D4.4.1 Tool for characterizing the aggregated flexibility of residential thermostatically controlled loads

November 28th, 2017

EcoGrid 2.0 is a research and demonstration project funded by EUDP (Energiteknologisk Udviklings- og Demonstrationsprogram). The 9 partners in the project are:

Main Authors:

Name/Partner	Email
Charalampos Ziras	chazi@elektro.dtu.dk
Daniel Esteban Morales Bondy	bondy@elektro.dtu.dk
Carsten Heinrich	cahei@elektro.dtu.dk

TABLE OF CONTENTS

1 Introduction

The goal of the EcoGrid 2.0 project is to develop and test new market and power system operation models, which facilitate the participation of demand response in offering power system services to the Transmission System Operator (TSO) and the Distribution System Operators (DSOs).

On a TSO level, regulating power is considered, because it is the fundamental means of balancing production and consumption of electric power in the Nordic area. Moreover, the high price volatility can provide higher revenues to the aggregators, compared to optimization of the consumption in the day-ahead markets. In contrast to the current practice in Denmark where regulating power is offered in single hourly bids, the project's proposed regulating market allows market participants to offer asymmetric bids in 15-min time intervals. An asymmetric bid consists of an initial response block, followed by a rebound block of the opposite direction, to allow the unit's energy state to return to its initial state (the one before providing the service). Such bids facilitate the participation of energy-constrained units, such as thermostatically controlled loads, electric vehicles and batteries, in the regulating market because they are energy-neutral. This is achieved by the nature of the regulating market clearing process, which accepts bids in their entirety, i.e. both the initial response block, followed by a rebound block. The authors of (N. O'Connell n.d.) have showed that such a market setup can improve the power system's performance and reduce the system balancing costs, by helping aggregators of distributed energy resources (DERs) to participate in the regulating market. On the DSO level, EcoGrid 2.0 proposes a market where aggregators of DERs can offer flexibility products to the DSO on a long-term scale (months), in order to improve distribution grid reliability, alleviate possible congestions or operational issues or defer investments. More details about the EcoGrid 2.0 market can be found in the D.2.2 EcoGrid 2.0 Market Specification project deliverable (J. Mehmedalic 2016).

In the project, aggregators of residential thermostatically controlled loads (resistive heaters and heat pumps) are considered as the flexible demand. These aggregators can offer flexibility products to the TSO and the DSO concurrently. A major project objective is to validate if aggregators can offer and provide such products reliably, to verify the offered services are delivered and test the whole process under real conditions (bidding, market operation and clearing, service delivery and verification). To do so, a large number of controllable loads located on the island of Bornholm are used, as well as an actual market platform.

Aggregators must be able to characterize and quantify their available flexibility, in order to offer services to the power system (both on a TSO and DSO level). The purpose of this deliverable is to describe how the aggregated flexibility of a pool of thermostatically controlled loads will be characterized in the project, under the realistic conditions of an actual system setup.

2 Flexibility model description

2.1 Introduction

Flexibility characterization is done by a series of tests upon the portfolio, where the aggregated loads' response and the associated energy payback (durations and magnitude) are measured. A saturation curve can then be obtained, an example of which is shown in Figure 1. This curve shows the capabilities of the aggregation in terms of adjusting its consumption using asymmetric blocks; this method of describing flexibility was introduced in (N. O'Connell n.d.). When heating loads offer some flexibility by changing their consumption, they need to afterwards recover by changing their consumption in the other direction. As seen in the figure, the curve on the negative power values shows the possible combinations of load decreases and the respective time durations that are allowed by the load without violating its energy constraints. Each point (response/duration) on the curve must be followed by a point of the second curve, which gives the allowed combinations for the rebound. A chosen combination of a point in the positive curve and a point in the negative curve mark an asymmetric block, which can be submitted to the regulating market.

Even though these continuous curves are convenient for illustration purposes, in practice they consist of discretized points, which is the form that will be used for the flexibility model in EcoGrid 2.0. Therefore, we refer to the response/rebound sets (power and duration for both parts) as **asymmetric blocks**.

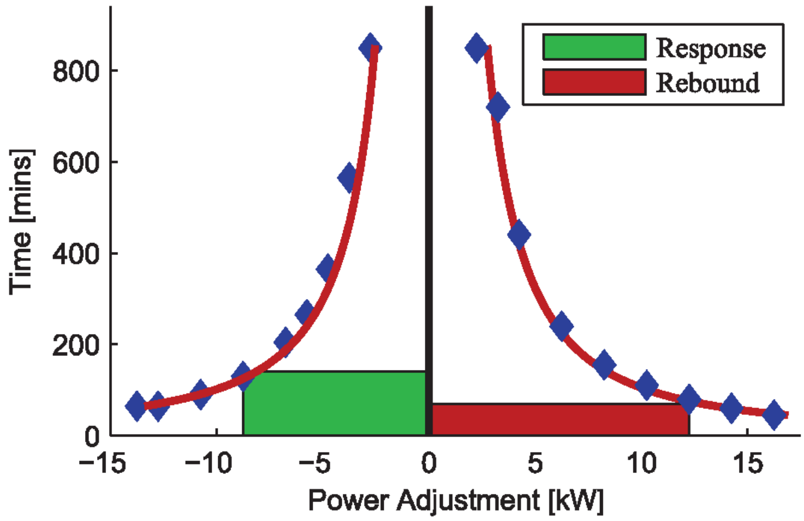

Figure 1: Example of a saturation curve taken from (N. O'Connell n.d.)

This characterization is dictated by 3 main considerations. First, the low controllability and observability capabilities of the controlled loads require the development of relatively simple bidding and control strategies and therefore a data-driven approach was adopted for the project. Instead of using analytical models to describe the load dynamics, their behaviour is observed under various tests and the relevant quantities are calculated from the metering data. Second, thermal loads are energy-constrained units and need to recover their state of charge (temperature) after increasing or decreasing their consumption. The markets developed in EcoGrid 2.0 facilitate asymmetric block offers, making saturation curves an appropriate and simple, yet not oversimplified, method of describing the aggregated flexibility. Third, the authors of (N. O'Connell, 2016) showed that using simple saturation curves and sets of asymmetric block offers obtained by experiments yields similar earnings compared to an analytical load model with full information.

The saturation curve obtained by (N. O'Connell n.d.) is based on a continuously-modulated refrigeration system and not an aggregation of on-off loads, as those available in EcoGrid 2.0. This poses a number of challenges and several factors, which are later discussed in more detail, have to be considered in order to derive the appropriate saturation curve. Moreover, in contrast to a refrigeration system where the ambient temperature does not vary considerably, the response of many thermal loads may be significantly affected by external temperature, solar irradiation and other factors, which requires an extension of the saturation curves.

One of the contributions of EcoGrid 2.0 is to propose and validate a way of describing the aggregated flexibility of thermal loads via asymmetric blocks under variable external conditions, such as ambient temperature and solar irradiance.

2.2 Overall setup and involved loads

The EcoGrid 2.0 project investigates the use of residential loads flexibility to provide both transmission and distribution system services, as described in the Introduction. The considered loads which are controlled to offer these services are thermal loads: electric heaters which are controlled by changing the thermostats' setpoint values and electric heat pumps, which can be only throttled, i.e. be forced to switch off when this is allowed by their local controllers.

A particular challenge stemming from actual implementation restrictions is that even though only the electric heating of a house may be controlled, it is the total net consumption meter data as measured by the smart meter that will be available. This includes the consumption of the controlled loads, all the uncontrolled loads of a house (e.g. lights, refrigerators, electric ovens etc.) and the photovoltaic panels production, if any are installed. The absence of a separate meter for the controlled loads may reduce the installation and maintenance costs but poses many challenges, especially in estimating the actual response of the loads and therefore verifying system service delivery.

2.3 Baseline estimation and forecasting

The baseline consumption of a group of consumers with DERs is the estimated total, aggregated load of the portfolio assuming that the DERs are not controlled and is provided by the **Baseline Estimation Tool**. It serves two purposes. The first is to forecast the uncontrolled consumption of a specified set of DERs based on a number of inputs, such as weather forecasts, time of the day, whether the day is holiday or not etc. Note that the baseline model uses weather forecast values for creating the underlying model, i.e. for training the neural network, and as an input when it is required to provide a baseline estimation. This forecast is useful for the aggregator in order to estimate the uncontrolled consumption of its portfolio. The second purpose is to estimate the baseline of the DERs when they are controlled, which corresponds to *the hypothetical consumption of the DERs if they had not been controlled*. This estimation is particularly important for evaluating the impact of the control actions imposed on the DERs, as well as for service verification.

The baseline provided by the tool is the expected value of the consumption in the period with a range of specific confidence intervals. A sketch of a service provision and a baseline is shown in Figure 1. For a more comprehensive analysis of how the baseline is derived and for more technical details, the interested reader is referred to the D.4.1.1 Toold for market interaction and service delivery verification project deliverable (C. Heinrich 2017).

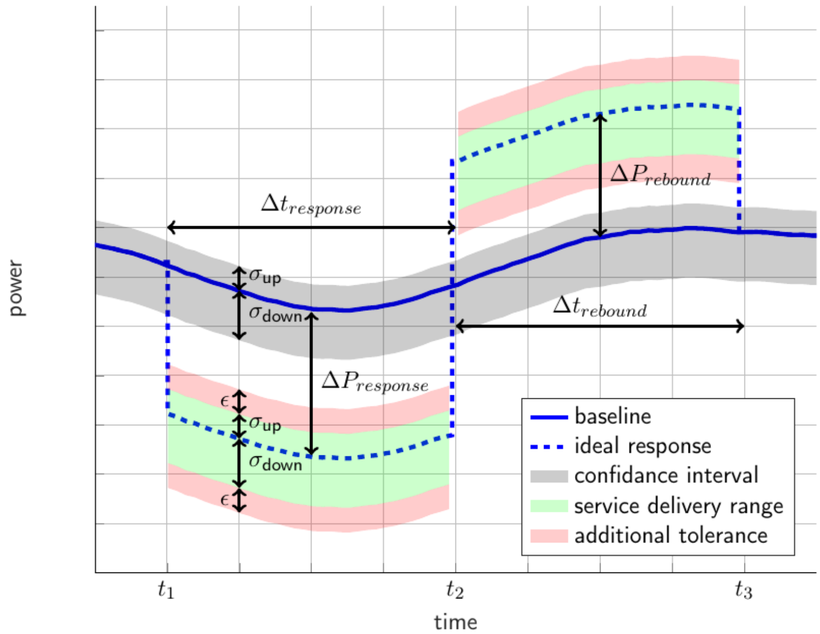

Figure 2: Example of a baseline with a confidence interval and a response/rebound block as an offered service

2.4 Tests description and evaluation process

The tests involve an initial load increase or reduction, both of which result in an associated rebound of decreased or increased consumption respectively. The rebound effect after the initial response is commonly referred to as the payback consumption. Given the time resolution of the smart meters and the flexibility market, which is equal to 15 minutes, multiples of this time interval will be used for the regulating power bids. As seen from the saturation curve, an *ideal* response consists of a constant response (-8 kW) and a constant rebound (12 kW) in terms of power. Indeed, the asymmetric block offers consist of response and rebound power values which are constant during the two phases (even though the response and rebound magnitude can be different).

A simple test could involve switching all loads off (by reducing their thermostat to a very low value) for one hour and then resetting their thermostats to their original values. The response of the aggregation is derived by subtracting the baseline consumption from the actual metered consumption.

By conducting a one hour load reduction test, it is likely to observe a load reduction which is not constant throughout the response period or a rebound which is also not constant in terms of power. The average power of the response and the rebound will be considered, in order to derive points in the saturation curve. While the response period is determined by the test itself, the rebound period must be calculated. The rebound period will be considered over when the observed load difference lies within the baseline uncertainty range, which means that a statistically important load deviation is not observed. A function, named **Test Evaluator,** processes the test results and evaluates if the response is acceptable or not, based on the services specifications, and provides the following values for each block: response and rebound power, as well as rebound duration.

From the described test evaluation it is evident that the produced blocks are estimates of the actual behaviour of the system. Even if the loads could be perfectly controlled, the presence of the baseline estimation introduces an uncertainty. *Therefore, the resulting blocks are simply observations of a stochastic process and should be treated as such*. It is thus important to repeat the tests many times, in order for the mean values of the observed blocks to converge to their real ones.

Remark*: if the response and rebound are considered unacceptable, then alternative control strategies can be tested, which will result in responses closer to the ideal ones. This may be relevant for the electric heaters, where an immediate and high rebound effect is expected. In that case gradual release of the loads or other control methods can be tested.*

2.5 Handling missing meter data and ICT failures

After conducting a test the meter data must be analysed in order to evaluate the loads behaviour and obtain one more set of response/rebound values. However, it is likely that meter data will be missing for some of the houses during the test. In order to evaluate the actual loads behaviour we **exclude** those houses. Moreover, there are ICT failures which affect the response of an aggregation during a test. Activation failures result in the underutilization of the aggregation's flexibility and a large failure rate will severely affect the capability of the loads to deliver their full potential.

There are two ways to handle ICT failures when using the test results to evaluate the loads behaviour. One is to embed these failure rates in the stochastic overall response of the loads. Since failures occur due to the imperfect ICT infrastructure, these can be considered as part of the response of the loads and their effect will be indirectly accounted for when calculating the response/rebound blocks in the saturation curve. Note that these blocks are merely observations of the tests; therefore, the reduced flexibility due to ICT issues will be **implicitly** captured by blocks of smaller responses (proportional to the failure rate).

While this procedure is straightforward, significant differences in the ICT performance over the testing periods (for instance due to some technical problems related to the communication media, databases, servers etc.) will result in a large variation in the response of the loads. Moreover, it is more difficult to estimate the performance of the loads under different ICT failure rates, for instance if failure rates are drastically reduced.

Therefore, ICT failure rates will be **explicitly** accounted for by excluding houses which failed to respond. Then the response and rebound will be calculated only for the houses which responded to the control signals and their number will be included in each calculated response block. Calculating the actual response of the loads will allow the use of extrapolation for scaling the response to the initial aggregation size; furthermore, it will be easier to estimate the loads response under different ICT failure rates, which may change over time. For each test the failure rate will also be calculated, so as to derive an average value of the ICT performance, which will be used in the control of the loads.

2.6 Test evaluation example

In Figure 1 the evaluation of a hypothetical test is illustrated, in order to demonstrate the various steps of the process. An initial population of 200 loads is tested, with 22 loads failing to respond, resulting in a failure rate equal to 11%. 30 houses are excluded since they have either failed to respond or meter data is not available. Note that in the general case a house without stored meter data may have also failed to respond due to ICT issues. Next, the baseline for the remaining 170 houses is calculated by using the baseline estimation tool (with any additional input required).

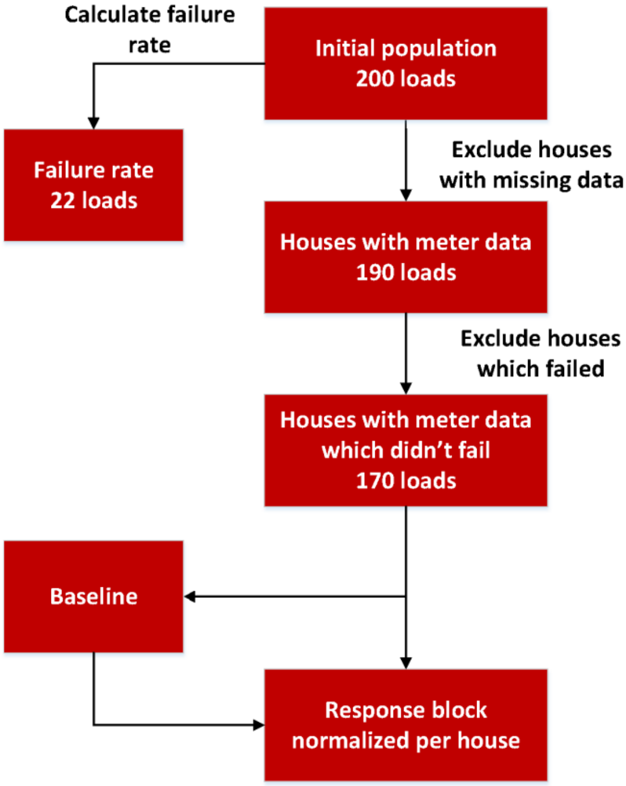

Figure 3: Process of deriving a block response from a test

Using the baseline, the response and rebound of the loads are calculated and normalized to the number of houses which didn't fail and delivered meter data, so that all tests can provide comparable results. This procedure will result in 4 values for each test:

- ✓ Response Power per house (kW)
- ✓ Response Duration (no of 15-min blocks)
- ✓ Rebound Power per house (kW)
- ✓ Rebound Duration (no of 15-min blocks)

There are, however, two more important aspects related to each test. One is the external conditions when the test was performed; for instance the ambient temperature, the time of the day and the solar radiation; these factors add extra dimensions to the saturation curves.

The second aspect is the control applied to achieve the response. A test may be more complicated than a simple throttle test, where loads are simply switched off for a specific duration and their rebound is observed. Due to the nature of the control (open-loop), the control applied to the loads in order to achieve a specific response may be derived by using a variety of methods. This can be difficult to formalize and be used in the control phase without any information from the flexibility model. Therefore, the applied control must be included in the derived asymmetric blocks and the flexibility model.

Once the response is obtained (by subtracting the baseline from the actual aggregated metered consumption), it will be an input to the "Test Evaluator" function, which will process the test and return the block's aforementioned 4 values (if the test is considered successful); the block will also contain the external conditions during the test and the applied control.

3 Flexibility maps and connection to other functions

3.1 General information

As mentioned in the previous section, each test that was considered **successful** will result in one block, with certain information fields. In Table 1 the structure of this information is shown.

Block No	Response Power (kW)	Response Duration	Rebound Power (kW)	Rebound Duration	Applied Control	External Factors

Table 1 – Information obtained by each test

Once more test data is available, it will be clearer if and which external factors affect the loads behaviour. Examples could be the ambient temperature, the baseline consumption, solar irradiance, time of the day etc. Moreover, if tests with more complicated control signals are carried out, for instance gradual activations and releases, the exact format of this control information needs to be determined and stored. This table is referred to as the **flexibility map** and will be used during the bidding and control phases.

3.2 Connection with the optimization strategy tool

The flexibility model is used to provide the blocks with which the aggregator will bid in the regulating markets. To do so, the optimization strategy tool will request the possible blocks from the **control bridge**, which will provide the blocks using the flexibility model and other input, which is closely tied to the "external factors" field of the flexibility model. This input will be used to choose the blocks from the flexibility maps which are the closest to the current conditions.

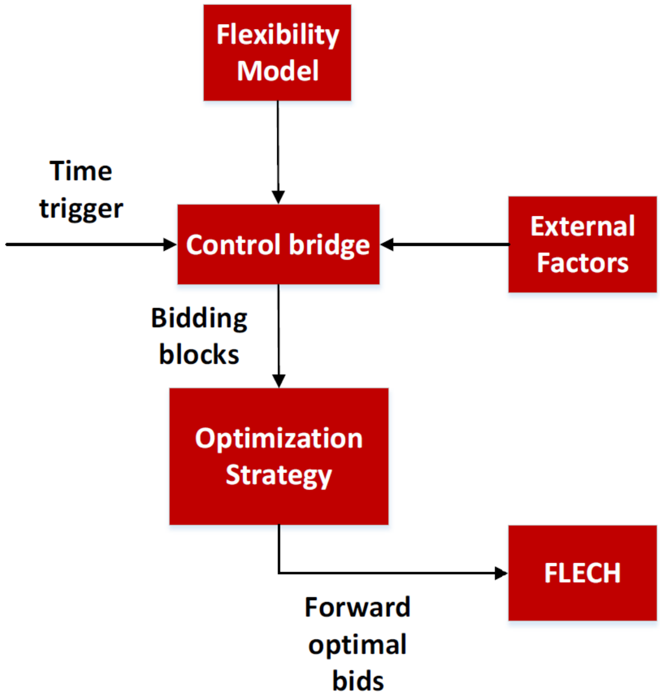

Figure 4: Use of the flexibility model for bidding in the regulating market and associated interactions

Block No	Response Power (kW)	Response Duration	Rebound Power (kW)	Rebound Duration
1	0.5	4	3	3

Table 2 – Block information passed to the optimization strategy tool with an example

3.3 Connection with the controller

Once a specific bid has been accepted, FLECH will pass the information of the bid acceptance to the aggregator interface (control bridge). Then, control logic will retrieve the information from the flexibility model in order to perform the service, i.e. how exactly to deliver the asymmetric block. This information is passed to the controller, which has the task of dispatching the control signals to the units.

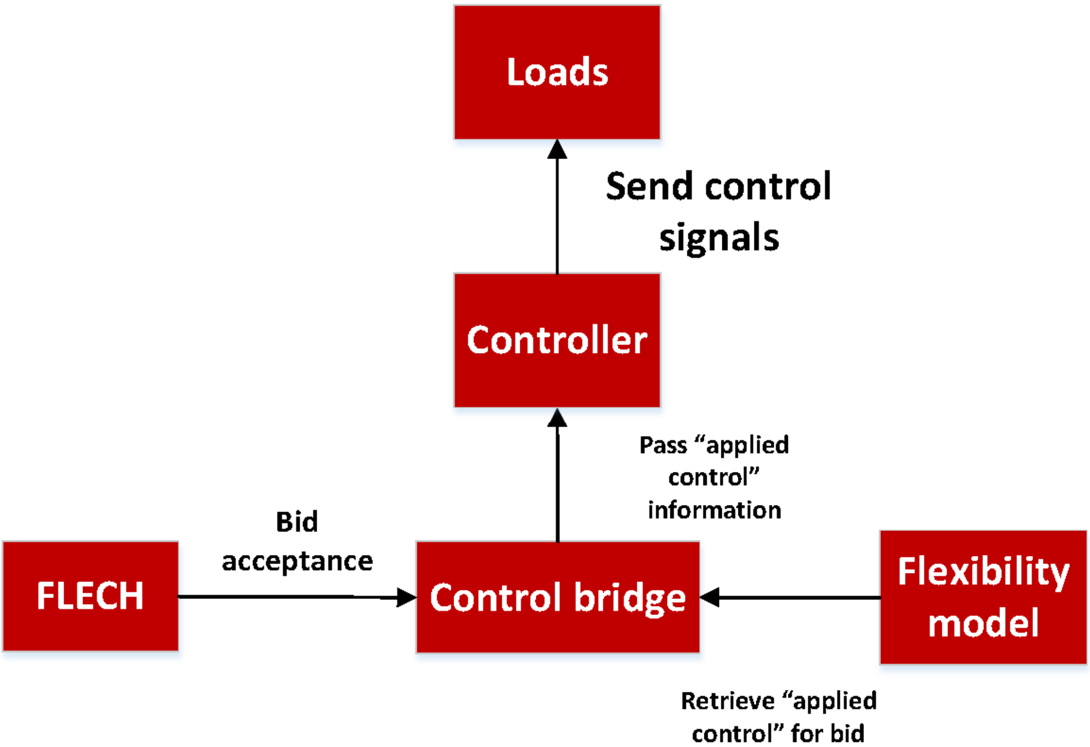

Figure 5: Use of the flexibility model for delivering a service

4 Tools description and development process

In this section the functions that are required in order to build the flexibility model are presented. First, data processing is required to fix time issues of weather/meter data and then to clean the data from any unrealistic or missing values. Afterwards, the updated clean data is produced, which will be used for training the neural network model for the baselines. In Figure 1 the data-cleaning process is outlined, where the orange boxes indicate data, the red box indicates the functions and the blue the functions' output.

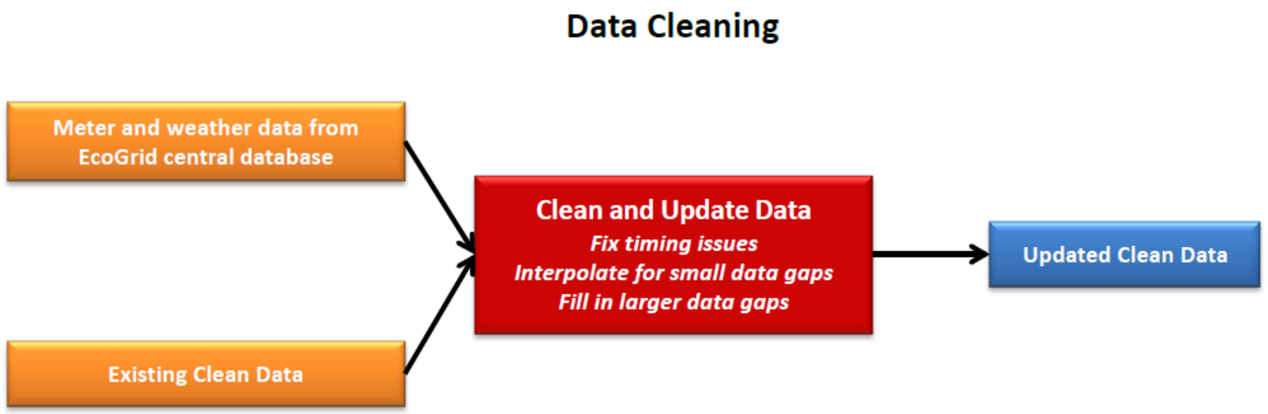

Figure 6: Process of cleaning and updating data

Once the clean data has been updated, it is possible to create a baseline generator. A neural network model is trained upon the clean data of a specified list of DSO numbers, as shown in the following figure. It is important to note that the baseline is created for a specific list of houses, which corresponds to the part of the portfolio whose response is to be quantified.

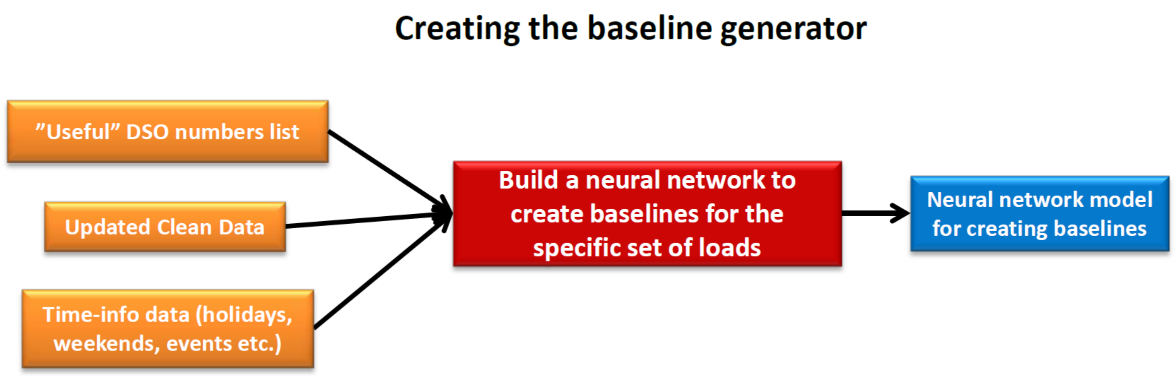

Figure 7: Process for creating a baseline model

Next, tests can be evaluated and the flexibility map can be updated. In the following figure the various steps of the tests evaluation process are illustrated.

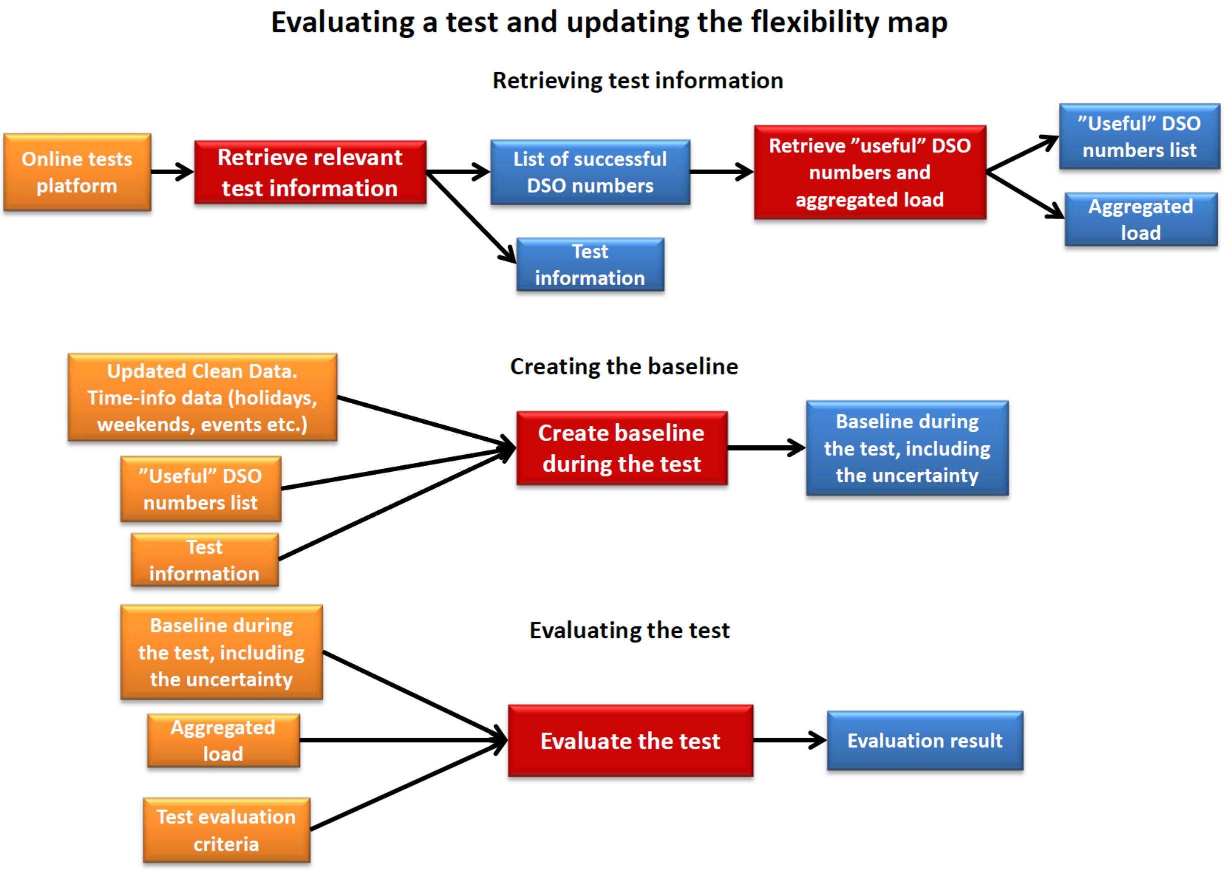

Figure 8: Process of evaluating tests and updating the flexibility map

First, all the relevant information regarding a test is retrieved from the online tests platform, such as the time period, the households which completed the test successfully etc. For these DSO numbers, a subset, referred to as "useful" DSO numbers (which do not have missing meter data during the test), is obtained and a baseline is created. The aggregated load during the test for these DSO numbers is retrieved from the central database and by using the baseline and test evaluation criteria the test is evaluated and the flexibility map is updated accordingly.

5 Bibliography

C. Heinrich, C. Ziras, E. M. Morales Bondy. »D4.1.1 Tool for market interaction and service delivery verification.« 2017.

J. Mehmedalic, E. Mahler Larsen, D. E. Morales Bondy, A. Papakonstantinou. »EcoGrid 2.0 Market Specification.« 2016.

N. O'Connell, P. Pinson, H. Madsen and M. O'Malley. »Economic Dispatch of Demand Response Balancing Through Asymmetric Block Offers.« *IEEE Transactions on Power Systems, vol. 31, no. 4, pp. 2999-3007, July 2016.*, u.d.

Read more at www.ecogrid.dk