groß and e. al, "ndustrial Peak Shaving with Battery Storage Using a Probabilistic Quantile Forecast," *Applied Energy,* vol. 327, 3033.

[26] NESO, "ESO Routes to Market Review for Demand Side Flexibility - Stage 1 Archetypes final," [Online]. Available: https://www.neso.energy/industry-information/flexibility/demand-side-flexibility-routes-market-review.

[27] NESO, "Markets Roadmap," 2025. [Online]. Available: https://www.neso.energy/document/358126/download.

[28] Modo Energy, "Modo Energy Portal," [Online]. Available: https://modoenergy.com/home.

[29] BEIS, “Combined Heat and Power - Environmental,” Department for Business, Energy & Industrial Strategy, 2021.

[30] NESO, “Carbon Intensity Balancing Actions Methodology,” [Online]. Available: https://www.neso.energy/data-portal/carbon-intensity-balancing-actions/carbon_intensity_balancing_actions_methodology.

[31] NESO, “Carbon Intensity API,” [Online]. Available: https://carbonintensity.org.uk/.