Chapter 1

Case study of Aran Islands: Optimal demand response control of heat pumps and appliances

Marko Jelić[1] Dea Pujić[1] Nikola Tomašević[1] Paulo Lissa[2,3] Dayanne Peretti Correa[2,3] Marcus Keane[2,3,4]

Demand response has proven to be a crucial mechanism in the process of flexibility exploitation on the demand side. Throughout the years, it has evolved and expanded, reaching more and more previously untapped potential sources. In that process, residential users have provided a significant buffering capacity for balancing energy production and demand, but this came with a few challenges. With more and more households transitioning from being purely energy users to smart homes and energy prosumers with distributed renewable energy generation, new possibilities have opened up for integrated optimisation approaches that make the best use of both locally generated and grid-supplied energy as well as energy storage systems.

1.1 Origins of demand response programmes

The global share of energy consumption, as analysed by the United Nations Environment Programme [1], can be disaggregated into different sectors. According to this report, 30% of final energy consumption and 28% of CO_2 emissions can be attributed to buildings. Interestingly, electricity consumption in building operation is said to represent around 55% of global electricity consumption. When specifically looking at residential buildings, they are reported to contribute 22% of final energy consumption and 17% of CO_2 emissions. Therefore, it goes without saying that any reduction in energy consumption or increase in energy efficiency in the residential sector goes a long way towards fighting the ongoing climate battle.

Several elements have been noted in literature as key enablers of a transition to a greener future by analysing the problem of decarbonisation of the energy system at different scales, from the entire energy sector [2] to small and isolated grids that can be found on geographical islands [3]. Of these, two are crucial for residential

[1]Institute Mihajlo Pupin, University of Belgrade
[2]College of Science and Engineering, National University of Ireland, Galway
[3]Informatics Research Unit for Sustainable Engineering (IRUSE), Galway
[4]Ryan Institute, National University of Ireland, Galway

energy use optimisation. First is the increasing prominence of distributed renewable energy generation (photovoltaic panels, wind turbines, etc.) and implementation of highly efficient sustainable technologies (like heat pumps) to replace legacy devices with large carbon footprints. However, with intermittent renewable sources becoming more commonplace, the delicate balance between the supply and demand has been jeopardised. Therefore, the second enabler is the utilisation of demand-side flexibility to aid in sustaining this equilibrium. Although a comprehensive classification of different mechanisms by which demand-side flexibility can be exploited is not yet well established within the related literature, it is generally well understood that demand-side management (DSM) and demand response (DR) are the crucial instruments in this domain. Although often used interchangeably by mistake, these two terms signify different approaches to load modification. DSM generally depicts long-term efficiency improvements that, overall, result in load reduction over time and aid the process of achieving full energy autonomy. On the other hand, DR refers to short-term load modifications that are made in reference to external impulses or incentives, thus helping maintain the stability of the wider power supply system.

As estimated by [4] at the time, 20% of power generation capacity was utilised only to fulfil peak demand levels which were present only about 5% of the time. This discrepancy has resulted in high operational costs of the power supply network as well as negative implications on emissions. However, by employing mechanisms such as DR, load levels should be able to exhibit more flexibility and, as a result, allow for easier balancing with the supply. Load modification techniques such as ”peak curtailment/shaving” and ”valley filling”, as review by [5], can produce more balanced load curves that are less challenging to match with appropriate generation facilities. DR is an especially important tool in this regard as it can guide load modifications using its several different variants, as will be discussed in the following sections.

1.1.1 Traditional (industrial) DR applications

The first implementations of DR programmes can be traced back to the latter half of the 20$^{\text{th}}$ century. At that time, attempts at DR integration were primarily focused on large commercial and industrial customers. They were selected mainly due to an already present high level of automation as this ensures easier control of assets without additional devices and retrofits. Furthermore, industrial customers were also able to provide a significant amount of flexibility which made contracting them much more efficient as opposed to individual residential consumers which require some form of aggregation to provide a noteworthy impact.

Initial approaches were based on direct load controls through so-called “explicit DR” which entails that the flexibility provider offers direct control over some of their assets at predefined time intervals and frequency while being offered monetary reimbursements in line with the provided capacity. This system has allowed utilities to make use of a portion of industrial demand levels as a buffer. In time, this process has evolved into a so-called “implicit DR” where the exchange between demand flexibility and monetary reimbursement is conducted via a variable energy price tariff. By alternating between low and high price intervals, the periods during which demand should be increased or decreased are implicitly encoded. When optimising their demand, users essentially attempt to reduce their operating costs by aligning the demand with the tariff profile, with resulting savings representing the previously directly agreed upon monetary reimbursement.

1.1.2 Transition towards the residential sector

Understandably, the concept of direct control over household appliances in the residential sector, even if the technical challenges of deploying the necessary equipment are ignored, is met with resistance by dwellers. Due to specific aspects of how human behaviour influences energy consumption habits, different attempts to utilise specifically price-based DR approaches have shown positive results in this domain [6]. Arguably, the most prominent implementation of implicit DR which has been in use for some time are time-of-use tariffs (also commonly implemented as night/day or peak/off-peak tariffs). However, the inclusion of distributed renewable generation as well as various controllable devices, especially as more and more households embrace the concept of smart homes, calls for an integrated approach to load modifications based on current conditions in order to maintain effectiveness in providing a balance between the supply and demand. One such solution, along with a set of results from a real-world use case, will be presented in this chapter.

1.2 RESPOND control loop and methodology

In order to provide a holistic solution to the problem of energy management for residential smart homes, the answer provided by the consortium of H2020 RESPOND project[5], depicted in this chapter, utilises a set of smart services in conjunction with edge sensors and actuators to facilitate efficient day-to-day operation. This section presents different components of the proposed platform which, through synergistic operation, aim to integrate DR-supported optimisation into the operation of appliances, storage systems and heat pumps.

1.2.1 IoT backend platform

Having in mind the various types of data that need to be processed and stored in order to facilitate the operation of an Internet of things (IoT)-based system for home energy management, a complex heterogeneous platform for data handling and management was deployed as one of the primary components for this system. The RESPOND IoT platform was composed of a set of various data repositories:

- Semantic repository which contained metadata regarding users, sensors, equipment such as: characteristics of the photovoltaic panels (total capacity, slope, etc.), energy storage (battery capacity, maximum charge and discharge rates, etc.).

[5] www.project-respond.eu

- Influx database (DB) which contained timeseries measurements from the field devices such as electrical demand, renewable production measurements, temperature measurements, etc.
- Relational database which was used as an intermediary log for interaction between different services, as a repository for data to be shown in the accompanying mobile phone app as well as for user management for platform access.

All of these data stores were connected to the field level equipment via a Message Queuing Telemetry Transport (MQTT) broker. Since the integration point between different components of the system is located within the cloud platform, each service is envisioned to be able to obtain the required input parameters from the platform, as well as to store back the outputs into the corresponding data repositories. Even external services like weather forecasting which are utilised within the system are integrated with the aforementioned data repositories.

1.2.2 Forecasting services

In order to be able to shift the demand and adapt it accordingly, as explained in the first section of this chapter, it is of utmost importance to have an estimations of the expected, baseline, renewable on-site production and demand in the first place. Hence, within the RESPOND project, both production and demand forecasting models have been independently developed with the goal of providing predictions for the same horizon and resolution which will be exploited by subsequent services like the optimisation. Since, at the time of development, historical production data was lacking while there were sufficient logs of previous demand measurements, the models were developed to provide 24 hour-ahead hourly forecasts with the demand forecasting model being data-driven and the production forecasting service realised using physical models.

The production forecasting service is envisioned to map predicted meteorological parameters to the expected production of available renewable energy sources (RES). In the particular use case that is presented in this chapter, the pilot site was equipped with solar photovoltaic (PV) panels. The inputs for the corresponding physical model, presented in [7], can be classified into one of the following two categories:

- **Dynamic parameters:** in order to be able to provide the expected production, external meteorological conditions are required. Therefore, as inputs for this model, global solar radiation and cloud coverage were necessary. Additionally, apart from the weather, temporal parameters are correlated with the production, as it is highly dependant on the instantaneous solar position. Therefore, current time and date are included as the inputs, as well.
- **Static parameters:** as usual when utilising physical models, apart from the dynamic parameters, which are also common for data-driven models, the following physical and geographical parameters were necessary: slope and azimuth of the PV cell surface, temperature coefficient and surface area of the PV cells, rated capacity of the PV array, nominal operating cell temperature, longitude, latitude and time zone offset.

When performance of the utilised model is considered, it achieved a mean absolute error of 21% and, as was expected, was not as accurate as machine learning (ML) models would be if historical data had been available. Nevertheless, it was precise enough to be utilised as an indicative input for optimisation purposes.

Regarding demand forecasting, a couple of different ML models were tested and the k-nearest neighbour (kNN) algorithm was chosen for utilisation since it had the most accurate predictions. Similarly to the production forecast, the demand forecasting service was designed to provide 24 hour-ahead forecast with an hourly resolution. As explained in more details in [8], the kNN model predicts expected demand depending on the previous demand and a set of the time variables extracted form the date and time such as day of the month, the season, the day of the week and a Boolean variable indicating whether it is a working day or not. The previous electrical load is obtained from the Influx DB, whilst the outputs of the service are stored in MySQL DB, as was the case with the production forecaster. Since the COVID pandemic has impacted the validation period of the system, the models has been adapted accordingly [9]. This improvement was necessary since household electricity demand has changed as a result of the fact that most of the users started working from home.

1.2.3 Optimisation services

Constantly analysing the expected renewable generation and attempting to align the demand in accordance to its profile, as well as various DR requests, is a cumbersome task that very few residential users want to constantly take upon themselves to resolve. In order to ensure cost-effective and energy-efficient operation, significant efforts have to be invested in order to make best use of, for example, varying energy prices, or local energy generation and storage. Therefore, the optimisation services within the RESPOND control platform was envisioned as an integrative component that would be capable of assessing multiple different aspects of the energy management problem and automatically providing the best course of action by analysing:

- The arrangement of the underlying energy infrastructure components including all relevant energy carriers, converters and storage systems of individual energy prosumers.
- The forecasted production profile from all locally available renewable energy sources obtained by the corresponding forecasting service.
- The forecasted energy demand as well as corresponding load flexibility constraints obtained by the corresponding forecasting service.
- The limitations of the grid connection and other components of the system.

In order to make the best use of contemporary computation power while simultaneously guaranteeing that all available resources are used in the most efficient manner, adequate models of the energy systems are built and optimised using an appropriate solver engine. The methodology that was utilised within the control platform is based on the Energy Hub modelling approach. Originally presented by [10] with various subsequent implementations in literature, as revised in [11], owing to its flexible nature, the Energy Hub can be utilised in a variety of problems, from single carrier electric energy systems [12] to complex multi-stage hybrid systems that involve, for example, both electric and thermal domains with adequate converters [13].

Following this concept, a corresponding model was developed for the use cases that will be discussed in the following sections of this chapter. Since the optimisation in this case focuses on the electric domain as it is supplemented with smart sensors and actuators that are integrated in the platform, the model depicts different electric energy sources, a battery storage system (where applicable), converters and loads, as illustrated in Figure 1.1 through different stages. In accordance with this structure, a corresponding set of variables, bounds and constraints can be derived and implemented as a mixed-integer linear programming (MILP) problem. This choice allows for the energy flow for each model to be efficiently evaluated in sub-second times facilitated by the simplicity of a MILP model that does not require the implementation of numeric solvers as this type of precision is generally regarded as unnecessary in similar modelling problems.

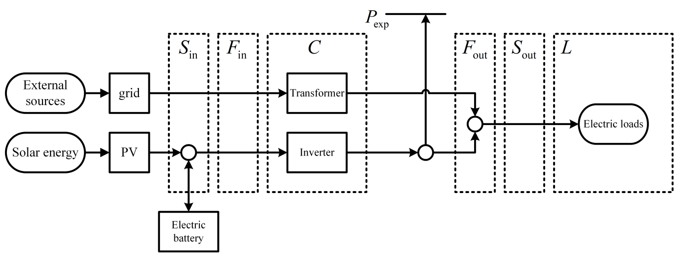

Figure 1.1 Illustration of the Energy Hub layout of a house in Aran Islands

The forecasting services, which are to be evaluated before the optimisation, define a portion of the variables of the model. For example, the RES forecast depicted previously defines the available energy from the PV array while the demand forecast provides a reference based on which some form of demand flexibility can be implemented. In line with general findings from [14], the overall flexibility in a community-oriented project is estimated at around 20% while a use case study depicted in [15] hits at the possibility of peak load reduction of 30% with extreme energy tariff manipulations. These findings provide a range of values that can provide context for the flexibility margin around the forecasted load profile in which upwards and downwards load modifications are made.

An important feature in the optimisations service is the integration of both explicit (direct) and implicit (price-based) DR. Namely, since the operation of each mode is optimised with operational costs set as the main criterion, the output reflect the most cost-effective solution with varying energy prices in mind. With implicit DR natively supported using this setup, an additional term in the criterion is added such that, if an external grid-side entity or aggregator wishes to request a certain demand level at a predefined time interval, the deviation of the output demand curve is separately and highly penalised. This is done such that the output optimal demand reflects the required profile as closely as possible, thus also facilitating explicit DR requests.

Finally, the outputs of the optimisation service are comprised of a set of optimal energy utilisation curves that reflect when and how much energy is to be stored, imported, exported and consumed. In accordance with these profiles, subsequent services depicted in the following sections will provide means of converting power consumption curves into concrete control actions depicting when to schedule appliance usage and how to set the references for heating, ventilation, and air conditioning (HVAC) devices.

1.2.4 Control services

After running the optimisation service, the system is provided with optimal energy utilisation curves as power values through time. However, this format of data usually cannot be considered to provide useful information without being processed first. This process entails the conversion of these curves into discrete ”turn on”/”turn off” instructions for appliances and set-point values. These two processes will be further discussed in the following sections.

1.2.4.1 User recommender service and appliance controls

The first component of the control service mainly focuses on the electric domain and the conversion of optimal loading profiles into appliance use schedules. An approach for solving this problem, outlined in [15] and further explored in [16], makes use of a heuristic tabu search method. Namely, by discretizing the appliance usage schedule, the algorithm looks for the best arrangements of their activity in time such that this schedule, in conjunction with the fixed demand curve, results in a closest match to the demand curve that is deemed optimal. However, this approach only tackles the electric domain while applying it to the thermal one would be a much more challenging task as it would require the use of complex models. This issue is precisely what will be further discussed in the following section.

1.2.4.2 Building models and HP controls

To create a simulated building model, a site survey has been conducted to gather data related to construction characteristics, such as type of the walls, windows and roof. Furthermore, sensors have been installed, measuring indoor temperature (°C), total electricity consumption (kWh) and heat pump electricity consumption (kWh). The collected data was used to develop a detailed and calibrated white-box model, using the Integrated Environmental Solution Virtual Environment (IESVE) software. Next, simulations have been carried out to identify the main parameters and heating transfer dynamics necessary to build a reduced grey-box model. The parameters extracted from the white-box model were indoor air temperature increase and decrease rates for both, domestic hot water (DHW) and indoor temperatures, considering their behaviour during stationary conditions (system off) and when actions are performed (indoor heating or DHW on). Additional information about the house parameters and white-box model calibration, including the validation metrics applied, can be seen in [17]. Once developed, the building simulator is able to receive the optimisation and forecasting services as input and translate them to optimal heat pump control actions, as can be seen in Figure 1.2.

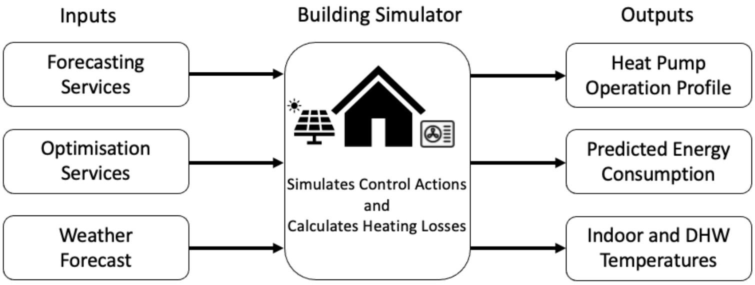

Figure 1.2 Building simulator process.

The building simulator reads the current environment state every 5 minutes and estimates the next DHW and indoor temperature values, calculated based on the rates established in the white-box model. This new environment allows for a more flexible framework where different control techniques can be tested. The forecasting and optimisation services, along with the weather forecast, are used as inputs, so they can be part of the DR control strategy. For instance, a control action can be scheduled to activate the heat pump when the PV production forecast will be higher in the next day. The output of the building simulator is a heat pump control operation profile, which indicates the control actions be to executed the next day. Furthermore, it calculates the predicted energy consumption and the expected indoor and DHW temperatures over the day.

1.3 Use case setup

The location of the experiments is Inishmore, the largest of the three Aran Islands, in Galway, Ireland. With a population of approximately 800 people, the island itself is very exposed to the weather elements, particularly during the winter months as it has very little shelter. The islands are connected to the mainland through a sub-sea cable, in which 1.855 MWh of electricity was imported in 2017 [18]. In 2016, a fault in the sub-sea cable resulted in a power outage on two Aran islands, that lasted for four days and affected approximately 400 residents [19]. On that occasion, some islanders had to rely on local diesel-powered generators. This event showed the islands' vulnerability and dependency on the main island generation, leveraging the need for new reliable on-site solutions.

1.3.1 Pilot installations

In the Aran Islands, there is a potential of 450 dwellings that share similar characteristics in terms of construction materials. Moreover, as they are geographically close to each other, the external environment conditions do not vary considerably across the buildings, hence the heating losses dynamics tend to be similar over the day. A total of 9 houses have been selected to be part of the test cases. They already had individual PV production for self-consumption, a heat pump system for indoor heating and DHW, and appliances such as washing machines and tumble dryers. To allow DR capabilities and to take benefit of the services provided by RESPOND, a new set of devices from Energomonitor[6] were installed in each of the houses. The new architecture added smart capabilities for the legacy equipment, allowing for individual load measurement and control, besides bringing monitoring of room temperatures and CO_2 concentration. The description of the deployed devices and their application can be found immediately below:

- External meter interface (Energomonitor Optosense): It measures electricity consumption or production by reading the optical impulse output of a digital electricity meter. Application: Electricity meter.
- Electricity wire sensor (Energomonitor Powersense): It measures electricity consumption or production by induction coils installed on 1 or 3-phase wires leading to the main breaker cabinet/panel. Application: Heat pump, PV, and electric vehicle charger.
- Temperature sensor (Energomonitor Thermosense): It is a thermometer for indoor or outdoor use. Application: Room temperature sensor.
- CO_2 and humidity sensor (Energomonitor Airsense): It monitors complex air quality in the room – carbon dioxide (CO_2) concentration, temperature, humidity and noise level. Application: User comfort level measurements.
- Smart plugs (Energomonitor Plugsense): It measures consumption over concrete appliances and can be used to switched them on and off remotely. Application: Individual load control (e.g., dishwasher, washing machine, tumble dryer).
- Gateway (Homebase): It is the heart of the solution Energomonitor, wirelessly picking up data from up to 30 transmitters in the house through encrypted radio protocols. Application: To provide communication between the previous devices and services.

The final list of measurement points per house can be seen in Table 1.1, where each appliance, equipment or sensor can be related to the aforementioned Energomonitor devices, following their specific application. The data gathered from the devices were utilised as input for the services described in sub-section 1.2 and also for user's verification and control, through a mobile application.

Finally, to assess the demand response capabilities and support the validation process, an average dwelling, that was built in the 1970's and has a total floor area of $110\,m^2$, was modelled following the process found in sub-section 1.2.4.2. In recent years, the dwelling has been upgraded, including additional insulation to the walls and roof and installation of an 8.5 kW Mitsubishi heat pump along with a PV panel array consisting of 8 panels, with a total nominal power of $2\,kW_p$. The heat pump connects to a 170 L hot water cylinder which is used to store hot water for both space heating and DHW.

[6]https://www.energomonitor.com/

Table 1.1 Number of measurement points per house.

House Number	**01**	**02**	**03**	**04**	**05**	**06**	**08**	**10**	**12**
Dishwasher	1		1				1	1	
Electricity meter	1	1	1	1	1	1	1	1	1
EV charger							1		
Heat pump	1	1	1	1	1	1	1	1	1
PV panel	1	1	1	1	1				1
Tumble dryer	1	1	1	1	1		1		1
Washing machine	1	1	1	1	1	1	1		1
Temperature sensor	3	5	5	4	5	5	5	5	5
Humidity sensor	2	5	2	1	2	2	2	2	2
CO_2 sensor	1	1	1		1	1	1	1	1

1.3.2 User interface

In RESPOND, a mobile app has been deployed to increase the user's participation in the DR strategies. The app enables users to visualise energy-related consumption and generation, to check comfort matters and status of the devices. Moreover, information from the forecasting and optimisation services can also be visualised, helping user's to make informed control actions. Through the app, users can receive notifications asking to consume more or less energy according to the DR event. Some of the screens available in the user interface can be seen in Figure 1.3.

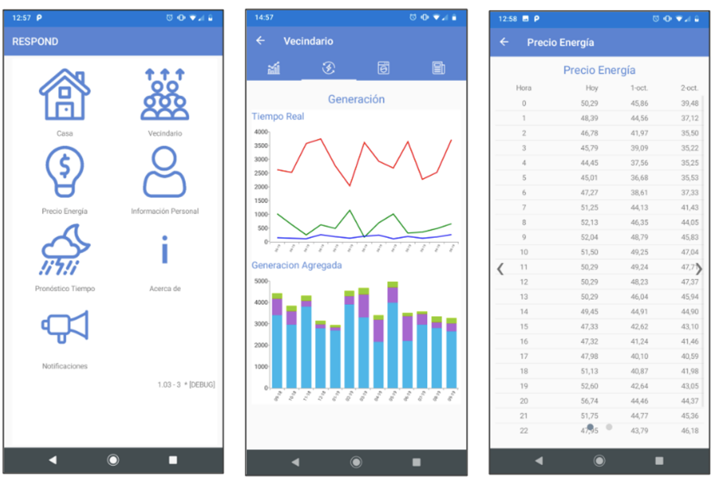

Figure 1.3 RESPOND App - General User Interface.

Detailed information about the mobile app can be found in the RESPOND report [20]. Starting with the main page screen, left screen of Figure 1.3, is shown once the RESPOND mobile app is loaded. This is the main screen where users can navigate and select other different screens to visualise the information available.

For instance, the energy consumption screen shows the recent and historic energy consumption at a dwelling and neighbourhood levels. Furthermore, dwellers can visualise hourly, daily, weekly or monthly energy consumption information. The energy generation screen (centre screen of Figure 1.3) shows users the recent and historic values of energy coming from the PV panels, also with hourly, daily, weekly or monthly resolution. This information may help users to realise the levels of PV production available, and combined with their energy consumption, to raise awareness in the potential reduction of energy coming from conventional non-renewable sources. The energy prices screen, right screen of Figure 1.3, shows hourly tariffs of the energy for the current and upcoming hours and days, so users can decide the best time to use or not some appliance.

Another important information found in the app is the comfort screen, which helps users to verify the indoor environment quality. It provides the mean temperature, humidity and CO_2 levels, which are considered the basic indicators of comfort. Moving to the device list screen, users can monitor and control devices within their houses. Users can then select each of these devices and take different actions. For example, for comfort devices, a user can check the temperature, humidity and/or CO_2 measurements. For appliances, users can check their current and historical energy consumption, as well as activating or deactivating them. The Weather Forecast screen is aimed at showing users the expected weather for the upcoming hours and days. This information, combined with other features such as the energy price, enables users to strive towards more environmentally-friendly and energy-saving behaviour. For example, knowing that the next day is going to be hot and sunny, a user can decide to hang their clothes instead of using the tumble dryer.

Finally, the notifications screen presents the notifications received. Users can receive different types of notifications such as recommendations or even alarms or warnings. These notifications will be received in the mobile app in the form of push messages instantly and also be available in the notifications screen.

1.4 Case studies and assessment

Four different test cases have been designed to assess different types of DR within the RESPOND project in the Aran Islands, considering implicit and explicit models. The test cases aim to exploit as much as possible the benefits of the architecture deployed in the pilot and the optimisation services available, presented in the previous sections.

The test cases results were calculated using the DEXMA Energy Intelligence Software [21]. The DEXMA platform enables real time energy management, with a Measurement and Verification (M&V) tool that contains an automatic baseline calculator [22] fully compatible with the International Performance Measurement and Verification Protocol (IPMVP) [23].

1.4.1 Test case #1

The objective of the first test case was to analyse the impact of the RESPOND smartphone app on user consumption behaviour. The main idea was to understand if, after having access to the information described in sub-section 1.3.2 (e.g. energy consumption per appliance, PV production, etc.), consumers changed their energy consumption pattern in a voluntary manner. For the assessment, comparison of the the period before and after the RESPOND app release has been compared. This test case started to be applied on April 13th, 2020 which was the day participants received their RESPOND app passwords. The final date of verification of voluntary behaviour change was May 30th. As this experiment aim is to verify users' willingness of changing their consumption, there was no notification or other kind of intervention asking them for some energy reduction or increase in the period of evaluation.

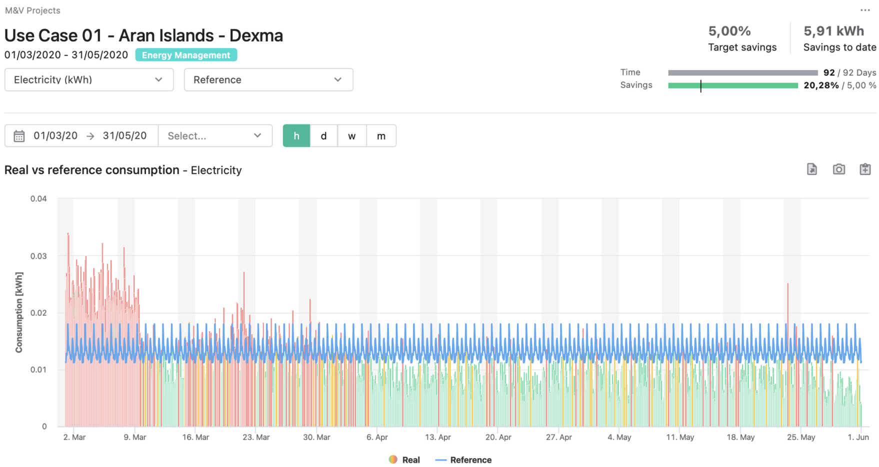

Figure 1.4 Example of M&V output from DEXMA Platform (blue line - baseline, green bar - lower than baseline consumption, orange bar - close to baseline consumption, red bar - higher than baseline consumption)

The baseline period used in this analysis was from March 1st, 2020 until April 12th of the same year. Considering the accumulated values of all dwellings, the results of this test case presented a reduction of 20.28% in energy consumption in the period, compared to the baseline. Figure 1.4 shows the results of the M&V project created inside the DEXMA platform to calculate the energy savings key performance indicator (KPI) in the first test case. The blue line represents the expected consumption over the days (baseline), while the bars are the real consumption. Green bars are values where the real consumption is lower than the baseline, red is the opposite, and yellow means that they are close. The reduction of greenhouse gas emission is also estimated in the platform and in this test case, there was a total of 2.51 t CO_2e avoided in the discussed period.

Performance of the communication infrastructure is a very important topic for guaranteeing the reliability of the test case. With this in mind, some houses were excluded from the validation process due to a lack of sufficient data. The criteria for exclusion in this test case was that houses should have at least 60% of data available in the discussed period, hence three houses that presented lower values were excluded.

1.4.2 Test case #2

The objective of the second test case is to maximise PV self-consumption during periods where there is a peak in energy production. On the day preceding the DR event, the forecasting services estimate the hourly PV production for the next day. With this information, a notification is sent to the participants informing them about the best time to consume energy, if a pre-defined threshold is achieved. Users can get energy savings and also help to reduce peak load in the grid. In the Aran Islands, if PV production is not consumed, it is directly injected into the grid, and users do not receive any payment from the energy provider.

This test case was applied when the prediction of energy production achieves a specific target. The first step was to define which houses were able to participate. Although there are 9 houses in total, only houses 01, 02, 03, 04, 05 and 12 have PV production, therefore messages were sent only to this group. The Irish language, or Gaeilge, is unique to Ireland and it is, therefore, of crucial importance to the identity of the Irish people [24]. To better engage the participants in the actions, the notifications were sent to the participants in English and Gaelic:

"Tomorrow between HH:MM-HH:MM your PV panels are expected to have a period of high production. Try to use your appliances during this period to save money and energy."

"Amarach idir HH:MM-HH:MM meastar go mbeidh do phainéileacha fotavóltacha ag ginniúint roinnt mhaith leictreachas. Déan iarracht do chuid fearais tí a úsáid i rith an am sin chun airgead agus fuinneamh a shábháil."

The application of this test case is classified into one of two experiment periods. During the first experiment period, a message was sent to the customers if the predictions achieved the threshold of 900 W for at least an hour. The demand response events started sending the messages on May 31$^{\text{st}}$. Since messages were not being sent due to the weather conditions in Ireland at that time (PV predictions rarely achieved 900 W), an analysis was performed to define a new threshold that could result in sending notifications 2-3 times per week on average. In August, the new value was then defined as 600 W for houses 01, 03, 04, 05 and 12 and 1100 W for house 02, which has a PV system with higher capacity, effectively defining the second experiment.

The KPI calculations showed a decrease of 6.11% in energy usage from the grid in the first experiment period, and 21.81% in the second period, considering the aggregated values of the participant houses. In the total test case duration, the final result is 17.89% of energy savings. The reduction of greenhouse gas emissions is the energy savings total converted to greenhouse gas emission equivalent and it is estimated that this test case avoided 49.66 t CO_2e over the entire period.

The total renewable energy consumption KPI shows the ratio of the total amount of renewable energy produced and the demand at the event period. Table 1.2 sum-

Table 1.2 Total energy production/consumption during DR events

Date	PV production [kWh]	Consumption [kWh]	% of PV usage
June 1st	7.55	5.66	75%
June 3rd	4.91	3.28	67%
June 5th	5.16	3.10	60%
June 8th	1.48	1.32	89%
Experiment one avg.	4.78	3.34	72.7%
August 6th	2.68	2.73	100%
August 7th	8.49	5.23	62%
August 8th	7.90	5.85	74%
August 9th	2.69	2.30	86%
August 10th	3.39	2.02	60%
August 11th	2.76	2.93	100%
August 17th	2.06	1.68	82%
August 19th	1.79	2.55	100%
August 20th	2.39	2.17	91%
August 22nd	7.61	4.24	56%
August 28th	3.84	1.77	46%
August 29th	0.91	2.02	100%
Experiment two avg.	3.87	2.96	79.6%

marises the amount of PV energy production consumed in each day of first and second experiment. The analysis considers the aggregated consumption and production value of the participant houses. As a result, PV production was 72.7% consumed on average over the first experiment. The second experiment presented an even better performance, where 79.6% of PV energy produced during the event was consumed on average, and sometimes reached 100% of usage.

The baseline used for this use case was from March 1st, 2020 to April 4th, 2020. The main calculations were realised using the DEXMA platform. According to the IPMVP methodology, data backfilling is not allowed [25], and in line with this, periods with missing data were excluded during the calculation process. Outliers, such as accumulated values due to communication issues, have also been removed.

The rescheduled demand KPI aims to verify if the use case helped to move demand into the event period. For instance, if the PV production was higher from 15:00 to 17:00 and the user had received a message, it was expected that a greater consumption during this period and less activity before and after the event would be observed. After the period of the experiments, the real measurements were compared with the baseline. According to the baseline analysis, it was expected that 6.70% of the daily load would be in the event period for experiment one. However, after applying the demand response events, the real data showed 8.71% of the load in the period, which represents a demand increase of around 30% in the event hours. On the other hand, experiment two did not present the same performance, with a 1% of load decrease during the event, compared to the baseline.

The economic savings KPI compares the difference between the average baseline energy cost and the energy cost during the DR event. It considers the amount of energy consumed from the grid, so using appliances when the PV production is higher during the event reduces the final costs. As a result, imports from the grid represented only 4% of the total necessary energy during the event, which is 20% less than the expected baseline. It is important to note that the costs are very related to the way that end users distribute their load over the day. For instance, if the PV production is high and achieved 2 kW, the user has to be aware and try to avoid exceeding this value by controlling the amount of load that uses a certain time to optimise the usage. Otherwise, the amount of energy bought from grid can be greater than expected as all loads are concentrated during the same time interval.

1.4.3 Test case #3

The aim of the third use case is to maximise PV self-consumption by generating an optimal usage profile for heat pumps. This use case can be performed in different ways, such as through fully-automated operation of the heat pumps, which can achieve better energy savings and does not rely on the user behaviour, or manually operated, following a similar methodology as outlined in Test case 2 where users have to play their role and perform the actions when necessary. For remote operations, the installation of an additional device, that is not available in the selected houses, is needed. To assess the potential of this test case, tests have been performed in the simulated environment described in sub-section 1.2.4.2.

The overall methodology is similar as Test case 2, where at the night before the event, the PV production prediction is checked to identify the best period to perform actions (when production is expected to be high). Actions are simulated and the comfort parameters verified in the building simulator. As a result, an optimal heat pump operation profile is generated. In this test case, the focused was placed on optimising the DHW temperature regulation by changing the operation mode and boosting the tank temperature when PV production is higher. This will avoid DHW heating actions during peak consumption hours by creating a buffer of hot water. In the Mitsubishi heat pumps available in the houses, operation mode 1 is the standard mode, where the tank temperatures range from 40 °C to 50 °C. In mode 2, the maximum threshold is increased and temperatures can go up to 55 °C.

As inputs for the building simulator, 10 days in August 2020 were selected where PV production prediction was considered good enough to aid the heat pump system operation. The forecasting services informed about the peak hour of PV production for each of the days. The model was then simulated, prioritising DHW actions mode 2 in these peak generation periods. As a result, an optimal operation mode to be set for each of the hours was generated. The operation modes can then be set manually by the user or autonomously, if technology for remote control is available. Figure 1.5 shows two days of simulations, where the blue line is the tank temperature and the dotted green line is the PV production. The grey bars are the control action for heating the tank. In the first chart, both actions are in periods of

PV production, which occurs only once in the second chart. The reason is that the optimised model also takes into account the minimum setpoint to keep the user's comfort as the main premise, so regardless of the PV production, if the temperature drops to below 40 °C, the system performs a DHW action.

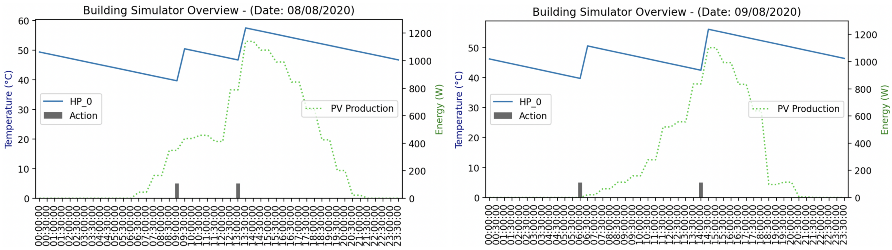

Figure 1.5 Building simulator dashboard.

If only operation mode 1 is considered, the heat pump would heat the tank to 50 °C and then stay on hold until the temperature drops to 40 °C, heating again to 50 °C and so on. This mode can be costly, as it does not verify the best time to perform the actions, which may be when there is no PV production or when the energy demand peak is high. The optimised profile provided by the building simulator checks the PV production schedule and anticipate actions to achieve economic savings, without adversely impacting users comfort.

The results from the building simulator were compared with the heat pump real consumption (baseline). Almost 30% of energy from the grid used for DHW could be saved if the actions had been performed as the optimal profile generated, mostly due to using the heat pump when the PV energy production was higher. This performance could drop to around 25% because the real world scenario can face more uncertainties related to users behaviour. For instance, the simulations consider an ideal profile of DHW usage, while the real user can suddenly decide to use all the water at once, thus making the control system activate more times over a day.

Another important metric is the rescheduled demand, which was calculated considering the total demand consumed inside the period of higher PV production incidence (10:00 to 18:00) and out of it, considering the average of the 10 analysed days. Looking at the collected data, only 37% of the heat pump consumption was inside the PV event range, while the optimised model increases this value to 57%, which effectively demonstrates an increase of 20%. Figure 1.6 presents the average consumption for both optimised (blue line) and real data (orange line), including also the average of PV production (dotted green line). Note that one of the benefits of the optimised model is the peak reduction and load shifting, as the peak load that was originally between 18:00 and 20:00 in the real data was moved a few hours before in the optimised version.

About the utilisation of PV production, the amount of real PV consumed and the optimal PV consumption was compared. For experimental purposes, that PV production was considered to be used exclusively for heat pump actions. For each

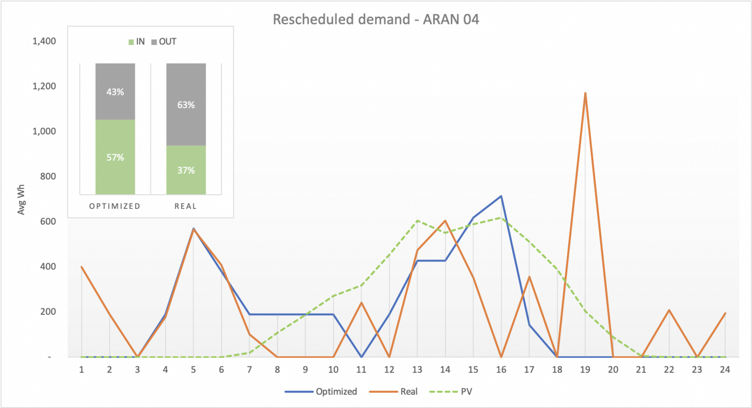

Figure 1.6 Rescheduled demand test case #3.

of the days, it was calculated the amount of renewable energy used and, on average, the optimal profile model improved renewable energy usage by 39.14%, by concentrating DHW action in periods with higher PV, as previously presented in Figure 1.5.

1.4.4 Test case #4

The aim of the fourth test case is to verify the changes in customer behaviour after receiving a message asking to turn off some appliances at a specific hour of the day with the main objective to decrease carbon emissions. In these events, the idea was to ask the participants to not use electric energy, without offering any direct financial incentives. The message was sent one hour prior the event to all the participants of the pilot through an app notification in both English and Gaelic languages, as follows:

"Electricity consumption of Ireland peaks within the next few hours, which means higher CO_2 emissions. Turn off some of your appliances between HH:MM and HH:MM - and help us with saving the climate."

"Buaicfidh tomhaltas leictreachais na hÉireann sna cúpla uair amach romhainn, rud a chiallaíonn astaíochtaí CO_2 níos airde. Múch cuid de do chuid fearais idir HH:MM agus HH:MM - agus cuidigh linn an aeráid a shábháil."

As opposed to the other test cases where that the focus was on the individual and the hour of the event was chosen based on their specific demand or production information, in this case the timing of the event was based on grid data provided by EirGrid [26]. The Aran Islands are connected to the electricity grid in Ireland, and during the test case period the highest demand for energy identified occurred between 17:00 and 18:00. This test case could also be applied using other inputs, such as periods of high consumption or lower production identified in the forecasting services. The baseline period for the test case was the first two weeks of July, 2020 and the DR events happened in August, 2020. Although the average peak in Ireland is between 17:00 and 18:00, before sending the message this range was confirmed on the demand system prediction provided by EirGrid [26] for each event day.

A reduction of 14.73% CO_2 was observed in aggregate during the days of the event compared to the baseline period. Considering the consumption of individual appliances, it was also verified that an additional 4.37 tCO_2e could have been avoided in the period if all the customers had done the action. However, the hourly analysis of the users' consumption behaviour showed no significant reduction in the peak load during the event period compared to the baseline. Considering the communication performance, houses 02, 04 and 08 were excluded and not considered in the validation process due to the missing data over the baseline period.

Conclusion

In summary, this chapter presents a platform for integrated management of residential energy systems by incorporating demand response events into a measure-forecast-optimise workflow for automated and semi-automated control of appliances and heat pumps. Through smart use of a set of deployed sensors and actuators, as well as synergistic relation between different services within the platform, the proposed system takes into account both generation and demand-side constraints in order to provide the most cost-effective and energy-efficient scenario for energy management. Different methodologies that were utilised for different components of the system are outlined, followed by a set of four thoroughly analysed use cases focusing on the adaptation of electric loads and utilisation of heat pumps in relation to specifically generated DR events.

Various scenarios are depicted in the discussed test case results with them portraying, in line with the applicable time periods and baseline data, effective savings of 2.5 t CO_2 emissions by providing information to users regarding their energy use and slightly below 80% of renewable energy self-consumption achieved once adequate messages are sent to denote periods with high expected production levels. Furthermore, an estimated 25% reduction of grid-imported energy for DHW temperature regulation through heat pump usage optimisation is demonstrated, followed by a case showing a reduction in CO_2 emissions of over 14% when responding to DR messages intended to shift appliance activations away from peak times for the grid.

Finally, it should be noted that this chapter provides a quantitative analysis based on the data that was collected and processed through the presented platform. The results presented are based on evaluations of the absolute available numerical data and are derived based on the selection of an appropriate baseline estimation methodology. Since energy use, especially for the residential case, is a complex multidisciplinary problem, there are also behaviour-related factors that influence the way in which energy is managed. Therefore, as a complementary addition to the presented results, related studies pertaining to the user experience domain [27] should also be considered.

Acknowledgement

The research presented in this chapter is partly financed by the European Union (H2020 RESPOND project, Grant Agreement No.: 768619 and SINERGY project, Grant Agreement No.: 952140) and the Ministry of Education, Science and Technological Development and the Science Fund of the Republic of Serbia (AI-ARTEMIS project, #6527051).

References

[1] Programme UNE, Construction GAfBa. 2020 Global Status Report for Buildings and Construction: Towards a Zero-emissions, Efficient and Resilient Buildings and Construction Sector - Executive Summary;. Accepted: 2020-12-15. Available from: https://wedocs.unep.org/xmlui/handle/20.500.11822/34572.

[2] Matt Golden, Adam Scheer, Carmen Best. Decarbonization of electricity requires market-based demand flexibility. The Electricity Journal;32(7):106621. Publisher: Elsevier. Available from: https://www.sciencedirect.com/science/article/pii/S1040619019302027.

[3] Barney A, Polatidis H, Jelić M, et al. Transition towards decarbonisation for islands: Development of an integrated energy planning platform and application. Sustainable Energy Technologies and Assessments;47:101501. Available from: https://www.sciencedirect.com/science/article/pii/S2213138821005129.

[4] Farhangi H. The path of the smart grid. IEEE Power and Energy Magazine;8(1):18–28. Conference Name: IEEE Power and Energy Magazine.

[5] Lund PD, Lindgren J, Mikkola J, et al. Review of energy system flexibility measures to enable high levels of variable renewable electricity. Renewable and Sustainable Energy Reviews;45:785–807. Available from: https://www.sciencedirect.com/science/article/pii/S1364032115000672.

[6] Faruqui A, Sergici S. Household response to dynamic pricing of electricity: a survey of 15 experiments. Journal of Regulatory Economics;38(2):193–225. Available from: https://doi.org/10.1007/s11149-010-9127-y.

[7] Solar Engineering of Thermal Processes. John Wiley & Sons, Ltd; 2013. Available from: https://onlinelibrary.wiley.com/doi/book/10.1002/9781118671603.

[8] Gómez-Omella M, Esnaola-Gonzalez I, Ferreiro S. Short-Term Electric Demand Forecasting for the Residential Sector: Lessons Learned from the RESPOND H2020 Project. Proceedings. 2020;65(1). Available from: https://www.mdpi.com/2504-3900/65/1/24.

[9] Gomez-Omella M, Esnaola-Gonzalez I, Ferreiro S. Short-term Forecasting Methodology for Energy Demand in Residential Buildings and the Impact of the COVID-19 Pandemic on Forecasts. In: Proceedings of 40th SGAI

International Conference on Artificial Intelligence. vol. 12498; 2020. p. 227–240.

[10] Favre-Perrod P. A vision of future energy networks. In: 2005 IEEE Power Engineering Society Inaugural Conference and Exposition in Africa;. p. 13–17.

[11] Mohammadi M, Noorollahi Y, Mohammadi-ivatloo B, et al. Energy hub: From a model to a concept – A review. Renewable and Sustainable Energy Reviews;80:1512–1527. Available from: https://www.sciencedirect.com/science/article/pii/S1364032117310985.

[12] Jelić M, Batić M, Tomašević N, et al. Towards Self-Sustainable Island Grids through Optimal Utilization of Renewable Energy Potential and Community Engagement. Energies;13(13):3386. Number: 13 Publisher: Multidisciplinary Digital Publishing Institute. Available from: https://www.mdpi.com/1996-1073/13/13/3386.

[13] Batić M, Tomašević N, Beccuti G, et al. Combined energy hub optimisation and demand side management for buildings. Energy and Buildings;127:229–241.

[14] Achieving energy efficiency through behaviour change: what does it take? — European Environment Agency [Publication];. Available from: https://www.eea.europa.eu/publications/achieving-energy-efficiency-through-behaviour/.

[15] Esnaola-Gonzalez I, Jelić M, Pujić D, et al. An AI-Powered System for Residential Demand Response. Electronics;10(6):693. Number: 6 Publisher: Multidisciplinary Digital Publishing Institute. Available from: https://www.mdpi.com/2079-9292/10/6/693.

[16] Alfageme A, Esnaola-Gonzalez I, Díez FJ, et al. Metaheuristics for Optimal Scheduling of Appliances in Energy Efficient Neighbourhoods. In: Marreiros G, Melo FS, Lau N, et al., editors. Progress in Artificial Intelligence. Lecture Notes in Computer Science. Springer International Publishing;. p. 151–162.

[17] Lissa P, Deane C, Schukat M, et al. Deep reinforcement learning for home energy management system control. Energy and AI. 2021;3:100043. Available from: https://www.sciencedirect.com/science/article/pii/S2666546820300434.

[18] Energy Master Plan 2018 Árainn and Inis Meáin [Publication];. Available from: https://156.234.107.34.bc.googleusercontent.com/community-energy/sustainable-energy-communities/tools-and-resources/energy-master-plan/Sample_Energy_Master_Plan.pdf.

[19] Clean Energy for EU Islands - Aran Islands (Ireland) [Publication];. Available from: https://www.euislands.eu/island-details/25.

[20] RESPOND D5.4 - Desktop dashboard and smart mobile client demonstrator [Report];. Available from: http://project-respond.eu/repository/.

[21] DEXMA Platform - AI powered energy savings tool [WebPage];. Available from: https://www.dexma.com/what-is-dexma-platform/.

[22] DEXMA - Automatic Baseline Calculator [WebPage];. Available from: https://support.dexma.com/hc/en-gb/articles/360013577059-Apps-Market-Automatic-Baseline-Calculator-ABC-.

[23] International Performance Measurement and Verification Protocol (IPMVP) [WebPage];. Available from: https://evo-world.org/en/products-services-mainmenu-en/protocols/ipmvp.

[24] Ceallaigh TJO, Dhonnabhain AN. Reawakening the irish language through the irish education system: Challenges and priorities [Journal Article]. International electronic journal of elementary education. 2015;8(2):179–198.

[25] Gallagher CV, Leahy K, O'Donovan P, et al. IntelliMaV: A cloud computing measurement and verification 2.0 application for automated, near real-time energy savings quantification and performance deviation detection [Journal Article]. Energy and buildings. 2019;185:26–38.

[26] Eir Group webpage [WebPage];. Available from: http://www.eirgridgroup.com.

[27] Toke Haunstrup Christensen, Henrik N Knudsen, Agustina Yara, et al.. RESPOND Deliverable 6.3 - User engagement assessment. The RESPOND Consortium;. Available from: http://project-respond.eu/repository/.