Working in partnership

Flexibility Baselining Tool – Mathematical Specification

Open Networks

February 2022 | Version 1.0

TABLE OF CONTENTS

Contents

Introduction 3

Core algorithm overview 4

Baselining Algorithms 5
Historic Baseline 5
Eligible Days 5
Selection of Assessed Days 5
Unadjusted baseline calculation 6
Same-day adjustment 7
Assigned Baseline 7
Nomination baseline 7
Zero baseline 7
Response 7

Assumptions within the tool 8
Sign convention 8
Input data 9
Calculations 9

Introduction

This technical note provides a detailed explanation of the assumptions and calculations contained within the baselining algorithms that form part of the baseline tool developed by TNEI and the SSEN TRANSITION innovation project, in collaboration with the ENA Open Networks Project. The TRANSITION project, led by SSEN and with ENWL as a partner, is an Ofgem NIC funded innovation project that aims to inform the development of flexibility markets at scale and support the DNO to DSO transition – more information is available here[1].

The algorithms that underpin the TRANSITION/TNEI baseline tool have been developed following the recommendations in the report prepared for the Energy Networks Association by DNV-GL in 2020[2]. Where appropriate, for example if there are elements that have not been specified in the DNV-GL report, further details for the algorithm have been implemented following the method described in the New Thames Valley Vision (NTVV) report[3].

We provide a detailed description of the underlying core algorithms, and then a shorter description of the key assumptions that were required to translate the methodological recommendations into implementable algorithms.

This document is accompanied by a Functional Specification, which outlines the requirements of the tool, as well as a Demonstration and Documentation report that details how to use the tool for carrying out baselining calculations.

[1] SSEN Transition

[2] ON20-WS1A-P7 Baseline Methodology Assessment 2020

[3] TRANSITION Oxfordshire Programme Commercial Arrangements July 2020

Core algorithm overview

- The algorithm calculates baselines for a single day at a time, with defined start and end-times for a single flexibility event on that given day. Together, these define an event day $\hat{d}$ and a set of times-of-day within that flexibility window F.
- The core model is set to work with a time-series of power readings $\boldsymbol{p}$, and is agnostic as to whether these are demand or generation – the tool that wraps around this will convert from energy to power as required.
- This time-series comprises several individual power measurements $p_{d,\tau}$, with a timestamp defined by a combination of indices for days d and times-of-day τ. For example, the time-series for a specific day (say, 1st January 2022 with half-hourly granularity) would be defined as:

$$\boldsymbol{p} = \begin{pmatrix} p_{d=01/01/2022,\tau=00:00} \\ p_{d=01/01/2022,\tau=00:30} \\ \dots \\ p_{d=01/01/2022,\tau=23:30} \end{pmatrix}$$

  The algorithm can work with any time-series data with a consistent resolution, but has been designed with some specific resolutions in mind, including half-hour, one-minute, and intermediate resolutions such as 10 minutes, 5 minutes and 6 minutes.

- $b_{\hat{d},\tau}$ is then the baseline for time-of-day τ on event day $\hat{d}$.
- Baselines are determined for the full day from 00:00 until 23:59. The algorithm expects series to be provided in “British Summer Time”, with adjustments for clock changes included, and not in “Greenwich Mean Time”, and no further specific adjustments are made for daylight savings time.
- The algorithm can compare a series of actual power measurements, $p_{d,\tau}$, with the baseline for each day.

Baselining Algorithms

The tool has two categories of baselining algorithms:

- **Historic baseline:** this refers to the methods where recent historic data is used to calculate a rolling baseline. Default values of parameters have been set based on the recommendations made by DNV GL. The most significant option available is whether to apply a same day adjustment within the historic baseline.
- **Assigned baseline:** this is an umbrella term to describe several approaches where the baseline is assigned without the tool considering the historic data. It includes both the zero baseline (where the assigned value is always zero) as well as the nomination baseline (where the assigned value is given by the provider). This reflects similarities in their implementation within the tool.

Historic Baseline

Historic baselines are calculated using an "x-in-y $days$" approach, where historical data from the set Y of the last y eligible days is gathered, and then the subset X of x of those days is selected and assessed for calculating a baseline.

The algorithm works with a time series of power measurements p_t (either demand or generation) for each timestamp t. These are assumed to be the average power between the timestamp for each measurement, and the timestamp for the subsequent measurement. As with the baselines, we can define these based on day d and a time-of-day τ, working with a series of power measurement $p_{d,\tau}$.

Eligible Days

For every day for which a baseline is calculated, the series of power measurements is filtered to only include eligible days.

- Days for which there was previously a flexibility event are not eligible – power readings from these days are excluded entirely.
- For weekdays, weekends and bank holidays are not eligible – these power readings are excluded.
- For weekends and bank holidays, weekdays are not eligible – these power readings are excluded.

From the remaining data, all the power readings from the most recent y of these eligible days are retained, and all other data is discarded.

By default, $y = 10$ for weekdays and $y = 4$ for weekends and bank holidays, following the recommendation made to the ENA Open Networks 2021 WS1a P7 Baseline team.

Selection of Assessed Days

The recommendation to Open Networks was to use the middle x out of these y days, where $x \leq y$. For example, for the recommended "Mid 8 of 10" approach, the days with the lowest and highest load would be excluded, and the remaining 8 days would be assessed for calculating the baseline.

The algorithm has two different approaches for ranking the demand on each day to select which are the "middle" days to be included, and which are the lowest and highest to be excluded, as this was not specified in the recommendation to the ENA. Both are based on calculating a metric for ranking each of the days, defined as $p_d^{(rank)}$. The two approaches are:

- Calculate average demand for each day during the flexibility event:

$$p_d^{(mean)} = \underset{\tau \in F}{\text{mean}}\left(p_{d,\tau}\right)$$

- Calculate peak demand for each day during the flexibility event:

$$p_d^{(max)} = \max_{\tau \in F}\left(p_{d,\tau}\right)$$

The choice of the ranking approach dictates which of these two metrics is used:

$$p_d^{(rank)} = \begin{cases} p_d^{(mean)} & if\ mean\ approach\ used \\ p_d^{(max)} & if\ peak\ approach\ used \end{cases}$$

The algorithm initially discards the data from the $\frac{y-x}{2}$ days with the largest value of $p_d^{(mean)}$or $p_d^{(max)}$, where $\frac{y-x}{2}$ is rounded down if necessary. The model will then discard the $\frac{y-x}{2}$ days with the smallest values of $p_d^{(mean)}$or $p_d^{(max)}$, where $\frac{y-x}{2}$ is rounded up if necessary, such that x days are left, forming the set X.

Note that this means that, if $y - x$ is not an even number, then the algorithm will discard one further "small" $p_d^{(mean)}$or $p_d^{(max)}$ day than "large" $p_d^{(mean)}$or $p_d^{(max)}$ day. For example, if $x = 7$ and $y = 10$ then the model will discard the day with the largest value of $p_d^{(mean)}$or $p_d^{(max)}$, and the two days with the smallest values of $p_d^{(mean)}$or $p_d^{(max)}$. This would tend to lead to estimates of higher baselines which will be less onerous for providers, although this will only arise as an issue if using an odd number for either x or y, which could be avoided in practice.

Note that an alternative option is built into the tool which considers the absolute difference between the total energy consumption outside of the flexibility event for each of the set of Y historic days and the day for which the baseline is being calculated, keeping the subset X of days for which the absolute difference is lowest. This option was not part of the recommendation to Open Networks but was used within the baseline methodology for the NTVV project.

This results in x full days of power measurements, p_t, for the set of X assessed days.

By default, $x = 8$ for weekdays and $x = 2$ for weekends and bank holidays, following the recommendation made to Open Networks.

Unadjusted baseline calculation

The unadjusted historical baseline, $\bar{b}$, for each time-of-day τ within the event window defined by the set of times F is then calculated as the *average* power for that time-of-day τ for the set of X historic days that is being considered.

$$\bar{b}_{\hat{d},\tau} = \frac{1}{x}\sum_{d \in X} p_{d,\tau}$$

Same-day adjustment

If a same-day adjustment is not applied, then the final historical baseline, b, for each time-of-day is just equal to the unadjusted historical baseline:

$$b_{\hat{d},\tau} = \bar{b}_{\hat{d},\tau}$$

If a same-day adjustment window is defined, then the times τ in the w hours prior to the flexibility event will be within that adjustment window, where W is the set of times τ which are in the adjustment window.

The baseline adjustment is the average difference between the measured power on the day of the flexibility event, $d^{(event)}$, and the unadjusted historical baseline during the adjustment window, which is calculated as described above:

$$A_{\hat{d}} = \frac{1}{w}\sum_{\tau \in W}(p_{\hat{d},\tau} - \bar{b}_{\hat{d},\tau})$$

Then, the final historical baseline is calculated as the unadjusted historic baseline with this adjustment added on:

$$b_{\hat{d},\tau} = \bar{b}_{\hat{d},\tau} + A_{\hat{d}}$$

The recommendation made to Open Networks was to use a value of $w = 2$, i.e., a 2-hour same-day-adjustment window, and this is currently a fixed value within the calculations when the adjustment is applied.

Note that the algorithm is not currently designed to work with adjustment windows that extend back into the previous day.

Assigned Baseline

There are two different types of “assigned” baseline within the tool: nomination baseline, and zero baseline.

Nomination baseline

The algorithms do not carry out any calculations when using a nomination baseline. Instead, values of $b_{\hat{d},\tau}$ are uploaded by the user, for every event day $\hat{d}$ and time-of-day τ.

Zero baseline

For a zero baseline, $b_{\hat{d},\tau}$ is set to 0 for every event day $\hat{d}$ and time-of-day τ.

$$b_{\hat{d},\tau} = 0$$

Response

For all of the times-of-day during the event day $\hat{d}$ and during the flexibility window F, the response to the flexibility event is calculated as the difference between the baseline and the measured power.

$$r_{\hat{d},\tau} = b_{\hat{d},\tau} - p_{\hat{d},\tau} \text{ for } \tau \in F$$

The algorithm calculates the total energy supplied during the flexibility event, $E_{\hat{d}}$:

$$E_{\hat{d}} = \sum_{\tau \in F}\frac{r_{\hat{d},\tau}}{\theta}$$

where θ is the hourly frequency of the time-series in $1/hours$. For a half-hour series, $\theta = 2$, and for a 1-minute time-series $\theta = 60$.

If a series of required (contracted or dispatched) responses is uploaded, then the algorithm will test whether the delivered response achieved this requirement in every time-step:

$$r_{\hat{d},\tau} > c_{\hat{d},\tau} \times \varphi$$

where $c_{\hat{d},\tau}$ is the required response, and φ is a threshold level for a minimum acceptable response, which can take any value between 0% and 100% and has a default value of 95%. The algorithm records the flexibility response as "full delivery" if this requirement is met for every time-step within the window:

$$Success = \begin{cases} 1 & r_{\hat{d},\tau} \geq c_{\hat{d},\tau} \times \varphi \quad \forall \tau \in F \\ 0 & otherwise \end{cases}$$

If the flexibility response observed does not meet this threshold, then the event is considered a "partial delivery". Furthermore, if a baseline calculation is not possible this will be indicated.

Note that this is only intended to provide an indication of the nature of the response. Formal determination of the success of flexibility responses may be carried out separately by DNOs in accordance with their own settlement rules.

Assumptions within the tool

The tool implements a "wrapper" function that calls the core algorithm as required, calculating the baselines for several days and displaying these on the front-end.

A number of additional assumptions have to be made within this "wrapper" function in order to perform the baselining calculations.

Sign convention

In order to characterise the response expected from a flexibility provider we first need to define the direction of power flow that we associate with the asset type. The tool adopts the convention used by the wider TRANSITION project. Since demand is considered to be a withdrawal of power from the grid (the demand asset is consuming energy), this is denoted with a negative value. Conversely, generation is an injection of power to the grid (the generator is producing energy), which is therefore denoted with a positive value. Storage can act as either generation or demand, and can therefore be either positive or negative.

We also need to understand the type of grid constraint that needs to be eased by the response of a flexibility provider. The constraint can either be an import constraint, where power withdrawn from the grid must be limited, or an export constraint where power injected onto the grid must be limited.

The tool's "wrapper" function will process the data as required (specifically, if the constraint is an import constraint then all values will be multiplied by -1) before passing the power/energy series into the core algorithm. This ensures that increases or reductions in the magnitude of demand or generation will be properly recorded as a positive or negative response.

	Impact	Inclusion
Demand (withdrawal)	Data is processed to be positive: reduction in demand leads to positive response.	Data is kept as negative: increase in demand leads to positive response.
Generation (injection)	Data is processed to be negative: increase in generation leads to positive response.	Data is kept as positive: reduction in generation leads to positive response.

Note that this is handled automatically within the tool and does not alter the expected convention that demand will be negative, and generation will be positive.

Input data

If a time-series of the required response is provided, then all data (including metered data and any nomination baselines) are resampled to have the same frequency as the required response series. For example, if the required response is defined on a half-hourly basis, then all other calculations will be done with data resampled to a half-hourly granularity.

If no time-series of required response is provided, then there is an option to convert all calculations into a half-hourly granularity. This option is enabled by default.

Calculations

This wrapper function calculates a baseline for all days for which there are any timestamps within the metered dataset, even those that are not actually flexibility events. In such case, the algorithm will take the most frequent start and finish times of those flexibility events that are defined as the default start and finish times for days without a flexibility event. This allows for baselines to be examined for all days, even if no flexibility was ultimately required.

Energy Networks Association
4 More London Riverside
London SE1 2AU
t. +44 (0)20 7706 5100
w. energynetworks.org
@EnergyNetworks

Energy Networks Association Limited is a company registered in England & Wales No. 04832301
Registered office: 4 More London Riverside, London, SE1 2AU