REPLICATE	Project no. 691735 REPLICATE PROJECT Renaissance of Places with Innovative Citizenship And Technology	This Project has received funding from the European Union's Horizon 2020 research and innovation programme under Grant Agreement N° 691735

REPLICATE PROJECT

REnaissance of PLaces with Innovative Citizenship And Technology

Project no. 691735

H2020-SCC-2015 Smart Cities and Communities

Innovation Action (IA)

D5.3 ENERGY DEMAND PLATFORM DEPLOYED TO MONITOR ENERGY GENERATION AND DEMAND

Due date of deliverable: 31/01/2019

Actual submission date: 30/01/2019

Start date of project: 01/02/2016 Duration: 60 months

Organisation name of lead contractor for this deliverable: Bristol City Council

Status (*Draft/Proposal/Accepted/Submitted*): Submitted

Project co-funded by the European Commission within the 7th Framework Programme		
Dissemination Level		
PU	Public	X
CO	Confidential, only for members of the consortium (including the Commission Services)	

Editor/Lead beneficiary :	Bristol City Council
Internally reviewed by :	SPES

Index of contents

1. EXECUTIVE SUMMARY

As part of the REPLICATE project we have successfully developed and rolled out an Energy Demand Management System (EDMS) which connects with Smart Home and Smart Mobility devices using a newly developed Smart City Platform.

The general functionality of the EDMS is monitoring and control of energy consumption and production. As a monitoring system, the EDMS accesses data from a wide range of equipment for generation, transport and consumption of energy, being able to flexibly create views of the data and making predictions and recommendations based on the observations. As a control system, the EDMS hosts energy management control logic to influence supply and demand patterns in order satisfy one or multiple technical or economic objectives, including relieving the transport and distribution networks, maximising local usage of locally produced energy, maximizing monetary benefits, and/or minimising the environmental footprint.

The EDMS also acts as a marketplace, where multiple stakeholders can create energy management programs, each having its own set of requirements for participation, its own visualisations, and its own business models including the incentives for participants. On the other side, consumers or producers can become participants of programs, benefitting from the incentives offered.

We have developed an innovative and highly sophisticated approach to modelling the behaviour of a large number of independent energy consumers using Reinforcement Learning to seek to demonstrate theoretically that it is possible to optimise energy use across a range of actors.

We are developing an advanced Demand Side Response trial which will design and test a range of consumer-acceptable energy interventions.

REPLICATE has developed a wide range of innovative approaches to bring together Smart Home and Smart Mobility solutions using a custom-built FIWARE Smart City Platform. The technologies we have developed are highly flexible, expected to provide significant environmental and social benefits, as well as providing what we expect to be a very attractive commercial offer, namely the ability to bring together a wide range of diverse small energy users such that they can manage their energy use collectively and therefore benefit from the commercial benefits that are only currently available to large players.

This Project has received funding from the European Union's Horizon 2020 research and innovation programme under Grant Agreement N° 691735

2. REPLICATE

The main objective of REPLICATE project is the development and validation in three lighthouse cities (**San Sebastián** - Spain, **Florence** - Italy and **Bristol** - UK) of a comprehensive and sustainable City Business Model to enhance the transition process to a smart city in the areas of the energy efficiency, sustainable mobility and ICT/Infrastructure. This will accelerate the deployment of innovative technologies, organisational and economic solutions to significantly increase resource and energy efficiency improve the sustainability of urban transport and drastically reduce greenhouse gas emissions in urban areas.

REPLICATE project aims to increase the quality of life for citizens across Europe by demonstrating the impact of innovative technologies used to co-create smart city services with citizens and prove the optimal process for replicating successes within cities and across cities.

The Business Models that are being tested through large scale demonstrators at the three cities are approached with an integrated planning through a co-productive vision, involving citizens and cities' stakeholders, providing integrated viable solutions to existing challenges in urban areas and to procure sustainable services. Sustainability of the solutions is fostered in three areas: economic and environmental and finally, fostering transparency in the public management.

In addition, the Model features the replicability of the solutions and their scale up in the entire city and in follower cities, particularly in three follower cities (**Essen** - Germany, **Lausanne** - Switzerland and **Nilüfer**-Turkey) that are involved in the project and therefore, have access to know-how and results achieved on the project so they can apply the developed model. At the moment, there are 2 observer cities, Guangzhou (China) and Bogota (Colombia).

This Project has received funding from the European Union's Horizon 2020 research and innovation programme under Grant Agreement N° 691735

3. INTRODUCTION

3.1 Relation to Other Project Documents

This Deliverable is one of the deliverables as set out in the Grant Agreement and sits alongside Deliverables 5.1 to 5.9 that relate to the Bristol Pilot. It also links to a number of the REPLICATE cross-cutting deliverables.

3.2 Reference documents

This document is based in the following projects level documents:

Ref.	Title	Description
REPLICATE Grant Agreement signed 240713.pdf	Grant Agreement	Grant Agreement no. 691735
DoA REPLICATE (691735)	REPLICATE Annex 1 – DoA to the GA	Description of the Action
REPLICATE Consortium agreement signed December 2015 (7th December version)	Consortium Agreement	REPLICATE project – Consortium Agreement
REPLICATE Project Management Plan	D1.1 Project Management Plan (v.1) (29/04/2016)	REPLICATE Project Management Plan
REPLICATE District Management Plans	D1.4 District Management Plan San Sebastian D1.5 District Management Plan Florence D1.6 District Management Plan Bristol	REPLICATE District Management Plans
REPLICATE	D11.1 Communication Plan	REPLICATE

This Project has received funding from the European Union's Horizon 2020 research and innovation programme under Grant Agreement N° 691735

Where there are contradictions, the documents listed above supersede this deliverable. The Grant Agreement is the contract with the European Commission so takes precedence over all other documents.

3.3 Abbreviations list

BIO	Bristol Is Open
CA	Consortium Agreement
CPMS	Charge Point Management System
DoA	Description of the Action
DSM	Demand Side Management
DSR	Demand Side Response
EC	European Commission
EDMS	Energy Demand Management System
IOT	Internet Of Things
GA	Grant Agreement
H2020	Horizon 2020
IOT	Internet of Things
OCPP	Open Charge Point Protocol
RL	Reinforcement Learning
SCP	Smart City Platform
PC	Project Coordinator
PL	Pilot Leader
PMP	Project Management Plan

SCP	Smart City Platform
TC	Technical Coordinator
VPN	Virtual Private Network
WP	Work Package
WPL	Work Package Leader

4. DELIVERABLE DESCRIPTION

This deliverable describes the approach we have taken to developing an Energy Demand Management System as part of the Bristol Pilot.

Section 5 sets out the requirements of the EDMS, the architecture we have adopted, as well as exploring the links with the Smart City Platform and user devices.

Section 6 covers EDMS monitoring functions and EDMS management functions, as well as outlining Demand Side Response proposals and EV Algorithm Optimisation work.

Section 7 covers innovations, impacts and scalability.

Conclusions are set out in **Section 8**.

This Project has received funding from the European Union's Horizon 2020 research and innovation programme under Grant Agreement N° 691735

5. ENERGY DEMAND MANAGEMENT SYSTEM OVERVIEW

5.1 EDMS Requirements

In this section, requirements on the Energy Demand Management System (EDMS) are described. These requirements have been derived from the project grant agreement and from discussions within the partners of WP5 over the course of the project.

The foreseen role of the Energy Demand Management System is to act as the community hub for managing supply and demand of energy. Energy in this context is not restricted to electricity but can include the gas, heating and cooling domains as well. The motivation for unifying the monitoring and control of these domains in a single system is that these sectors are becoming more interdependent. Modern power plants deliver both thermal and electrical energy, and in an increasing number of buildings several possibilities for space heating are combined, for example being connected to a heating network but still having some legacy electric heaters available that can utilise excess solar energy harvested on the building's rooftop.

The general functionality of the EDMS is **monitoring** and **control** of energy **consumption** and **production**. As a **monitoring system**, the EDMS accesses data from a wide range of equipment for generation, transport and consumption of energy, being able to flexibly create views of the data and making predictions and recommendations based on the observations. As a **control system**, the EDMS hosts energy management control logic to influence supply and demand patterns in order satisfy one or multiple technical or economic objectives, including to relieve the transport and distribution networks, to maximise local usage of locally produced energy, maximising monetary benefits, and/or minimising the environmental footprint.

The EDMS also acts as a **marketplace**, where multiple stakeholders can create energy management programs, each having its own set of requirements for participation, its own visualisations, and its own business models including the incentives for participants. On the other side, consumers or producers can become participants of programs, benefitting from the incentives offered. Programs can be as simple as regularly sending energy saving hints to the participants, or they can include complex control logic using data from smart meters and controlling devices of the participating users.

Requirement R1: Monitoring of Energy Production and Consumption. The EDMS shall offer the ability to monitor both the current status and history of all producers and consumers that are connected to it. This requirement has been derived from the description of Task T5.1.3.4 "Energy Management System – Holistic monitoring and analysis of demand and supply".

Requirement R2: DSM strategies for various business cases. The EDMS shall be able to actively influence the supply and demand patterns using Demand-side Management (DSM) algorithms to optimize the KPIs of the respective business cases. This requirement has been derived from the description of Task 5.1.3.5 "Energy Management system - DSM Strategies". More details on this approach are provided in Section 6.2.4.

Requirement R3: Ability for demand management with highly fluctuating loads. Private households are a much more challenging business case for DSM than large industries, due to the high variation and limited predictability of loads of the former. Furthermore, the responsiveness of private households to demand-management signals is highly stochastic. The DSM algorithms shall be designed to deal with such levels of fluctuation, coping with the limited reliability by making use of machine learning with probabilistic models. This requirement has been derived from discussions with domain experts within WP5.

Requirement R4: Compatibility with the FIWARE based Smart City Platform. The EDMS shall be able to communicate with producers and consumers via the FIWARE platform. As FIWARE is the technology used throughout the REPLICATE project's ICT infrastructure, the EDMS shall be able to read and understand data in FIWARE format, and write such data back into the ICT platform. This requirement follows from the REPLICATE approach to use FIWARE as the base ICT technology.

Requirement R5: Ability to communicate with proprietary cloud systems. The EDMS should be able to read the status of and control devices via the cloud system of device vendors. This requirement follows from the choice of Task 5.1.3.5 "Procurement and installation of home automation units" to procure smart appliances from Samsung, combined with the observation that the Samsung system is not yet FIWARE compatible, which is beyond what the project can influence. Thus the EDMS should be able to directly communicate with that Cloud system, using the FIWARE based platform here only to retrieve contextual information (e.g. credentials for access).

Requirement R6: Ability to manually control devices. The EDMS shall offer an interface for manual control of connected devices. While automated demand management is preferable, manual control is often required as a fallback or for debugging purposes. This requirement has been established from practical needs of the smart home use case, and from the requirements of demonstrating the control capabilities of the system.

Requirement R7: Nonaggressive management. When the device owner manually overwrites management decisions of the EDMS, the management logic of the EDMS should not attempt again to manage the device for an appropriate time range. What is appropriate depends on the device type. For example, for a washing machine the appropriate time range is until the current washing cycle has been finished. This requirement has been derived from the setup of the planned DSR-trial.

5.2 EDMS Architecture

As shown in Figure 1, the EDMS conceptually distinguishes between two **user roles** (excluding technical administrator roles, which are not shown here for the sake of simplicity). **EDM program managers** are users which create and administrate **EDM programs.** In contrast, **Consumers** or **prosumers** are the users which potentially participate in such programs. They can own **devices**, which they register at the EDMS and manage the connectivity and other properties. A device can be a monitoring system like a smart meter, a control system like a thermostat, or a device that performs both monitoring and control like a smart pug. When consumers or prosumers participate in an EDM program, they can authorise some of their devices to be used by the program by means of monitoring and control. Devices can themselves be energy management subsystems, e.g. managing supply and demand on the level of an individual household, building, or neighbourhood.

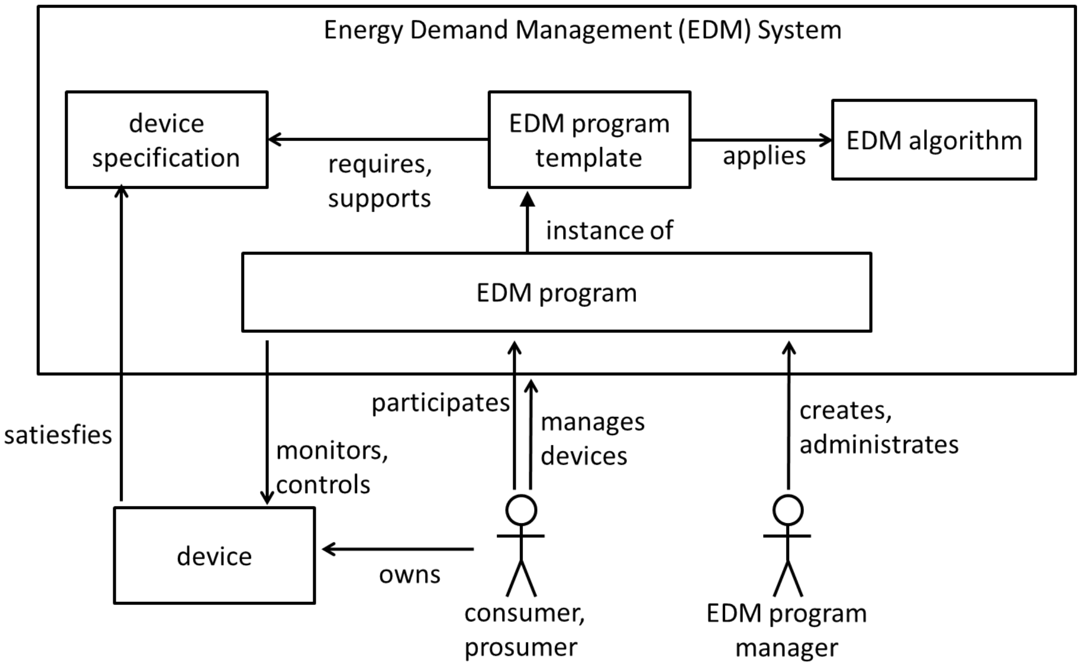

Figure 1: Conceptual model used by the EDMS.

This Project has received funding from the European Union's Horizon 2020 research and innovation programme under Grant Agreement N° 691735

The management logic of EDM programs is determined by an **EDM program template**. The template determines which **EDM algorithm** it runs, which **device specifications** it supports, and which device specifications are required for participation in the program. A consumer or prosumer can only participate in a program if for each required device specification she has registered such a device and authorises it to be used by the program. Furthermore, an EDM program typically comes with an incentive model for the consumers/prosumers and a business model for the program manager.

In EDM programs, the participating users can have different program-internal roles. For example, the program might distinguish between small and large consumers, or between producers and consumers. Each program-internal participant type might have different device requirements to fulfil.

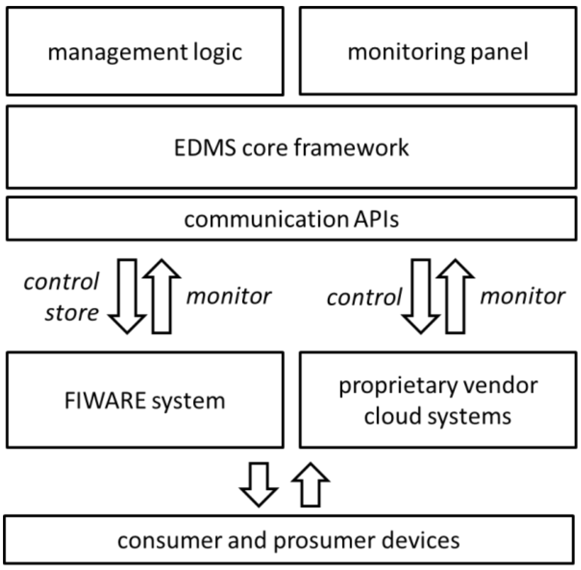

Figure 2: EDMS architecture

Figure 2 depicts the architecture of the EDMS and its interplay with other ICT components. The components on top represent the four building blocks of the EDMS. The management logic is a component that is separated from the main framework in order to address Requirement R2. Depending on the requirements of the energy demand management business model, the management logic can be instantiated with the required control logic. The monitoring panel addresses Requirement R1; providing data visualisation of not only for the connected consumers and prosumers, but also aggregated views for the energy management use case under consideration. Data exchange with external components is realised by the communication APIs. The communication with the FIWARE platform takes place via the NGSI interface of FIWARE (Requirement R4). The FIWARE instance is used for retrieving data about consumers and prosumers connected to FIWARE, and control commands can be issued by a publish/subscribe mechanism. Furthermore, the EDMS also exports data to FIWARE for publishing as open data or other usages, using the same interface. Among the communication APIs are also interfaces to communicate with the non-FIWARE cloud systems of device vendors (Requirement R5), which is necessary when those vendors are not FIWARE compatible and no adapter is available to enable FIWARE usage.

5.3 Links to Smart City Platform

The City of Bristol provides a smart city infrastructure, which is a ring of optical fibres connecting multiple locations together. In these locations there is server infrastructure, network switching and controllers, distributed computing and data storage. The data reaches our network from the cloud via VPNs connecting to our main firewall at the University of Bristol. The FIWARE instance is located on virtual machines within the network nodes.

Bristol Is Open has provided the infrastructure platform onto which data is backhauled and stored.

5.3.1 Smart Homes

The smart homes send data back to our FIWARE instance via Loxone MiniServer connectivity and over a VPN linking individual users with the Smart City Platform. Individual machines record and report their energy use, and data is collated in the FIWARE backend.

A demonstration of the data backend was given at the EU General Assembly in Bristol on 24 October. This showed energy consumption data from the Smart Home being fed into and stored on our instance of FIWARE. The data is compiled by the Loxone MiniServer and sent as a string of code from each home. It can be configured to poll for data as often as is required. For the demo the data updated every 5 minutes.

This Project has received funding from the European Union's Horizon 2020 research and innovation programme under Grant Agreement N° 691735

5.3.2 Smart Mobility

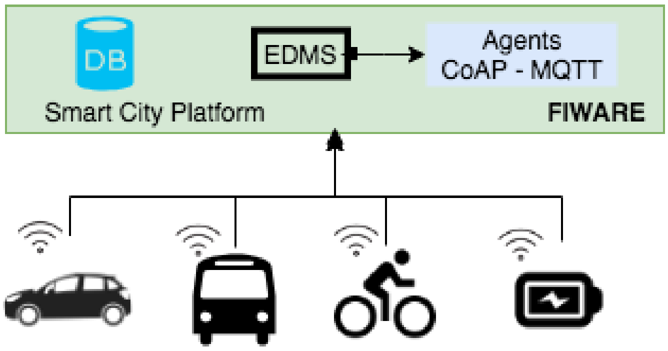

Figure 3: Mobility Links to SCP

The connection of the mobility pilot with the SCP is depicted in Figure 3. To accommodate different IOT data protocols we deployed FIWARE IOT agents component in order to allow diverse data coming to the platform. The FIWARE platform is a cloud deployment where the data are being collected by the appropriate agent and then data is available through the context broker or the historical database. To guarantee the security of the data, the platform maps every user/device to a FIWARE service and then isolates the data access to only authorised users.

5.4 Links to Consumer and Prosumer Devices

5.4.1 Smart Connected Homes

Smart Connected Homes involved 150 homes in the project area receiving smart home equipment including smart appliances and related home hub and monitoring devices.

See Appendix A for a summary of devices to be installed in each household and Appendix B for the Smart Homes Information Pack sent to participants. We worked with a number of suppliers to provide specialist equipment. Loxone supplied the smart home hub and smart plugs. We opted for smart appliances provided by Samsung as these had high environmental credentials as well as having advanced smart functionality. A Raspberry Pi was supplied by Narec to provide a secure data connection to the Bristol Is Open data store.

Recruitment

Households were recruited with the support of a Community Engagement Group led by Bristol Energy Network and supported by Knowle West Media Centre. The aim of this group was to recruit and support households to participate from a wide demographic to reflect the project area. There was a particular focus on involving groups who might not traditionally engage with technology projects of these types. An Asset Based Community Development approach has been taken utilising the 'Bristol Approach'.

A criterion for involvement was set to ensure equity in the offer to the community. For example only older 'A' rating or below machines would be replaced due to the environmental impact of replacing new machines. Households were asked to sign terms and conditions and were supported during the installations on how to operate the equipment.

Installation

Installations were managed by Narec Distributed Energy and involved older inefficient machines being donated to The Sofa Project, a local social enterprise, for either recycling or reuse amongst lower income households.

Data flow

Energy data in the households was captured via two primary devices. Firstly a whole home meter reader attached to home electricity meters. Secondly a smart plug behind the appliances to record the usage of the device themselves. This data is both stored locally on the Loxone home hub whilst being transferred to Bristol Is Open via the VPN created with the Raspberry Pi. Once with BIO this is then supplied to UWE for evaluation purposes.

A flagging system is in place to monitor the status of the VPN connection to households that contains an alert system should any errors be detected. Equally a flagging system to detect anomalies in data quality is being developed with UWE to ensure quality of data during the monitoring period.

The monitoring panel is the interactive component of the EDMS. It is designed as a versatile tool for freely configurable visualisations of the available data. Additionally, manual control of devices and managing EDM programs is supported.

This Project has received funding from the European Union's Horizon 2020 research and innovation programme under Grant Agreement N° 691735

5.4.2 Smart Mobility

"Deliverable D5.7 - Transport Infrastructure Adaptation Including EV Charge Point Installation" set out in detail how Electric Vehicles will be charged using Smart charge points. It detailed how modern charge points are being installed around the REPLICATE area for use with new Co-Wheels vehicles. These charge points are typically 22kW and at least Open Charge Point Protocol (OCPP) version 1.5 compliant - this means they are able to facilitate a variety of smart grid functionality.

The way we have set up the system in Bristol, as demonstrated at the 2018 General Assembly, is that the charge points communicate with the existing "Charge Your Car" Charge Point Management System. This system then allows real time and historic data to be produced on charge point usage for a number of functions. The real time data is automated via an API call to send data to the Smart City Platform using its FIWARE functionality. This data can then be used by the NEC Energy Demand Management System and the Route Monkey Optimisation Algorithm, as detailed in later sections.

Bristol City Council is in the process of upgrading its Charge Point Management System (CPMS) so that it provides a wider range of smart functionality as one of our key findings in this area is that it is the CPMS that dictates which charge point functionality can be implemented. The lessons we have learnt through REPLICATE have greatly helped Bristol City Council to ensure it can require the most modern CPMS standards.

This Project has received funding from the European Union's Horizon 2020 research and innovation programme under Grant Agreement N° 691735

6. EDMS PILOT DEPLOYMENT

6.1 EDMS Energy Monitoring

The EDMS architecture as described in the previous section has been successfully implemented such that data now flows from the consumer/ prosumer devices via the Smart City Platform to the Energy Demand Management System. This infrastructure then enables the monitoring and management of smart devices.

A simplified view of the current deployment of the EDMS in the city of Bristol is depicted in Figure 4. What has been mainly demonstrated in 2018 is the interaction of the Energy Demand Management System (EDMS) with devices in the Smart Homes. The endpoints on the side of the Smart Homes are smart plugs and meter readers, as well smart appliances. The default way to communicate with the EDMS is via the REPLICATE ICT platform, which is based on FIWARE. However, in the case of the smart appliances, technical restrictions required us to interact with them via the vendor cloud (the Samsung cloud system in this case).

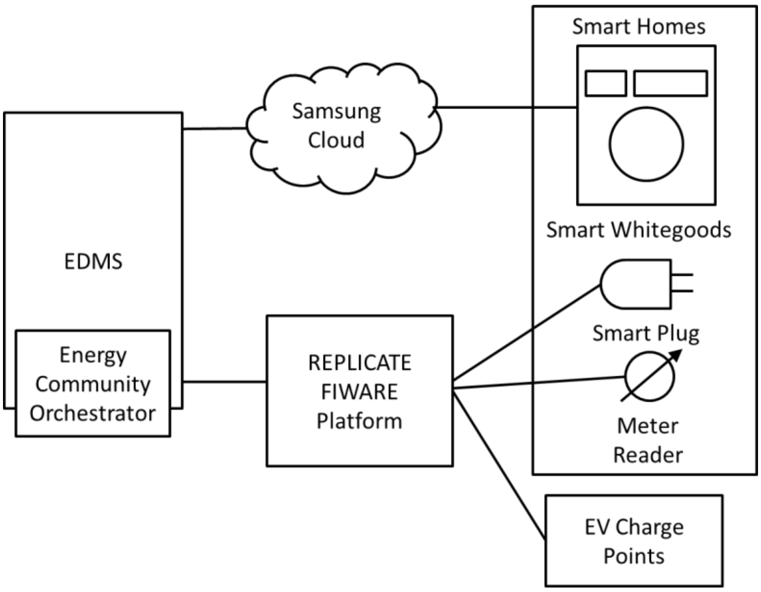

Figure 4: Current EDMS deployment in Bristol

Furthermore, the EDMS is accessing data from EV charge points. This data is accessed via the FIWARE platform that is hosted in the datacentre of Bristol Is Open.

On the EDMS side, the point of user interaction is the Energy Community Orchestrator, which is the name of the graphical user interface for monitoring and control. More information about this tool is provided in the subsequent section.

6.1.1 Map view

Figure 5 shows an example of a map view. Physical energy devices (e.g. smart homes, charge points) are shown on the map at a specific longitude/latitude. We remark that the exact longitude/latitude shown may not always reflect the actual physical location. For example, if the logged-in user is the home owner, she will see her home at the right location, but from another user's point of view (e.g. a program manager), the individual homes are shown at random locations for privacy reasons. The charge points, however, are not assumed to be privacy-sensitive and are always shown at the correct location.

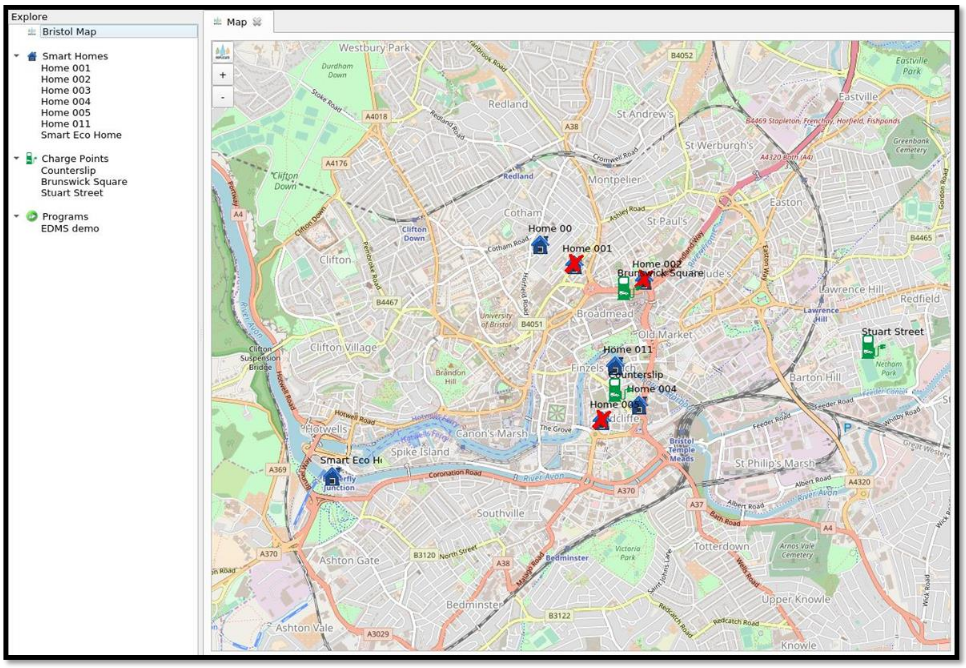

Figure 5: EDMS monitoring – map view

The users have the possibility to zoom and pan the map using the familiar mouse wheel and drag gestures, or they can use the top left icons or the keyboard shortcuts (ctrl+plus, ctrl+minus).

The icons indicate additional information about the particular device they represent. Whenever the information backend is not reachable (e.g. FIWARE, or the Samsung cloud), a red cross will be displayed on top of the icon, as shown in the screenshot above (Figure 5). Additionally, charge points are normally green but will turn red when a car is currently charging. Similarly, homes will have a "play" icon when the corresponding Samsung device (e.g. washing machine) is currently running.

Another convenient way to obtain more information about particular homes or charge points is by hovering the mouse cursor over them, which will bring up a popup with exact values describing the associated devices.

If the user requires more detailed information, she can double-click on a specific icon to bring up one of the dedicated views as detailed in the next sections. Finding an exact icon for a particular device might be difficult on the map though. For this reason, there exists a list on the left hand side of the monitoring panel, showing all devices in an ordered and searchable fashion. Double-clicking an item from this list has the same effect as on the map, namely, opening one of the views described below.

6.1.2 Consumer View

Figure 6 shows the screen obtained after having clicked a home icon.

Static information about this particular home is available on the top and as above, privacy-sensitive information (name/address) is only displayed if the home owner is the logged-in user.

The next section lists all available devices in this particular home. Most households will have a Loxone Meter Reader, measuring the overall electricity consumption of the household. The EDMS communicates with the FIWARE platform to obtain the real-time status and measure of this reader. Furthermore, an historical graph of this consumption is shown which the user can interact with by zooming and panning.

Most households will additionally be equipped with at least one Samsung appliance and some will have more than one as it is the case for this particular household. The EDMS communicates with the Samsung cloud in real-time to obtain information on the appliances, also shown as an interactive graph.

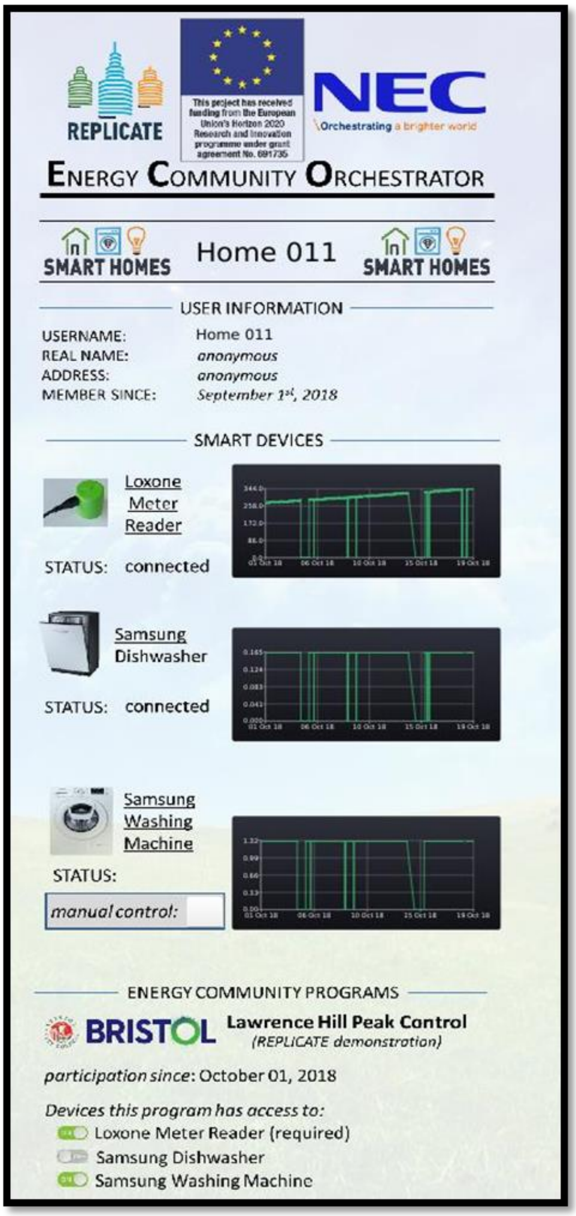

Figure 6: EDMS monitoring – consumer view

Samsung appliances are controllable through the Samsung cloud and the EDMS exposes this functionality to home owners. Users can use the "manual control" section to control (start/pause/resume/cancel) operation of e.g. the washing machine at any time. Care has been taken to not interfere with the intentions of the home owner, e.g. the EDMS will actively pause appliances to save electricity but will not touch specific households on days where it detects a manual override (i.e. force run the appliance) by their respective owner.

The last section of this view presents a list of all programs that this particular household is enrolled in. For each program, the home owner (or program owner) may individually specify which appliance may be controlled.

6.1.3 Program View

The figure on the right-hand side shows an example of the screen obtained by double-clicking one of the programs in the list shown on the right-hand side of the monitoring panel (see Figure 7).

The overall program description is shown on the top half, followed by an administration section offering the opportunity to change some of the relevant parameters. This particular program is a simple demo that will throttle households' appliances if their combined electricity consumption gets above a specific peak threshold that the program owner may modify at any time.

The next section of this screen presents some statistics pertaining to the program, namely the cumulated consumption. The last section presents the individual participants that may be added or removed individually.

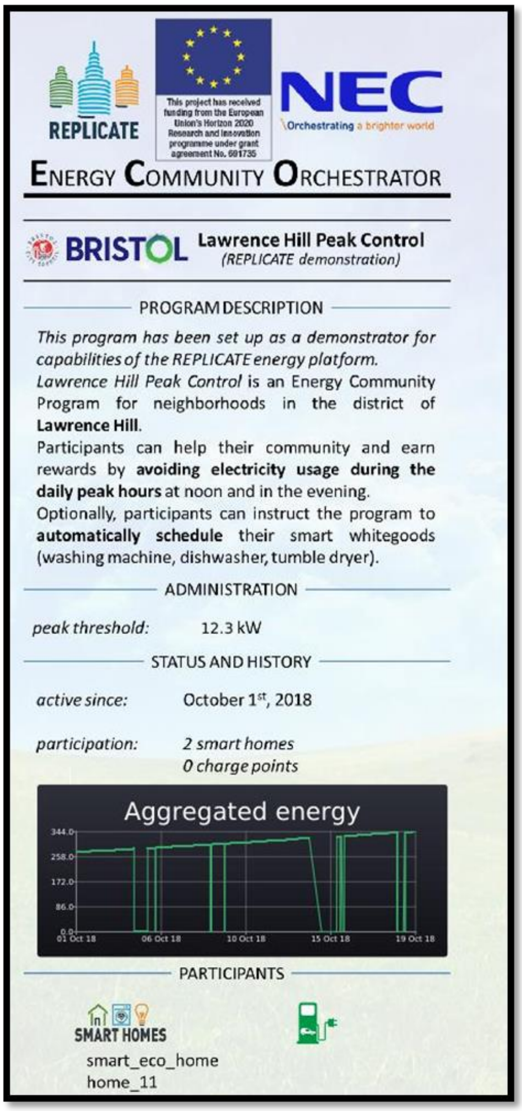

Figure 7: EDMS monitoring – program view

Note that this screen represents an example program and such programs can be created by writing up a (python) script for the device control logic, following a simple API. The scripts are currently developed manually by the EDMS team, but a future version of the EDMS may offer a graphical development environment for new programs for power users. Such programs also generally require a custom user interface for administrative/reporting purposes and this possibility exists accordingly, as described in the next section.

6.1.4 Custom Views

One of the most essential features of the EDMS monitoring panel is the ability to create custom data views. Views that have been created can be saved for later use and restricted to specific users.

From the complete list of available data sources on the left-hand side of the monitoring panel (Figure 5), the user can drag-and-drop data into the visualisation area. The resulting data curves can be interactively zoomed in and out both horizontally and vertically; they can be annotated with titles, free text, as well as static images and logos. The graphs are updated in real-time as soon as new data becomes available for the corresponding FIWARE entity.

Data sources may also be csv-files created by some external process (e.g. the python program scripts) to make it possible to present post-processed data, potentially aggregated from multiple individual sources. An example is the cumulated electricity consumption of the previously shown demo program that is calculated by the python script, used internally, and saved on disk for visualisation purposes.

Future versions of the EDMS may expose integrated development facilities to create such aggregated data sources, either visually through linking blocks of data together, or by enabling users to write simplified python code.

This Project has received funding from the European Union's Horizon 2020 research and innovation programme under Grant Agreement N° 691735

6.2 EDMS Energy Management

Automated management of demand and supply is a core functionality for the EDMS according to Requirement R2. The ambition of energy management at the Bristol pilot is to leverage the demand-response potential of private households, combine them with other city assets like PV, district heating, and electric vehicles, to build up a critical mass of demand flexibility that can be monetised on the energy market, helping to offer better energy tariffs to citizens.

The challenge of energy demand management for small consumers is two-fold. On the one hand, load curves of private households do not follow an easily predictable daily pattern. In fact, a large proportion of the load depends on factors that are unknown to the management system and thus have to be treated as randomness. Even when taking into account contextual information like the weather, day of the week, and demographic factors such as household income or family status, a large part of the observed load remains nondeterministic, and it is difficult to even tell whether an observed load reduction is due to a management signal that has been sent or due to coincidence.

On the other hand, it is also not practical to require a guaranteed level of demand elasticity from private households. Whether or not demand-response requests are satisfied, and the amount of load-shifting or load shedding, remains a variable that can only be estimated. Also the effects of monetary incentives or time-of-use tariffs have such characteristics.

For this reason, the algorithmic framework that has been developed for the purpose of energy management takes into account the following design goals:

a) Minimal assumptions and minimal prior knowledge on the consumer behaviour; in particular no prior assumptions on responsiveness of consumers to demand-response signals
b) Self-adaptation of system to responsiveness of consumers
c) Ability to optimize a global objective across all consumers

6.2.1 Overview of the approach for energy management

The algorithmic framework introduced in the subsequent paragraphs has a modular structure in order to be adaptable to a wide range of energy management use cases. There are two main modules. The first module tries to predict the behaviour of consumers, in particular their reaction to energy management signals (e.g. control signals for devices, or load shifting requests). The second module takes these predictions as input and computes optimal assignments of such signals to the different consumers, taking into account use case specific constraints.

Both modules can be instantiated with different implementations that address the need of the respective use case. For example, for the DSR trial with smart homes, the signal that is applied to each home represents a time interval for smart appliance pausing. The prediction module must estimate the probability for each household that (a) the smart appliance will be in operation during the time interval (so that the pausing action has an effect), and (b) the probability that the user does not manually overwrite the pausing action. Both (a) and (b) are dependent on how the pausing interval is chosen, and this dependency is represented in the prediction model. In this particular use case, the optimisation model selects a pausing interval for each home based on these predictions, in order to maximise the overall positive influence on the load curve, under additional constraints such as not resuming the operation of too many devices at exactly the same time and overlapping the pausing intervals with peak load times.

The instantiations of the modules that have been implemented and tested for the purpose of this document follow a more general purpose – they are designed as a generic energy management algorithm that can work with arbitrary signals (including price signals, direct device control, time-based reduction requests). The algorithms described below do not even need to know the meaning of the energy management signals – they apply those signals in a systematic way to find out how they influence the load curves of the consumers. We remark that this generality potentially comes at the price of efficiency – the generic algorithm might take longer to adapt to a specific use case than a use-case specific solution. However, in the experimental study we were able to demonstrate that, despite a rather large range of 24 unknown signals per consumer, the algorithm was able to converge within a few tests of every signal towards a good assignment of signals.

In the subsequent paragraphs, the theoretical basis of the framework, the algorithm implementation, and a simulation-based experimental evaluation are presented.

6.2.2 Theoretical basis for EDMS Framework (the Combinatorial Multi-Bandit Problem)

Energy management across multiple consumers/prosumers with initially unknown responsiveness can be modelled as a Reinforcement Learning (RL) problem. Instances of Reinforcement Learning problems are characterised by a state space and an action set. At each point in time, the RL agent can take one action from the action set. The effect of the action is that the environment goes into a new state, and a positive or negative reward is retrieved by the agent. Both the next state and the reward are potentially influenced by the action taken by the agent, but in a way that is initially unknown and has to be learned.

This Project has received funding from the European Union's Horizon 2020 research and innovation programme under Grant Agreement N° 691735

This abstract model for Reinforcement Learning is mapped to the energy management problem with n consumers as follows: There are n actors, one corresponding to each consumer. The action space consists of m different energy demand management signals that can be sent to each consumer. We make no specific assumptions on what these signals mean. In one context they could be price signals or other financial incentives, in other contexts they might be requests to reduce energy usage during some time interval, or even direct control signals for devices. The state of the environment is observed as the electric load curve of each consumer in response to the signal (for simplicity, the absence of any signal is considered as a special signal as well). Imagine for example that before the daily evening peak each consumer receives some signal; then the state of each consumer is the load observed during the evening peak time interval. Note that instead of load curves we directly consider *load reduction curves*, which are defined as the difference between the observed load curve and a baseline curve for every consumer.

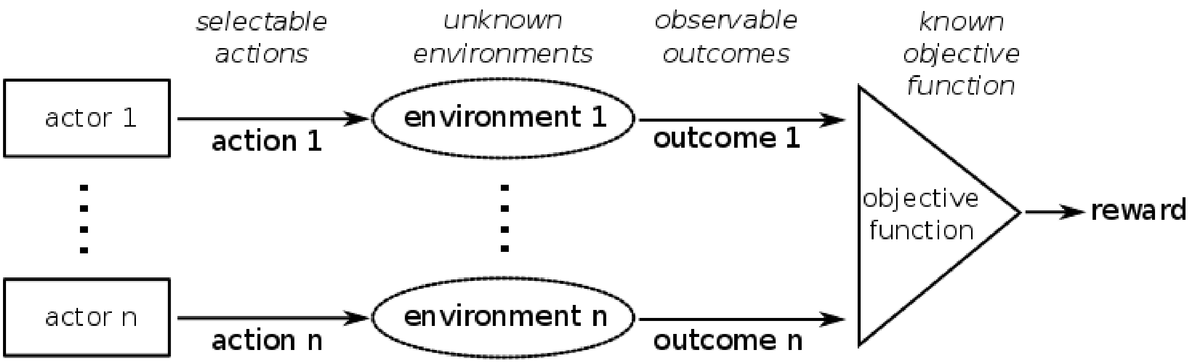

Figure 8: Combinatorial Multi-Bandit Problem

Figure 8 visualises the Reinforcement Learning problem under consideration, which we call Combinatorial Multi-Bandit Problem due to it resemblance to the well-known Multi-Arm Bandit Problem, but here having multiple bandits (i.e. actors) each having an own action set. When an actor takes an action, the result can be observed as the state of its environment, but it is initially unknown how each action influences the environment. There is a common reward function for all actors. This reward is computed from the combination of outcomes of the individual actors by a known reward function.

More formally, in the Combinatorial Multi-Bandit Problem we are given a multitude of n actors, indexed by i= 1 …n. Each actor i can perform a number k^i of different actions, indexed by $a_i = 1 \dots k^i$. Selecting an action $a^i$for actor i results in a probabilistic output $o^i \in O^i$. Unlike in most of the work on Multi-Arm Bandit Problems, the outputs o^i do not directly represent scalar rewards but can be elements of any space O^i over which probability measures can be defined. For each actor i the distribution $P(o^i)$ over the possible outputs is unknown and only depends on the most recent action a^i that has been chosen for the actor. What is known is the *objective function* $r: O^1 \times \dots \times O^n \to R$ which combines the individual actor outputs into a scalar *reward*.

Time proceeds in discrete episodes $t = 1 \dots T$. In each episode an *action assignment* $a_t = (a_t^1, \dots, a_t^n)$ has to be made, assigning to each actor i an action a_t^i. We also use vector notation $o = (o^1, \dots, o^n)$ to denote the outputs of the n actors observed at the end of each episode.

The objective is to maximize the *total expected reward* $\sum_{t=1}^{T} E[r(o) \mid a_t]$.

As the outcomes only depend on the actions chosen in the same episode, there exists an optimal action assignment $\tilde{a} = (\tilde{a}_1 \dots, \tilde{a}_n)$ that maximizes the expected reward independently of the current episode t. If $\tilde{a}$ was known, then a trivial optimal policy would be to select $\tilde{a}$ in each episode.

As $\tilde{a}$ and the outcome probabilities are unknown, an optimal or near-optimal assignment of actions has to be learned from experience using a strategy that balances *exploration* (choosing actions for the purpose of learning their outcome distribution) and *exploitation* (choosing action assignments known to perform well).

6.2.3 Modelling the Energy Demand Management Problem (the algorithm implementation)

When modelling our application as a multi-bandit problem, each consumer corresponds to an actor, and the action space consists of signals that can be sent to energy consumers e.g. on a daily basis. The signals incentivize consumers to re-schedule or curtail their energy consumption in various ways. At the end of each episode the load curves of the consumers are observed. A metric, defined on the aggregated load curve of all consumers, measures the extent to which the energy management goal has been fulfilled. In an experimental setting we consider the target to obtain the maximal load reduction that can be sustained during a given *target time interval*, measured against a given baseline curve. This specification models a common use case in energy demand management, where the owner of a management unit agrees with the network operator to provide load reduction as a service.

Formally, there are n actors, here each having the same number k of actions. Having applied an action, the outcome is observed as an H-dimensional discrete *load reduction curve* $l^i = (l_1^i, \ldots, l_H^i)$. The objective function r is defined as

$$r(l^1, \ldots, l^n) = min_{h=1\ldots H} \sum_{i=1}^{n} l_h^i$$

Algorithmic framework

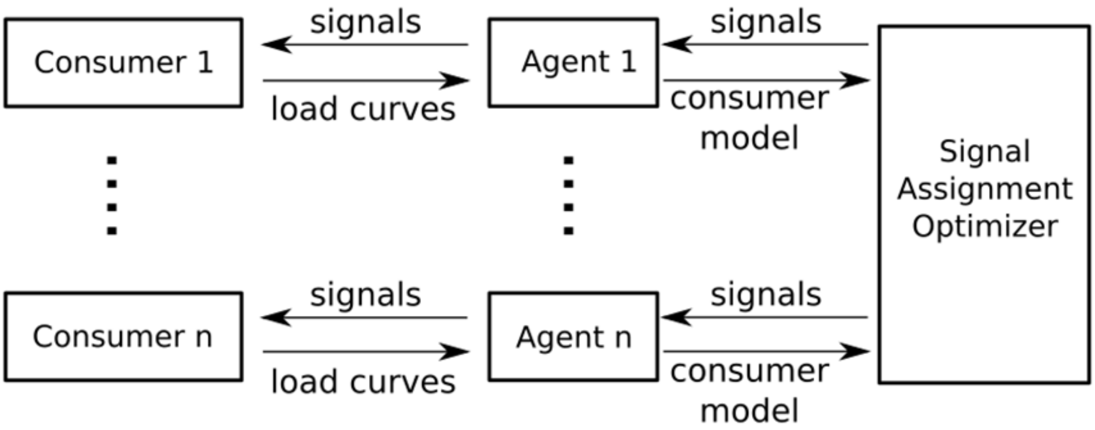

Figure 9: Visualisation of Algorithmic Framework

For the energy management application of our learning problem we have designed a range of algorithms that combine combinatorial optimization with exploration strategies for bandit problems. In the next paragraphs we describe the optimization module used by all algorithms, and subsequently we explain how the different algorithms apply the module.

Optimization module

Assume that complete information about the environments for $N \geq 1$ *sample episodes* is available, that is, for each t = 1...N a known function $\ell_t(i, j, h)$ specifies the load reduction that will be achieved at any daytime h when applying any action j to any actor i. Given this information, the optimization module computes an action assignment maximizing the *average* reward over the N sample episodes.

This computation is realized by formulating and solving the problem to maximize the average objective value as an integer linear problem. The program uses $n \cdot k$ binary variables b_{ij} to specify whether actor i will apply action j.

maximize $\sum_{t=1}^{N} M_t$

subject to $\sum_{i=1}^{n}\sum_{j=1}^{k} \ell_{t(i,j,h)} \cdot b_{ij} \geq M_t$ for $h = 1 \ldots H, t = 1 \ldots N$

$\sum_{j}^{k} b_{ij} = 1$ for $i = 1 \ldots n$

$b_{ij} \in \{0,1\}$ for $i = 1 \ldots n,\ j = 1 \ldots k$

Each summand M_t of the objective function represents the reward obtained in one episode t. For any episode t and any daytime h within that episode, an upper bound on M_t is given by the total load reduction of all actors at that daytime. Thus, for each episode t there are H such upper bounds on M_t. These $N \cdot H$ constraints are represented by the second line of the program. The third line represents the n constraints that only one action can be chosen for each actor $i = 1 \ldots n$.

We remark that the mathematical program cannot be solved in polynomial worst case runtime unless P=NP due to equivalence to the NP-complete SAT problem. Nevertheless, state-of-the-art solvers for integer programs can solve fairly large instances within reasonable runtime. When computation time runs out, solvers can also return the best suboptimal solution found so far, including an upper bound on the remaining difference to the optimal solution.

Single Episode algorithm

We present a *Single Episode (SE)* algorithm using the optimization module with N=1 to compute an optimal action assignment for a single episode. This episode, specified by function $\ell := \ell_1$, has to represent the population of all possible episodes.

A simple way to realize such a function is to use

$\ell(i,j,h) := \bar{l}(i,j,h)$,

where $\bar{l}(i,j,h)$ is defined as the average load reduction that has been previously obtained at daytime h by actor i after having applied action j. In case no experience is available for action j of actor I, an initial value is used, so

$\ell(i,j,h) := I(i,j,h)$,

where $I(i,j,h)$ is a user-defined optimistic initial value. The reason why the initial value is chosen optimistically as a high load reduction is to incentivize the optimizer to initially explore all actions until some real experience about them is available. To control the amount of exploration, we introduce a parameter called *initial value weight* $\beta \geq 0$, which determines how much the optimistic initial value is taken into account in later episodes. In the computation of the average, each sample from real experience is weighted with 1 while the initial value is weighted with β.

As an alternative means of exploration we also use the well-known ϵ-greedy scheme and apply it to each actor independently, that is, in each episode each actor ignores with probability ϵ the assignment computed by the optimization module and applies instead an action selected uniformly at random.

A final parameter driving exploration for SE is the number $\tau \geq 0$ of *initial exploration episodes*, during which each actor applies random actions to obtain some initial experience.

Multi Episode algorithm

A limitation of single episode algorithms is that in general it is impossible to adequately represent the population of all episodes by a single one. Solving the linear program using the expected values of ℓ as parameters does in general not result in the action assignment maximizing the expected reward due to the nonlinear objective function r.

One way to deal with this issue is to compute an assignment that is optimal for an ensemble of episodes. Given bias-free sample episodes $1 \dots N$, the average reward of assignment a is a bias-free and consistent estimator of the expected reward of a. Thus, the assignment with maximal average reward among the N samples almost surely converges against the optimal assignment as N grows. Motivated by this observation, our *Multi Episode (ME)* algorithm constructs a set of N sample episodes and applies the optimization module accordingly.

Let $S(i,j)$ be the set of samples collected for actor i and action j, where each sample represents an H-dimensional vector of load reduction values that have been observed from actor i having applied action j. If no such sample has been collected yet, the set is defined using the (optimistic) initial values $I(i,j,1) \dots I(i,j,H)$.

A sample episode can be constructed by setting $\ell_t(i,j,\cdot)$ to some value chosen from the sample set $S(i,j)$. Note that for any i,j,i',j' the values $\ell_t(i,j,\cdot)$ and $\ell_t(i',j',\cdot)$ possibly have been collected at different episodes. Thus, information on the correlation of the two curves is discarded by this construction method, and in general the sample episode is only bias-free under assumption of stochastic independence of the outcomes of all actor/action combinations.

We use the number of sample episodes N as a hyper-parameter of the Multi Episode algorithm, and we construct the N episodes by randomly picking from each $S(i,j)$. After an element of $S(i,j)$ has been selected, it is temporarily removed from the set until all other elements of $S(i,j)$ have been selected (and removed) as well. This way we ensure that the available sample curves are distributed among the N sample episodes as evenly as possible, which reduces the correlation between the episodes and thus decreases the variance of the expected reward estimators.

Offline computation of optimal assignments

The regret of a reinforcement learning algorithm is defined as the difference between the algorithm's outcome and that of the optimal algorithm which knows which action combination maximizes the reward. The regret is a standard metric to evaluate reinforcement learning algorithms.

To estimate the regret of the energy management algorithms described above we implement a method to compute offline an assignment very close to the optimum. A true sample episode can be obtained from simulated actors by querying their hypothetical behaviour for each action. The assignment is then computed by the optimization module using a sufficient number of such samples.

6.2.4 Simulation-based experimental evaluation and consumer modelling

To evaluate our experimental framework, we simulate a scenario where energy consumers can be requested to reduce their energy usage during a time interval that is specified in the request. This model corresponds to a hierarchical demand-response system, where the global management unit communicates with management subsystems (e.g. on the level of individual households), and these subsystems in turn are responsible to do the low-level control of devices and potentially communicate with the energy user. In the simulations it is assumed that each of the consumers has some load that can be shifted to a different time (e.g. operation of a washing machine), as well as some smaller amount of load that can be shedded (e.g. turning off some unnecessary lights). Furthermore, the base load of every consumer is fluctuating in a stochastic way. The target of the management system is to coordinate the load reduction requests such that throughout a pre-specified target time interval (which could for example represent the evening peak hours) as much load as possible is avoided. Here it is essential that the aggregated load reduction is sustained over the whole target time window. This means that if some of the consumers finish their load reduction too early, this should ideally be compensated by other consumers starting their load reduction at that time. Technical details are described in the remainder of this section.

The n actors in our setup are modelled as n energy consumers. The overall objective for the management system is to sustain the largest possible load reduction over a given daily *target time interval*, which is discretized into daytime slots $1, \dots, H$. The actors interpret each action as the request to reduce energy usage during some *request interval* which is a sub-interval of $[1, H]$.

Out of the $\Omega(H^2)$ possible sub-intervals we define an action set A of size $O(H)$ as the smallest set with $[1, H] \in A$ and $[h_1, h_2] \in A \Rightarrow [h_1, [0.5(h_1 + h_2)], [0.5(h_1 + h_2)], h_2] \in A$. This limitation is to restrict the amount of exploration needed for each individual actor. In our experiments we work with H=12, and the above scheme results in |A|=24.

Our consumer model mixes three types of loads. Assuming that for every consumer a baseline load curve has been fixed, we directly model load reductions against that hypothetical baseline.

We model *unconditional load reductions* as a discrete Gaussian process that is independent from applied actions. For each daytime h and actor i, the unconditional load reduction u(i,h) is a Normal-distributed random variable with mean 0 and standard deviation σ_u .

For some fixed $z \in [0,1]$ and each daytime h>1, u(i,h) is generated by $u(i,h) = z \cdot u(i,h-1) + (1-z) \cdot N(0, \sigma_u^2)$. From the linearity of mean and standard deviation of Gaussian variables follows that $u(i,h) \sim N(0, \sigma_u^2)$ for any $h = 1, \ldots, H$. In our experiments we work with fixed z=0.5 but with varying values of σ_u. The standard deviation σ_u will serve as a parameter to control the difficulty of the learning problem; a higher deviation means a more difficult instance.

In addition, each consumer is assumed to have some *curtailable load*, which represents energy consumption the consumer will suspend during the request interval. This curtailable load is generated uniformly at random as a constant value between 0 and 200W for each consumer.

Finally, every consumer has some *shiftable load*, which is specified as a load with a specific daily starting time, length, and magnitude. The magnitude is generated as a uniform random number between 0.5 kW and 1.0 kW. The length L is a uniform random number between 0.25 H and 0.5 H, and the start time is generated uniformly at random between time -0.5L and H-0.5L, such that up to 50\% of the shiftable load can take place outside the target interval [1,H]. This load can be freely shifted back and forth in time under the constraint that it remains within the union of its original time window and the target interval, and the consumer always tries to shift it such that its overlap with the request interval is minimized. Each consumer is with 50% probability cooperative enough to curtail the shiftable load when its overlap with the request interval cannot be reduced.

Summarizing the model, each actor aggregates three distinct types of load reduction potential with different characteristics. We emphasize that the only prior information about this model available to the learning agent is an upper bound on the expected load reduction potential of 2kW per actor.

This Project has received funding from the European Union's Horizon 2020 research and innovation programme under Grant Agreement N° 691735

Experimental Results

All experiments were performed with a horizon of 365 episodes, corresponding to one year in our application. Linear optimization was done via the pulp library for Python using COIN-OR as the backend optimizer. A runtime limit for the optimizer was set to 30s per episode, which turned out sufficient for a remaining optimality gap almost always under 5%. For computing the optimal assignment offline, 20 sample days were taken from the simulation and the optimizer was given enough runtime for an optimality gap of less than 1%. The algorithms were assigned an initial value of 2kW as the load reduction potential for each actor, action, and daytime.

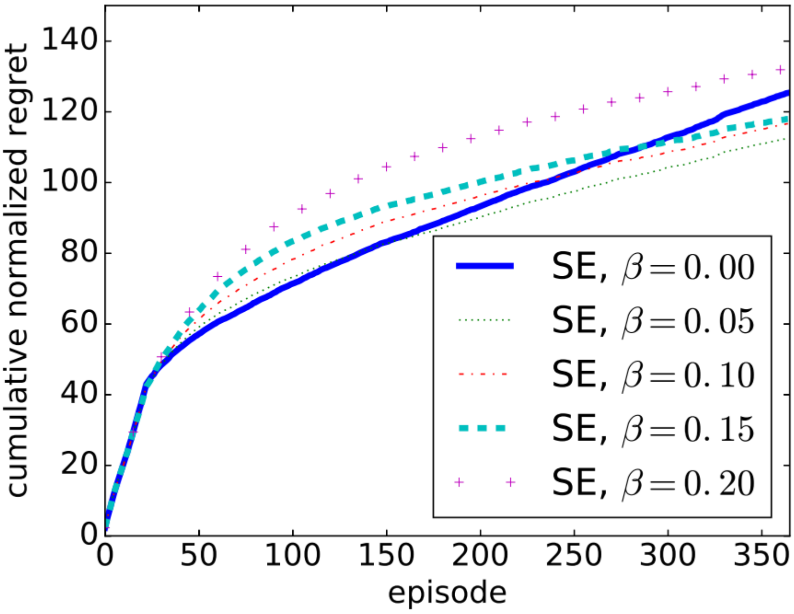

Figure 10: SE algorithm with $\beta \in [0, 0.2]$, $\tau = 0, \epsilon = 0$, *simulation with* $\sigma_u = 500W$.

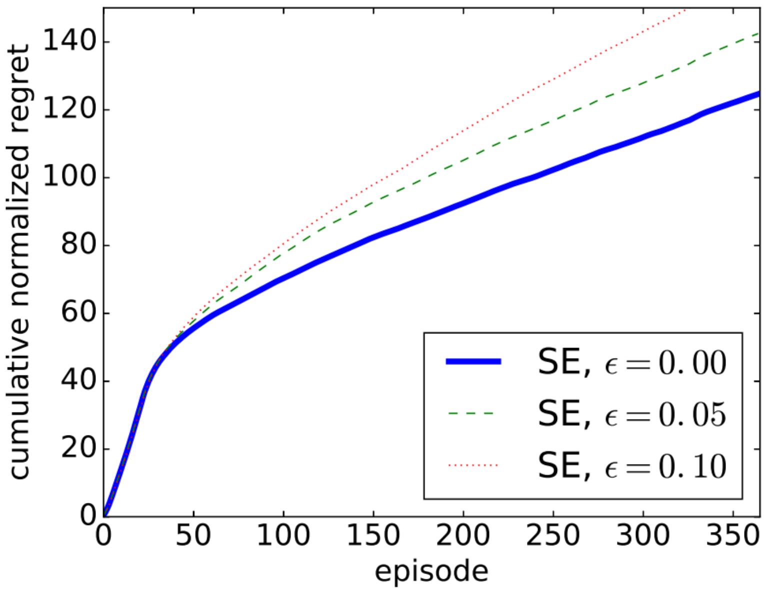

Figure 11: SE algorithm with $\beta = 0$, $\epsilon \in [0, 0.1]$, $\tau = 0$, simulation with σ_u=500W

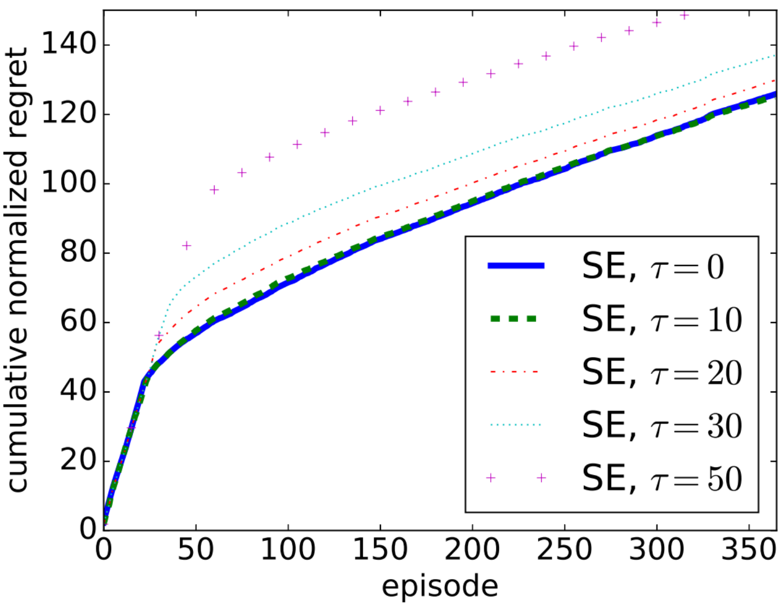

Figure 12: SE algorithm with $\beta = 0$, $\epsilon = 0$, $\tau \in [0, 50]$; simulation with $\sigma_u = 500W$.

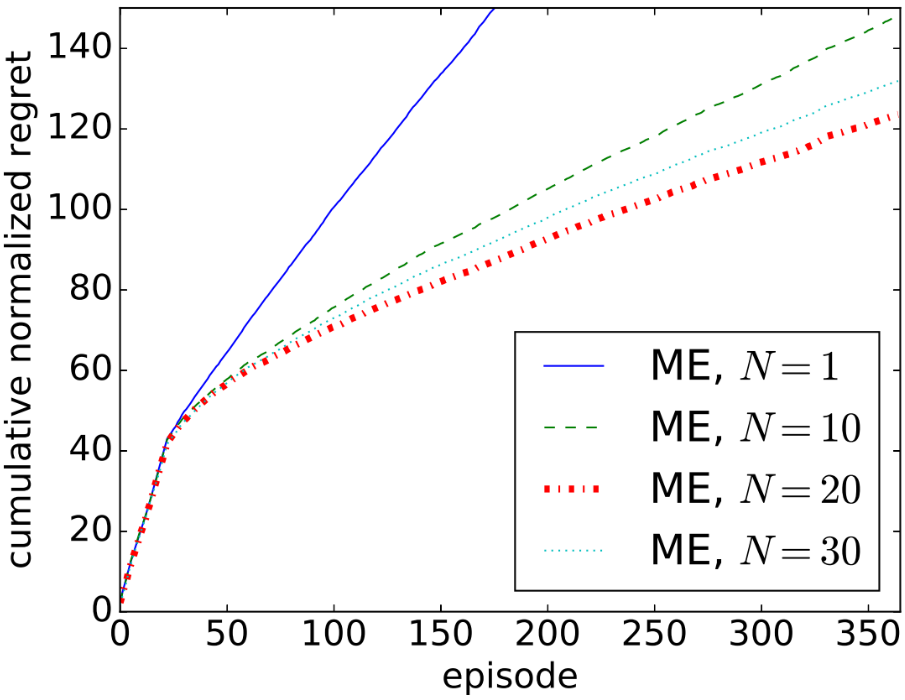

Figure 13: ME algorithm with $\boldsymbol{N} \in [\mathbf{1}, \mathbf{30}]$*; simulation with* $\boldsymbol{\sigma_u} = \mathbf{500}\boldsymbol{W}$*.*

In a first set of experiments the parameter spaces of the Single Episode (Figure 10) and Multi Episode (Figure 13) algorithm are explored under the condition of rather large daily load curve fluctuations with $\sigma_u = 500W$, making it difficult for learners to determine optimal assignments.

Due to the optimistic initialization, all 24 actions are tried at least once by each actor, resulting in a steep regret curve during that phase regardless of the algorithm and its configuration. The influence of parameter β, controlling for SE the amount of remaining optimism after an action has been taken once, exhibits the typical behaviour of a parameter trading off exploration and exploitation, as observable in Figure 10. Setting $\beta = 0$ is beneficial in early episodes, while values around 0.1 pay off later.

As Figure 11 shows, the well-known ϵ-greedy strategy does not improve exploration in the context of our experiments. Plausible reasons are that (a) ϵ-greedy does not distinguish between slightly suboptimal and clearly suboptimal actions, and (b) with 150 actors in each round the probability is high that some of them apply an action which does not combine well with the actions of the others. We set $\epsilon = 0$ in the remaining experiments. The final parameter

This Project has received funding from the European Union's Horizon 2020 research and innovation programme under Grant Agreement N° 691735

of SE is τ, the number of initial random exploration episodes; see Figure 12. Also here it turns out that only for moderate values the initially larger regret has the potential to pay off within a reasonable number of episodes.

For the Multi Episode algorithm we evaluate in Figure 13 the influence of the number of sample episodes N on the result. It turns out that the best performance is achieved for N=20. A possible reason why the performance again deteriorates for N>20 is the limited number of available samples per action. Building too many episodes from the same sample set at some point only increases the correlation among episodes and does not improve the reward estimator anymore. Furthermore, with more sample episodes it becomes computationally harder to determine a good action assignment within our runtime limit of 30s for the optimizer, thus the overall solution quality decreases.

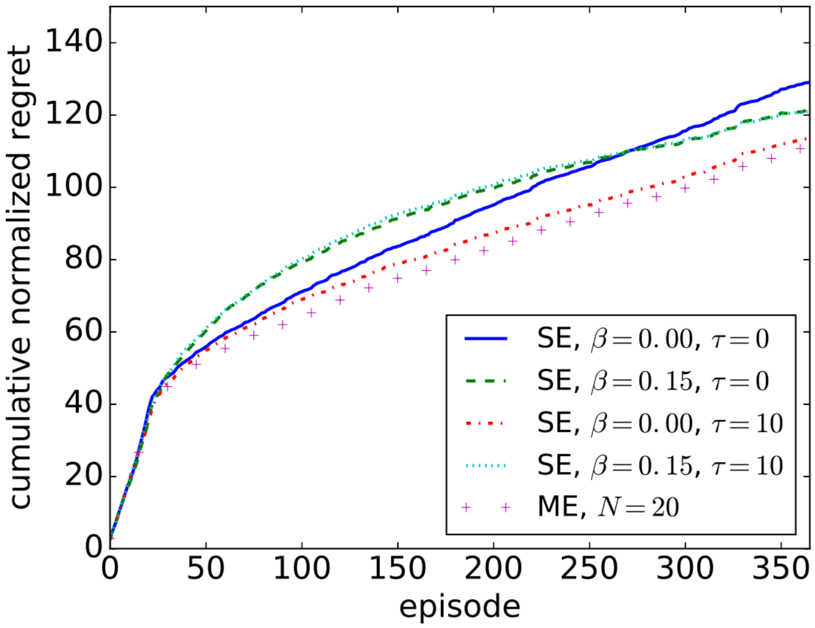

Figure 14: SE and ME algorithm applied to simulation with $\boldsymbol{\sigma_u = 500W}$.

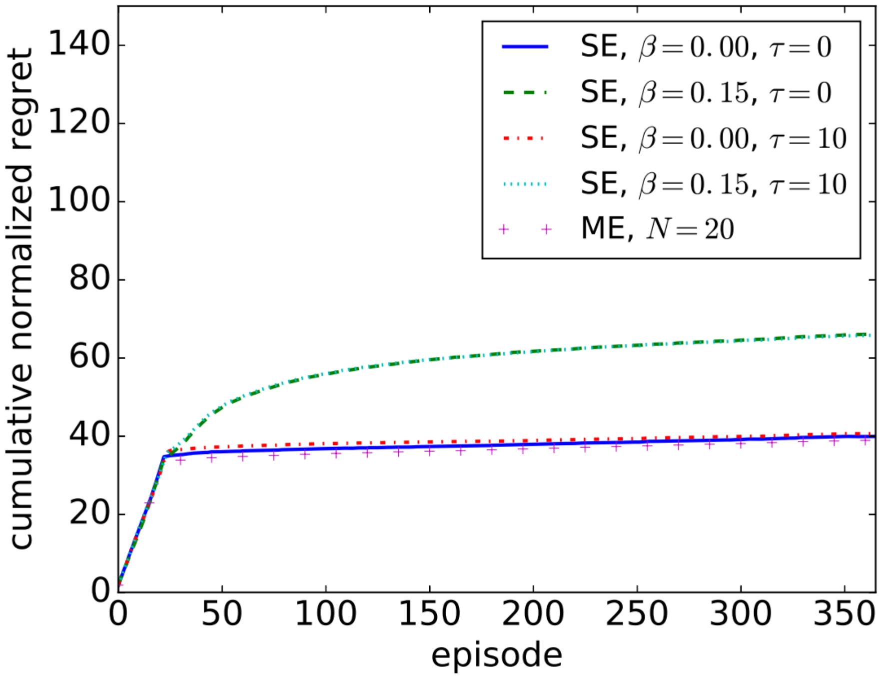

Figure 15: SE and ME algorithm applied to simulation with $\boldsymbol{\sigma_u = 100W}$.

In a final series of experiments we apply SE and ME using the most promising parameter configurations. The results for simulations with $\sigma_u = 500W$ and $\mu_u = 100W$ are visualized in Figure 14 and Figure 15, respectively. The first observation is that there is no fundamental advantage of using multiple episodes in our application; combining SE with some initial exploration seems to lead to a similar performance. An interesting observation can be made for the Single Episode algorithm using optimism ($\beta = 0.15$) to drive exploration. While slowing down convergence to the optimal assignment in case of low load fluctuations ($\sigma_u = 100W$), it seems to be the only mechanism which continues to significantly improve the assignment over the course of all 365 episodes. For $\sigma_u = 500W$, at the end of the curve in Figure 14, the curves with $\beta = 0.15$ have the least slopes among all tested algorithms, i.e. they apply the action assignments closest to the optimum. In general, Figure 15 shows that both ME and SE can find action assignments very close to the optimum in environments with little randomness.

This Project has received funding from the European Union's Horizon 2020 research and innovation programme under Grant Agreement N° 691735

6.3 Demand Side Response Trial

Demand Side Response (DSR) in the UK is currently limited to commercial premises primarily. There is no live marketplace for residential demand response. However in future with the UK's plan for a decentralised energy system, there will be an increasing role for residential level DSR as part of the smart grid.

A DSR trial is currently in design phase to roll out to a sample of the Smart Connected Homes. This design is being coordinated with colleagues at the University of the West of England and the University of Bristol with expert input from the local Distribution Network Operator (Western Power Distribution) and Bristol Energy.

The aim of the trial is to co-design a consumer-acceptable DSR programme, engage and recruit households to participate and, alongside this, deliver an ICT package of works that enables the technical delivery including linking control systems through the Smart City Platform and EDMS.

6.3.1 Objectives

The objectives of the DSR trial are to:

1) understand how and when a selection of householders currently use appliances in their homes and their engagement with programming/automation features
2) explore householders preferences and constraints for demand shifting
3) devise options for rewarding demand shifting behaviour (e.g. community spirit or 'green' citizen) from peak time energy consumption and engaging with demand side management (DSM)
4) design:

   a) a trial testing participants willingness to shift appliance usage and

   b) a set of user-informed DSM services
5) carry out a trial testing participants' willingness to shift appliance usage in response to reward/recognition packages

Dependent on development of Replicate Smart City Platform (SCP) and EDMS:

6) To trial user-informed DSM services in Replicate smart appliance households

6.3.2 Method

We are using the following method for the DSR trial:

1) Interview households including a set from the REPLICATE 'Smart Appliance' cohort and external to project to:
   a) Gather household demographic data
   b) Understand participants' appliance usage, engagement with programming/ automation features, and preferences, constraints and potential drivers/motivations for shifting appliance use
   c) Assess likely uptake of DSM measures
2) Analyse interviews to inform design of
   a) factors that would make it easier for participants to update a DSM service and/or abstract it
   b) DSM services
   c) rewards for participation
3) Carry out desk review of reward/recognition packages and household DSR trial designs in research, industry, business publications to inform participants at focus groups and trial design
4) Hold focus group(s) with participants, presenting findings from interviews as well as ideas from previous research and industry trials for discussion and to build consensus on reward/recognition packages, trial design and user-informed DSM services
5) Carry out trial with participants

6.3.3 ICT requirements

Figure 16 presents a summary of potential ICT requirements for a residential DSR trial. These are drawn from preliminary interviews carried out by University of Bristol on DSR services, previous DSR trial discussions and documents, and review of previous DSR trials. The interview and focus group(s) will aim to check the relevance of these requirements to participants as well as identify additional ones. It's recognised that some of these requirements may fall beyond what we can realistically include in a real-life DSR trial as part of REPLICATE but are likely to form part of future development of residential DSR/automation/smart energy & homes.

This Project has received funding from the European Union's Horizon 2020 research and innovation programme under Grant Agreement N° 691735

REQUIREMENT	DETAILS
Manual override of automation	Due to people not wanting or fearing loss of control
Override of data sharing	Due to people wanting control of what data is shared.
Real-time energy consumption information	For those interested in knowing more about their energy behaviour
Selective automation of different devices	People may have different preferences for control of different appliances e.g. shower versus washing machine.
Set automation context and profile for different devices	Preferences on days/times for automation/delays for different appliances
Set personal 'goal' for DSR service	For example, environmental or financial savings/targets and service is optimised relative to goal
Define automation preferences – day, times	Overall preferences for automation depending on personal/household schedule
Set 'comfort' levels	More applicable to heating/lighting settings than appliances – but relates to future work on 'Energy as a service'
Communication/interaction with friends/neighbours	To stimulate participation in DSR service or as part of 'community spirit' reward engagement for load shifting
Feedback on their environmental, energy consumption, grid, financial impact/savings	Recognised desire of interview respondents to receive feedback as part of DSR service – and necessary part of wider reward/engagement communication
Information on environmental, energy consumption, grid, peak time shifting etc for	For those interested in knowing more as part of DSR service and to test impact of information provision on demand shifting/whole trial engagement.

interest/education	
Opt-in or out option	Assumes automatic opt-in to trial to take away decision making from householders
Information on rewards (financial, points, prizes etc) received from trial	To provide information and stimulate continued engagement with trial
Information of dynamic pricing/tariffs	If possible to simulate as part of trial if there is no actual tariff – assumes future dynamic pricing/tariffs
Real time information on local distributed generation and storage	If possible to simulate as part of trial. Assumes future community/local energy schemes.

Figure 16: Initial set of DSR trial ICT requirements

6.3.4 Intended Outcomes and Indicators of Success

The study comprises of two elements:

Part 1: The interview and focus-group based research, where the requirements for the DSR services are elicited, and;

Part 2: Automated trial of the utility of the newly-delivered DSR service in shifting energy use form peak time.

The expected outcomes from the study are thus:

Part 1:

Outcome 1: Requirements Specification for a new DSR service to result from the interview and focus group study;

Part 2:

Outcome 2: Design and implementation of the new DSR service to be integrated into the REPLICATE Platform;

Outcome 3: Dataset on use of newly derived DSR service use to be collected on the REPLICATE platform;

This Project has received funding from the European Union's Horizon 2020 research and innovation programme under Grant Agreement N° 691735

Outcome 4: Report on the study of the efficiency of the newly delivered DSR services in shifting energy use to off-peak times.

The success of the study can be evaluated in terms of delivery of the above stated outcomes, which will imply that the following tasks are completed:

1) An interview study has been carried out with at 20 to 30 participants, from which data is analysed and formulated into a set of DSR automation service design requirements specification;
2) The requirements from Outcome 1 are integrated into the implemented REPLICATE platform (as per Outcome 2);
3) Data on smart appliance use under the "normal" (i.e., without use of the new DSR automation services) is collected at REPLICATE platform for a substantial time period (e.g. at least 3-4 weeks) to serve as a baseline for comparison against the automated DSR service provision.
4) Data on smart appliance use with the newly implemented DSR service use is collected for a substantial period;
5) Comparative analysis of the baseline and automated-DSR-use data is undertaken to establish the effect of the new DSR upon the energy use behaviour.

6.4 Smart Charging

Electric Vehicle charge-points represent one of the key flexible demand assets that will be managed by the EDMS. Management of EV charging is invariably termed 'smart charging', and involves managing the charging schedule (essentially, modulating the charging power levels over time) in a way that optimises one or more objectives subject to constraints. In this section we describe smart-charging algorithms that have been implemented by Route Monkey Ltd for the REPLICATE project, and we model their effectiveness using the data from approximately 120 charging sessions at REPLICATE charge-point sites that have been handled by the FIWARE system since September 2018.

6.4.1 Smart Charging: Preliminaries and Notation

The 'bread and butter' of smart charging is the charge session S, which has the following attributes:

- R, the amount of kWh required by the user;
- P, the power level, in kW, provided by the charge-point when it is switched on;
- Tstart and Tend, the start (plug in) and end (plug out) times of the session;

So, we can denote a session as S = (R, P, Tstart, Tend).

A charge schedule, SC, can then be modelled as, in general, a sequence of modulations M = (m1,m2,…mN), associated with a series of timeslots T = (t1, t2, …, tN). That is, SC = (M, T). From here on, we will make the common assumption that each timeslot is 15 minutes. Meanwhile, a 'modulation' is simply an instruction to the charge-point to modulate its power level; so we can model a single modulation as a proportion; e.g. '1' means fully on (e.g., where a 22 kW charge-point would deliver 22kW in that timeslot), and '0.5' means 50% (e.g. a 7kW charge-point would deliver 3.5 kW if modulated at 0.5). In most cases however, the only possibilities for modulation are 1 and 0; for simplicity, we will assume only 1 and 0 are available from now on, although the algorithms can easily be modified to account for more complex modulations.

6.4.2 Smart Charging: Charge Schedule Optimization

In a standard scenario without smart charging, the requirements of a session

S = (R, P, Tstart, Tend)

would be met by a charge schedule that looks like this:

(1, 1, 1, 1, 1, 1, 1, 1, 0, 0, 0, 0, 0, 0, 0)

This simply means that the charge-point will be turned on at *Tstart* and continuously charge at *P* kW for as long as is necessary to deliver the *R* kWh required by the vehicle, and then be switched off (or deliver very low 'trickle charge') for the remainder of the session until it ends at *Tend*.

In contrast, in a *smart charging* scenario, we will typically have a different distribution of modulations. This is in fact the essence of smart charging: instead of a simple 'leading 1s' charge schedule, the charging is distributed over time in order to optimise one or more objectives while meeting one or more constraints. In the above case, we may instead have a smart schedule as follows:

(1, 1, 0, 0, 0, 1, 1, 0, 0, 1, 1, 1, 1, 0, 0)

The shape of that schedule may reflect several factors at play: for example, the electricity price may be cheaper in slots 8—15. Hence, the hour of charging towards the end of the session

This Project has received funding from the European Union's Horizon 2020 research and innovation programme under Grant Agreement N° 691735

(slots 10—13) is driven by cost. But an initial half hour of charging is done for pragmatic reasons (despite the cost, ensuring that the driver has some charge if the vehicle is needed again sooner than expected). Further, some charging is done at slots 6 and 7 (rather than cheaper slots 14 and 15) may reflect the fact that the electricity cost is reduced in that period.

A central algorithm in smart charging is therefore the optimisation, subject to constraints, of the charge schedule according to one or more objectives. It is natural to represent an objective as a value per timeslot, aligned with the charge schedule. For example, the series of contrived costs below would correspond to the example scenario above:

(22, 22, 22, 22, 22, 22, 22, 15, 15, 15, 15, 15, 15, 15, 15)

However we may instead, or in addition, be interested in how 'green' is the electricity in those slots. The objective represented below, for example, may represent the percentage of renewable sources in the grid mix at each timeslot:

(18, 20, 21, 21, 20, 20, 20, 20, 18, 18, 16, 15, 12, 10, 10)

In this case the 'leading 1s' schedule would actually be the greenest, but also the most expensive.

More formally, if we have k objectives, each represented by a time-slot aligned vector of costs Ck, then the task of a smart charge optimisation algorithm is to minimise:

$$\sum_k wk\,(M \cdot Ck)$$

where wk represents the importance attributed to objective k, subject to the constraint:

$$P.\sum_k mk = R$$

i.e. that the sum of kWh delivered over the timeslots meets the requirement; and also subject to other constraints that may apply.

This optimization task in itself is not a challenging one – in most scenarios we can simply sort the timeslots by weighted cost and then choose the best-weighted slots until we have met the charge requirement. It becomes more challenging when intermediate levels of modulation are allowed, or where we cannot assume a 'flat' charge curve, however in general this is an optimization task that can usually be achieved in real time.

However, the much more challenging tasks involved in smart charging come from the uncertainties around the charge sessions and the objectives themselves. If we know when the session will start and how long it will last, and what the objective values will be throughout the session, then the optimization task is achievable, as indicated above. But in general (i) we do not know when a charge session will start, and (ii) we do not know how long it will last, and (iii) for many of the objectives of interest, we do not know for sure what the values will be in advance. This brings us to the next section.

6.4.3 Smart Charging: Forecasting of Charge Sessions and Objectives

In general, as indicated above, we do not know in advance when a charge session will start. In some scenarios, forecasting of charge sessions in advance is needed, in order to help balance overall energy management. For example, if we can estimate the charge session demand over the next 3 hours (much of which may come from vehicles not yet plugged in), this helps us understand what current management constraints or decision to make in respect of storage battery assets and other flexible loads. Note that this presumes that predicting the start of a session also involves predicting the length and the charge requirement of that session. Meanwhile, once we know that an EV has just plugged in, and have estimated the session length and charge requirement, we need to be able to forecast the objectives of interest. The example objective we will look at later is the grid mix, which is never known in advance, and is dependent in large part on what wind-speed and cloud-cover will be, at the full range of renewables sites feeding the local grid, over the duration of the upcoming session.

Broadly speaking, smart charging therefore requires two types of algorithms in addition to the basic charge schedule optimization algorithm: (i) forecasting algorithms for the data-streams associated with the objectives; (ii) forecasting sessions and session parameters.

The first type of algorithm essentially requires standard machine learning algorithms for multi-attribute time-series forecasting. All of the likely objectives - such as overall demand levels, grid mix, price - are time series of values that are traditionally forecast with a well-known range of algorithms. Among the better-performing approaches, across a wide range of scenarios, are ensembles of regression trees. Gaussian process regression, and weighted variants of the k-nearest neighbour algorithm. Route Monkey's approach uses a mixture of these approaches, and relearns and adapts continually as new data are obtained.

The second type of algorithm is considerably more complex. Although EV demand can also be modelled as a time series of values, represented as a long string of zeroes when the vehicle is not plugged in, forecasting becomes much more effective if it models the 'session-based' structure driver behind the values. At the scale of large aggregations (e.g. EV demand across a region), the time series approach is appropriate, however when small numbers of EVs are involved it is more accurate to use a session-based approach. Work in this domain in the academic literature has used approaches based on queuing theory; by fitting parameters to historical data on arrival and departure (plugin and plugout) times, potential future sessions can be predicted by a simple parametric model. However, an alternative way to use historical session data is to feed it into a nonparametric simulation model, which essentially uses previous data (with appropriate arrangements to favour recency) to build transition matrices for Monte Carlo simulation models that can generate future forecasts. This method also allows probabilistic forecasting, essentially by running the model several times to obtain confidence bounds. This is the approach that we have used for session forecasting in a number of recent projects, and is the intended approach for REPLICATE.

6.4.4 Analytics of Charge Sessions at REPLICATE Charge-posts

In order for smart charging to be possible at all, whether or not we can accurately predict sessions or objectives, there needs to be flexibility in the charge session. That is, the charge session must involve more than enough time to satisfy the charge requirement. For example, if a vehicle needs 22 kWh of charge, and it plugs in to a 7.7 kW charge-point for 3 hours or less, it simply needs to be charged continuously, and there is no opportunity to optimise the session. However if instead it is plugged in for 6 hours, then we can choose any 3 of those 6 hours to charge; more flexibly still, we can consider the 6 hour session as 24 15-minute slots, and we can choose any 12 of those (there are 2.7M possibilities).

In Figure 17 we illustrate the charge sessions recorded at REPLICATE charge posts from September through December 2018, plotting a measure of the 'flexibility' of each session against its charge requirement. Here, flexibility is simply modelled as the number of 'spare hours' that remained in the session after it had already met the charging requirement. We can clearly see that there is a fair degree of opportunity for smart charging, but to get a clearer picture we focus on the more numerous shorter sessions in Figure 18. The majority of sessions seem to require around 5 kWh and have little spare time, but many sessions do seem to offer both higher kWh requirements and greater flexibility.

This Project has received funding from the European Union's Horizon 2020 research and innovation programme under Grant Agreement N° 691735

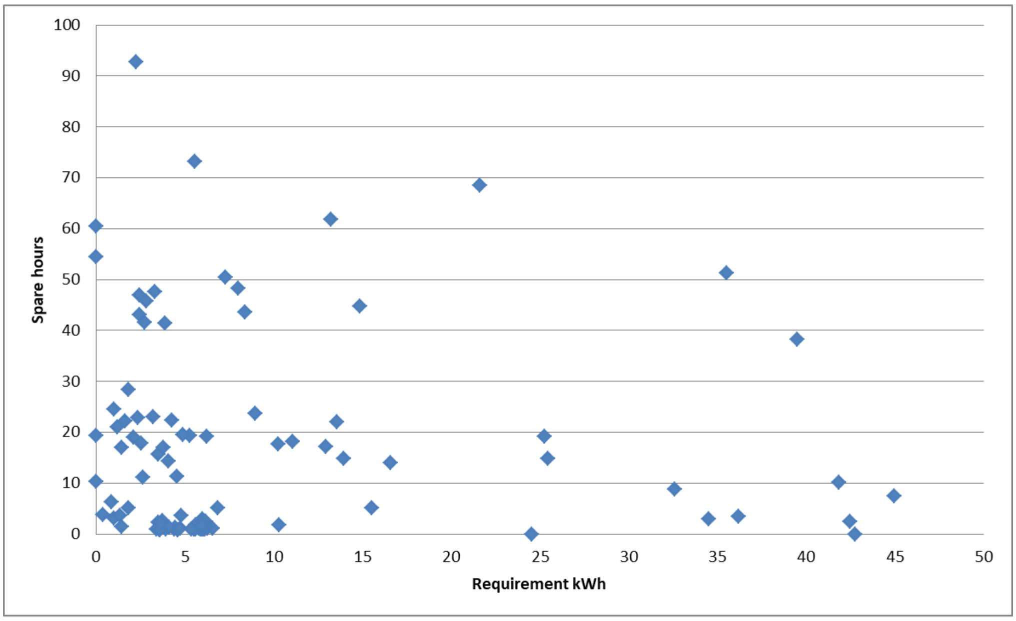

Figure 17: Flexibility of collected charging sessions

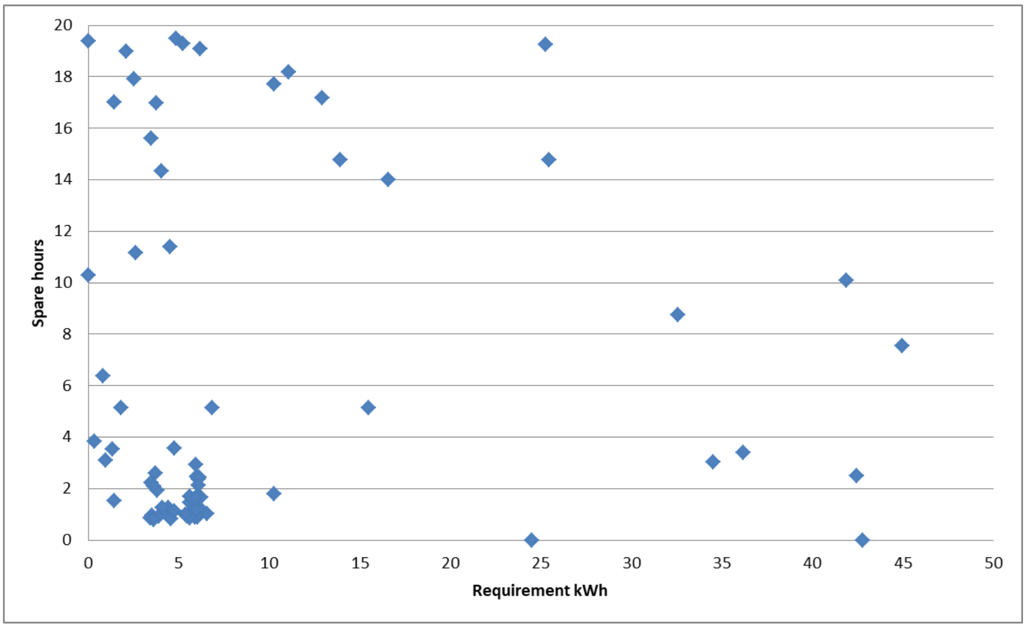

Figure 18: Flexibility of collected charge sessions; focus on shorter sessions

This Project has received funding from the European Union's Horizon 2020 research and innovation programme under Grant Agreement N° 691735

For a related angle on the smart charging opportunities, Figure 19 shows session lengths plotted against their starting hour. This figure also omits longer sessions, so we can better see the distribution of shorter sessions. It is fairly clear from this figure that sessions that start in the afternoon or evening tend to have a better chance of being longer, and hence allowing for smart charging opportunities (e.g. these will often be overnight sessions), however the fact that there are many short sessions that also begin at such times means that our confidence in predicting session length will likely be low, so that smart charging will need to involve some initial 'leading ones' to ensure some charge is supplied even if the objective values start out poorly.

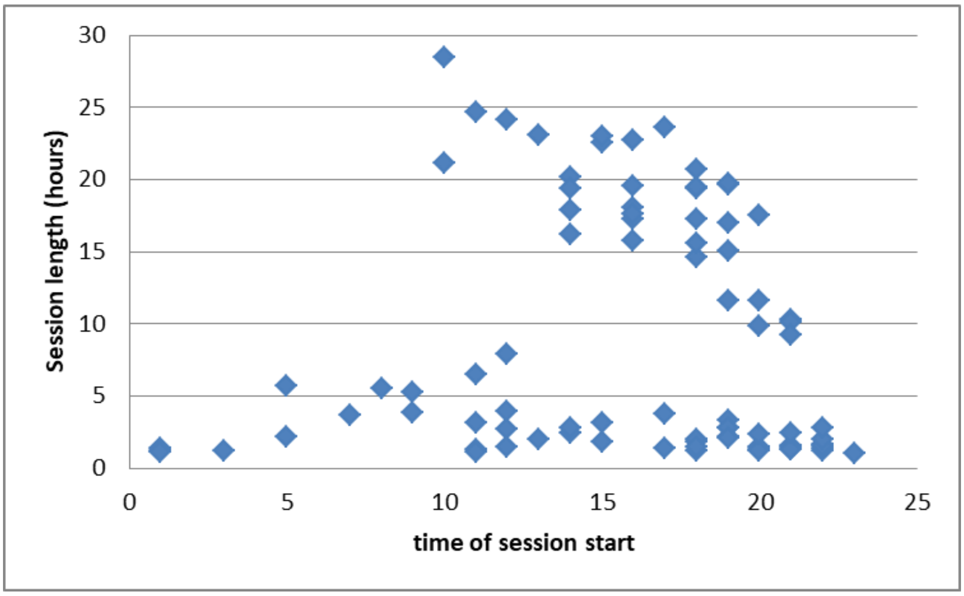

Figure 19: session length against start time; focus on shorter sessions

6.4.5 Modelled Smart Charging Outcomes

In this section we will quantify the potential for smart charging at REPLICATE sites using the data so far collected. As we have identified in the previous section, the many occurrences of sessions with 'spare hours' suggest there are clearly opportunities, but these may be reduced in practice given the difficulty of predicting charge session lengths (and hence we need to be cautious and provide an initial safety charge, which reduces flexibility). Error in forecasting of objective values introduces additional complication, however (from comprehensive experience), although we know that these errors can typically be kept at or below 10% by using state of the art machine learning approaches.

This Project has received funding from the European Union's Horizon 2020 research and innovation programme under Grant Agreement N° 691735

We will model outcomes under the assumption that our objective is to optimize the degree to which charging is done with 'green energy'. In Figure 20, the blue line shows the UK green mix (as a percentage) for the period from 6th to 16th December 2018. The raw data for this comes from http://www.bmreports.com. The red line models kW draw from REPLICATE recorded charge sessions, under the assumption that the draw is typically 80% of a charge-point's stated kW value, and that charging was continuous from plug-in until the point at which the kWh requirement was reached. The green line, however, represents the continuation of charge sessions beyond this point – i.e. it represents plugged-in vehicles that are not currently drawing power, and hence represents flexibility (since some or all of the charging could have been done during this time instead). It is clear from the figure that the green mix varies substantially within some of these plugged-in periods, suggesting opportunities for smart charging.

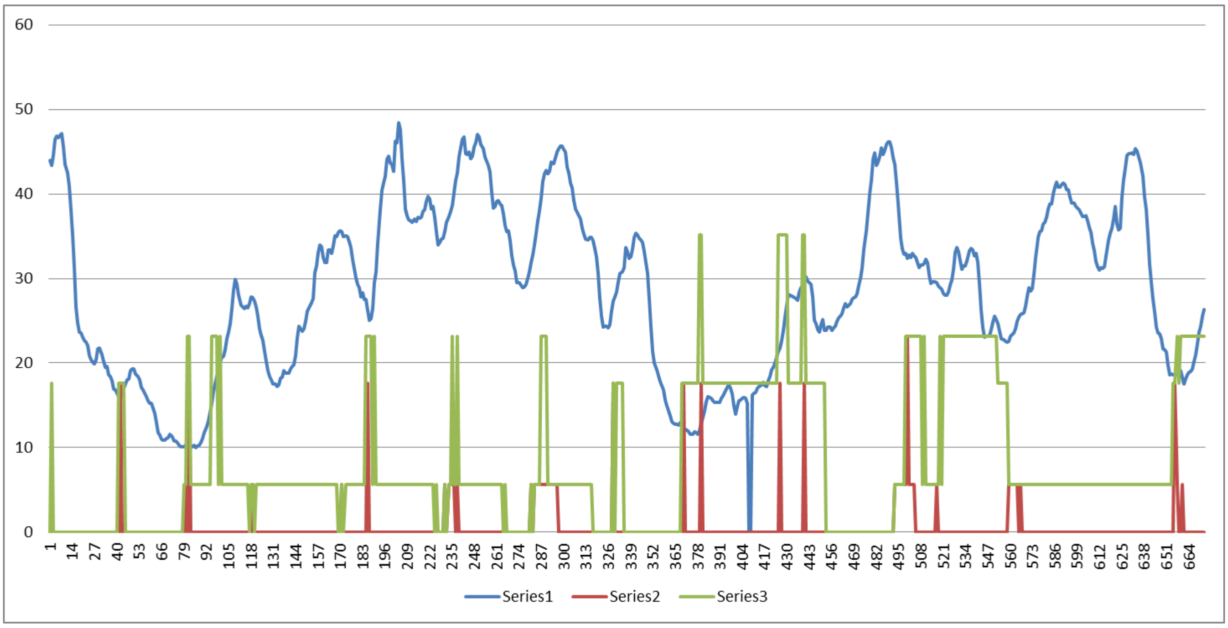

Figure 20: illustrating grid mix over a 12-day period in December 2018 (blue), contrasted with kW draw from REPLICATE EV sessions assuming 'leading-ones' charging (red). The green plot indicates times when one or more EVs remain plugged in, but are not actually drawing energy.

To obtain quantified estimates for how 'green charging' outcomes could be enhanced in the REPLICATE scenario, we took the real charge session data for September to December 2018, and used it in a simulation in which, at the onset of every charge session, the charge session length was predicted using Monte Carlo simulation based on previous history, and the charge schedule optimization was done on the basis of two parameters: an assumed forecast error e, and an initial safety charge time, m. For example, e = 10% and m = 60 represents a situation where the forecasts for the green mix were on average 10% in error, and the charge session started with 60 minutes of continuous charge, reducing the slots available later (if any) for smart charging. Each individual scenario was evaluated by calculating the total kWh in the sessions that could be attributed to renewable sources (based on the actual, rather than forecast, data).

Figure 21 summarises the outcomes of these experiments. The baseline figure is 122.7 kWh, which is the 'green kWh' consumed by the actual charge schedules (under the assumptions that no smart charging was in operation). If we estimated the green mix perfectly and did not schedule any initial safety charge, the outcome is 151.3 kWh (the leftmost bar in the leftmost group), representing a 23% improvement over the baseline. With parameters that are reasonable in practice – 10% forecast error and a 60 minutes safety charge – we get 131.7 kWh, a 7.3% improvement over the baseline. These are relatively promising figures considering the large number of short sessions in the data. The sharp drop in improvements within each group of bars, as we increase the initial safety charge, reflects the reduced leeway we have for smart charging that results by being forced to commit increasing numbers of slots to the beginning of the schedule. However, it should be noted that the need to predict session length, and also kWh requirement, are both worst-case for a smart charging scenario. A variety of schemes could be available for consideration, whereby EV users are encouraged to provide their expected plugout times and/or state of charge in advance via an app, and/or by making use of telematics, and so forth. Such arrangements reduce the need for safety charging and can lead to (as these data suggest) marked improvements in smart charging outcomes.

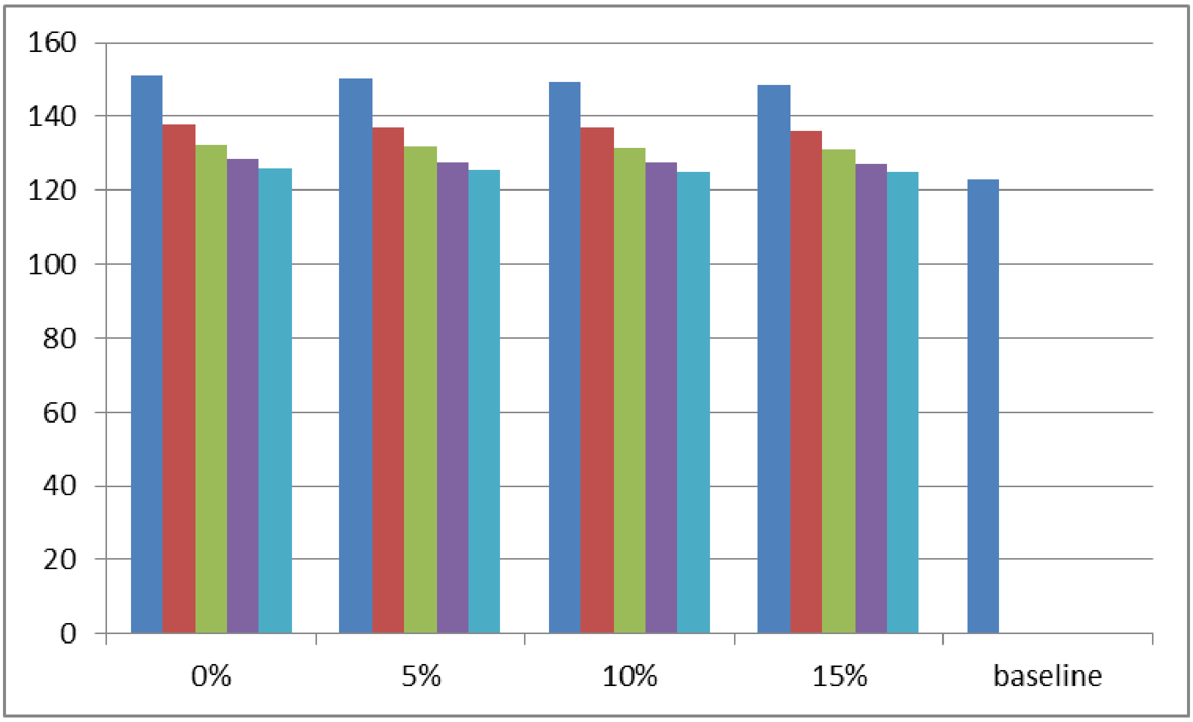

Figure 21: outcomes of smart charging simulations using REPLCIATE charge sessions; the vertical axis is 'green kWh', and the horizontal axis is forecast error. Each group of five bars, from left to right in the group, represent safety minutes of 0,30,60,90, and 120.

7. INNOVATIONS, IMPACTS AND SCALABILITY

7.1 Innovation solution

REPLICATE has allowed project partners to carry out a great deal of innovative work in order to facilitate the set up of the Energy Demand Management System.

REPLICATE has successfully connected a wide range of consumers and prosumers to the Energy Demand Management System via the FIWARE Smart City Platform and proprietary cloud systems. Different approaches were applied for different devices with a VPN solution controlled by a miniserver developed, tested and implemented for the Smart Homes. For Smart Mobility, existing systems were integrated using Bristol City Council systems to push data to the Smart City Platform. Once data was on the Smart City Platform it could be used for a variety of applications. NEC has demonstrated that the EDMS can see the state of devices in real time.

NEC has developed an innovative and highly sophisticated approach to modelling the behaviour of a large number of independent energy consumers using the Combinatorial Multi-Bandit approach to seek to demonstrate theoretically that it is possible to optimise energy use across a range of actors. This approach has been designed to be highly transferrable to different problems and does not depend on extensive knowledge of the system or actors.

NEC has also developed a range of tools that allow users and program managers to interact with their data.

NEC has demonstrated that it can remotely control domestic appliances by integrating with the smart appliance supplier's proprietary systems.

University of Bristol and Bristol Is Open have gained extensive knowledge of how to design and build a FIWARE based Smart City Platform that can accommodate a wide range of devices and allow applications to connect to them in real time.

The ongoing Demand Side Response trial will bring together all these elements into a more sophisticated trial that demonstrates how all these elements can be combined to produce prototype products and understand how attractive they will be to users.

Real, session-level EV charging data, with plug-in, plug-out times and kWh delivered, is notoriously difficult to obtain. Even with preferential access to some clients, Route Monkey has before now had access to very little in the way of such EV data from public (rather than fleet) users. The data obtained from REPLICATE so far (and continuing to accumulate) has been very useful for helping us rethink aspects of our approach to building our smart charging algorithms suite. These aspects include: (i) handling the many 'spurious' artefacts that occur in real charging event data streams – such as users who rapidly switch between different connectors on the same charge post, presumably attempting to deal with some physical connectivity issues; (ii) understanding the variability in public user charge sessions – this is rather wider than we had envisaged, with more in the way of very short and very long sessions; (iii) shifting some of our focus towards 'information poor' smart charging, in which state of charge, length of session, and other data (e.g. recent mileage from telematics) are not reliably available.

7.2 Social impacts

Whilst the KPIs and definitions of success for the EDMS have not been specified in as much detail as they have been for say building refurbishment or mobility, they are still expected to provide significant benefits.

On the social side it is expected that as consumers become more aware and in control of their energy use they will be able to reduce their environmental impact.

With the Energy Demand Management System there will be a number of types of changes to behaviour expected. For Smart Homes, the EDMS envisages the consumer could offer to make unconditional load reductions, curtail their load or shift their load. All of these could result in cheaper energy (or other benefits). Likewise, the DSR trial could reward consumers that are willing to change their habits and evaluate how successful different strategies are.

7.3 Environmental impacts

As for social impacts, the EDMS, DSR trials and Smart Charging could lead to environmental benefits. Again, there aren't specific KPIs for these. The nature of EDMS type interventions is that the benefits will be more complex because sometimes the energy use will be deferred rather than reduced. Whilst the reduction of peak loads and time shifting to periods of lower carbon intensity are highly desirable they will be much more complex to determine.

Peak loads are expensive and are likely to require the use of the least environmentally friendly energy sources so every peak kWh that can be shifted is of significant environmental benefit but it won't show up in simple measures such as total energy use.

As part of the monitoring we will consider what estimations we can make of these benefits but they won't feature in the core KPIs.

In pure energy terms the loads required for EVs are very considerable and so the potential savings in this area could be quite major.

Initial indications that we are on track in terms of environmental benefits are good. For example, the Counterslip Renault Zoe used 405kWh in December 2018 which equates (conservatively) to about 1500 miles travelled (based on 150miles on a full 40kWh charge). This is well above the 583 miles per month which it was estimated in the bid would produce a saving of 1.57 tonnes per vehicle per annum (based on typical Internal Combustion Engine emissions of 140g CO_2/km).

Very early indications from our optimisation work show there is substantial potential through smarter charging to shift EV charging into periods when the energy mix is greener leading to significant environmental benefits.

7.4 Replication and scalability potential

REPLICATE is enabling a Smart Energy and Smart Mobility trial within Bristol as one of three EU lighthouse cities. We have established connectivity between IOT devices such as smart appliances, electric car charging points and smart meters in order to monitor and control energy demand within the city. There has been a cultural shift in the user community both by allowing a third party to supply and control household devices and by contributing to this initiative which aims to improve air quality, lower energy demand spikes and improve individual carbon footprints. It also aims to give data to support the uptake of these devices and follow trends in their use.

We are supporting a number of academic studies into smart homes, IOT, energy demand management and sustainable living, together with increasing citizen engagement and collaboration between the local authority, university, industry and individuals.

Smart technology has given us the ability to control household appliances remotely. This allows households to opt into the control, both being able to switch devices on and off from their own smart phones, or to relinquish control and trust in the automated service to optimise their home energy demand and switch the device on at the time of lowest demand to prevent supply overburden. Where we feel our approach has been particularly strong has been in the ability through the use of the Loxone MiniServer to drill down into the data to attribute it to individual devices rather than the whole property as would happen with just a smart meter.

We have successfully connected the EV charge points to the EDMS via the SCP. This was done using the Open Charge Point Protocol and FIWARE. We chose the Open Charge Point Protocol as this showed the greatest promise as a format which will allow roaming between charge points.

We have found that the Charge Point Management System has been the weak link in the chain in terms of functionality. We would recommend that close attention is paid to using open protocols and ensuring that the charge points and back office system provide compatible levels of smart grid functionality.

The historic data is also allowing us to start to build pictures of different charging behaviours and profiles. Going forward understanding theoretical and actual load curves for different back office <-> charge point <-> cable <-> car combinations will be important as these will be key inputs in to the energy demand management system.

As noted in Deliverable 5.7, we are aware that there are currently significant gaps in our knowledge of vehicle status and driver intentions. By seeking to bring these diverse information sources together we intend to ensure that consumers can make better decisions about charging.

We have only just started to investigate the impact of new charging behaviours on the local distribution network. We are aware that our Distribution Network Operator has concerns about how the roll out of Electric Vehicle charging points will affect their network and we consider it very likely that the ability to offer Energy Demand Management Services should make developing new charge points easier and cheaper by using smart methods rather than having to plan for the worst possible case (which is often prohibitively expensive). The other part of this jigsaw will be Time Of Use Tariffs which are likely to have a significant impact on how willing consumers are to vary their behaviour.

The flexibility of the NEC EDMS system means that, with some customisation, it can be applied to a wide range of energy demand management situations. We note that, given that EDMSs will often need to interact with proprietary third party systems in order to be able to control devices remotely, it is unlikely that FIWARE will always be the whole solution, and so flexibility is important such that as wide a range of devices as possible can be incorporated.

The Route Monkey algorithm optimisation work has led to the interesting lesson that, even with a low prior information basis, positive smart-charging outcomes are quite achievable. This is particularly intriguing, since we can expect increasing opportunities for smart charging services where, although the chargers may be centrally controllable (by, for example, a city council), the variety of users and vehicles involved means that no system will be available to provide reliable information in addition to notifications of the plugin and plugout events themselves; nevertheless, it seems that positive outcomes could still be achieved from smart charging, even assuming a reasonable level of safety charge, and these outcomes will naturally scale. Lessons from this optimisation work in Bristol will help inform Deliverable "D7.5 – Report On Management Models V2".

7.5 Economic feasibility

We expect the DSR trial will provide clear evidence of the potential benefits that such systems could offer to consumers who are willing to offer flexibility.

As noted above, the scale of EV fleets is such that companies who control large EV fleets would be expected to be able to offer quite large potential energy shifts which would be expected to become commercially viable providing the regulatory frameworks and the technologies, such as the EDMS, allow the right level of granularity and trust such that aggregators can be confident about the savings they can offer and a value that can be placed on these. At this stage this is a hypothesis but we hope that the data we collect during the monitoring period will allow more quantification of the scale of benefits.

7.6 Impact on SMEs

We have worked with a range of small-to-medium, dynamic and innovative companies in developing the EDMS.

Trakm8, Route Monkey's parent company, is an SME and has been developing energy optimisation algorithms thanks to its involvement in REPLICATE.

Installations were managed by Narec Distributed Energy, a small specialised energy company.

As part of the rollout out of smart appliances older inefficient machines are being donated to The Sofa Project, a local social enterprise, for either recycling or reuse amongst lower income households.

Loxone Group, who produce the MiniServer that has been key to this project, are a young, dynamic and innovative company, although they have recently moved out of the SME category given that their 2017 turnover was €57.5million. Their knowhow has been key to the successful implementation of this project and their fast growth shows the vibrancy of the smart energy sector.

This Project has received funding from the European Union's Horizon 2020 research and innovation programme under Grant Agreement N° 691735

8. CONCLUSIONS

As part of the REPLICATE project we have successfully developed and rolled out an Energy Demand Management System (EDMS) which connects with Smart Home and Smart Mobility devices using a newly developed Smart City Platform.

The general functionality of the EDMS is monitoring and control of energy consumption and production. The EDMS also acts as a marketplace, where multiple stakeholders can create energy management programs.

EDMS Theoretical model

In developing a theoretical model of energy consumption in Bristol we aimed to be able to combine a number of assets to build up a critical mass of flexibility, whilst at the same time allowing for the fact that load curves of private households do not follow an easily predictable daily pattern and it is not practical to require a guaranteed level of demand elasticity from private households. For these reasons, we chose an algorithmic framework which requires minimal assumptions and minimal prior knowledge on the consumer behaviour, has self-adaptation of system to responsiveness of consumers and has the ability to optimise a global objective across all consumers.

We have developed a Reinforcement Learning approach using a Combinatorial Multi-Bandit Problem. In it, when an actor takes an action, the result can be observed as the state of its environment, but it is initially unknown how each action influences the environment. The objective is to maximize the total expected reward, which is the combination of outcomes of the individual actors by a known reward function. The optimal or near-optimal assignment of actions has to be learned from experience using a strategy that balances *exploration* (choosing actions for the purpose of learning their outcome distribution) and *exploitation* (choosing action assignments known to perform well).

In the simulations it is assumed that each of the consumers has some load that can be shifted to a different time (e.g. operation of a washing machine), as well as some smaller amount of load that can be shedded (e.g. turning off some unnecessary lights). Furthermore, the base load of every consumer fluctuates. The target of the management system is to coordinate the load reduction requests such that throughout a pre-specified target time interval as much load as possible is avoided.

This Project has received funding from the European Union's Horizon 2020 research and innovation programme under Grant Agreement N° 691735

We tested a number of strategies to optimise the parameters of the algorithm. We then ran Single Episode and Multiple Episode experiments using the most promising parameter configurations. We found that there is no fundamental advantage of using Multiple Episodes in our application; combining Single Episode with some initial exploration seems to lead to a similar performance.

Demand Side Response Trial

Demand Side Response (DSR) in the UK is currently limited to commercial premises primarily. There is no live marketplace for residential demand response. However in future with the UK's plan for a decentralised energy system, there will be an increasing role for residential level DSR as part of the smart grid. A DSR trial is currently in design phase to roll out to a sample of the Smart Connected Homes. The aim of the trial is to co-design a consumer-acceptable DSR programme, engage and recruit households to participate and, alongside this, deliver an ICT package of works that enables the technical delivery including linking control systems through the Smart City Platform and EDMS.

Smart Charging

Using early EV charge point data, Route Monkey's optimisation work has shown that there is significant potential to vary charging schedules whilst still allowing sufficient time to leave the EV fully charged for the next user. By combining this with real data on grid carbon intensity we have shown that varying the charging profiles can bring about substantial carbon savings.

The data obtained from REPLICATE has been very useful for helping us rethink aspects of our approach to building our smart charging algorithms suite. These aspects include handling the many 'spurious' artefacts that occur in real charging event data streams, understanding the variability in public user charge sessions and shifting some of our focus towards 'information poor' smart charging, in which state of charge, length of session, and other data (e.g. recent mileage from telematics) are not reliably available. The other area of potential future effort is to go from an 'information poor' to an 'information rich' scenario which would allow much better predictions of future behaviours. A variety of schemes could be available for consideration, whereby EV users are encouraged to provide their expected plugout times, expected journey length or state of charge in advance via an app, or by making use of telematics. We hope to explore these further as the project continues.

This Project has received funding from the European Union's Horizon 2020 research and innovation programme under Grant Agreement N° 691735

Innovation

A significant amount of innovation has resulted from the REPLICATE EDMS work.

We have developed an innovative and highly sophisticated approach to modelling the behaviour of a large number of independent energy consumers using the Combinatorial Multi-Bandit approach to seek to demonstrate theoretically that it is possible to optimise energy use across a range of actors. This approach has been designed to be highly transferrable to different problems and does not depend on extensive knowledge of the system or actors. We have also demonstrated that the EDMS can remotely control domestic appliances by integrating with the smart appliance supplier's proprietary systems.

We have gained extensive knowledge of how to design and build a FIWARE-based Smart City Platform that can accommodate a wide range of devices and allow applications to connect to them in real time.

Social Impacts

The REPLICATE EDMS trial is resulting in significant positive social impacts. It is expected that as consumers become more aware and in control of their energy use they will be able to reduce their environmental impact. They will increasingly be able to offer flexibility resulting in cheaper energy (or other benefits). Likewise, the DSR trial could reward consumers that are willing to change their habits.

Environmental Impacts

The REPLICATE EDMS trial is resulting in significant environmental benefits. The nature of EDMS type interventions is that the benefits will be more complex because sometimes the energy use will be deferred rather than reduced. Peak loads are expensive and are likely to require the use of the least environmentally friendly energy sources so every peak kWh that can be shifted is of significant environmental benefit. In pure energy terms the loads required for EVs are very considerable and so the potential savings in this area could be quite large.

Scalability and Replicability

A number of the findings of the REPLICATE EDMS trial will be particularly relevant to replicability and scalability.

We feel our use of the Loxone MiniServer to drill down into the data to attribute it to individual devices rather than the whole property has been key. On the EV side we have found that the use of an open protocol such as the Open Charge Point Protocol will allow more flexibility and innovation in the future. We would like to draw attention to the importance that back office systems, as well as the devices they control, have as much smart grid functionality as possible.

This Project has received funding from the European Union's Horizon 2020 research and innovation programme under Grant Agreement N° 691735

We have only just started to investigate the impact of new charging behaviours on the local distribution network. We are aware that our Distribution Network Operator has concerns about the impact Electric Vehicle charging points will have on their network and we consider it very likely that the ability to offer EDMS type services should make developing new charge points easier and cheaper by using smart methods rather than having to plan for the worst possible case (which is often prohibitively expensive).

The flexibility of the NEC EDMS system means that, with some customisation, it can be applied to a wide range of energy demand management situations. We note that, given that EDMSs will often need to interact with proprietary third party systems in order to be able to control devices remotely, it is unlikely that FIWARE will always be the whole solution, and so flexibility is important such that as wide a range of devices as possible can be incorporated.

Economic Impact

The REPLICATE EDMS trial and the DSR trial will provide clear evidence of the potential economic benefits that such systems could offer to consumers who are willing to offer flexibility.

The scale of EV fleets is such that companies who control large EV fleets would be expected to be able to offer quite large potential energy shifts which would be expected to become commercially viable providing the regulatory frameworks and the technologies, such as the EDMS, allow the right level of granularity and trust such that aggregators can be confident about the savings they can offer and a value that can be placed on these.

Impact on SMEs

We have worked with a range of small-to-medium, dynamic and innovative companies in developing the EDMS. Through the extensive teamwork of the REPLICATE team these SMEs have been able to develop and refine their products by testing them in a real city environment.

9. APPENDICES

Appendix A – Smart Home Equipment

Your Smart Home Equipment

Please find below details of the smart home equipment to be installed in your home. Please remember that these will be yours to keep providing you adhere to the Terms and Conditions of the project.

Smart Appliance

Your smart appliance will be an energy efficient Samsung washer, dryer or dishwasher. Further information regarding the models and specifications can be found below.

Smart Washing Machine

Samsung AddWash WW90K5410WW/EU Washing Machine, 9kg Load, A+++ Energy Rating, 1400rpm Spin, White.

Dimensions H85cm x W60cm x D55cm

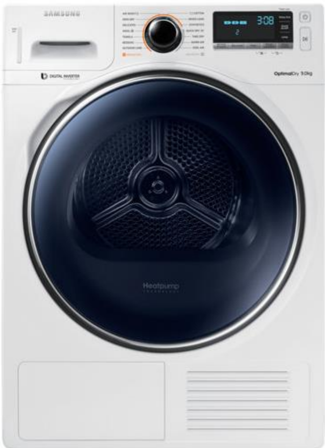

Smart Dryer

Samsung DV90M8204AW/EU Tumble Dryer, 9kg, A+++, White

Dimensions H85cm x W60cm x D65cm

Smart Dishwasher

Samsung DW60M9550BB/EU Dishwasher A+++ Full-size Integrated Dishwasher. PLEASE NOTE WE CAN ONLY PROVIDE INTEGRATED DISHWASHERS NOT STAND ALONE MODELS

Dimensions H80.5cm x 59.6cm x 558cm

If you have any technical questions about your Samsung appliance contact the Samsung support team for further assistance Telephone: 0330 7267864 or visit Samsung Support

Energy Monitoring

Your energy monitoring system will be provided using the Loxone Smart Plug, Energy Monitor and Miniserver..

Loxone Smart Plug 'Loxone Smart Socket Air' (to monitor the energy use of your smart appliance)

This plug records data from the appliance. It is not used to control the appliance.

Loxone Meter Reader 'IR Meter Reader Air' (to monitor the energy use of your home)

This device reads the electricity usage of the property, and sends it to the Miniserver Go. The device should have a USB cable plugged into a plug to give it power. In future, this may run off your electrical Distribution Board.

It is magnetic, and sticks on to the electricity meter. In some cases, the electricity meter in a home gives off encrypted signals, so we may have added an additional electricity meter which it can read.

The IR Meter Reader Air uses a very small amount of electricity, over a whole year it will use 10 pence of electricity

Loxone Mini Server (to transmit the monitoring data back to our secure data store)

The Miniserver Go is the heard of the Loxone system, it gathers all the data from the Smart Socket Air and IR Meter Reader Air. The information is then processed and sent to the mobile app that the householder uses

If you have any technical questions about your Loxone equipment contact the Loxone support team for further assistance.

Telephone: 01183 130 140 Email: support@loxone.co.uk

RASPBERRY PI 3B+

This is used to allow the Loxone system to communicate with the Bristol is Open (BIO) securely. The Raspberry Pi is a tiny computer, which the Loxone Miniserver is plugged into. The Raspberry Pi will communicate with the house internet via WiFi.

Appendix B - Smart Homes Info Pack

Smart Homes information pack

SMART HOMES

BRISTOL CITY COUNCIL

Smart Homes information pack

What's on offer?

Eligible households could receive a Samsung washing machine, dishwasher and/or tumble dryer **free of charge** to help save money on their energy bills.

What is a smart appliance?

A smart appliance can be linked to an app on your smartphone or tablet, allowing you to control its functions remotely. It can be programmed to work at certain times of the day but can also still be used as an ordinary appliance.

Alongside the appliances we would also provide *Loxone Smart Home* equipment that will help you better understand your energy use and provide our project team with information to help with the research we are undertaking. This will include an energy monitoring device and 'smart home' communication hub. You will be able to keep the appliance(s) at the end of the project (providing you remain in the project for its duration). We will also organise for any old appliances to be removed by the *Sofa Project*.

We will provide full technical support for the duration of the project to make sure we get you up and running again should you have any difficulties with the appliances. Relevant contact details will be provided in your welcome pack and any faulty items will be replaced as part of the product's warranty guarantee.

We would like to find out whether having appliances that can be automatically set to work at times when energy is cheaper can save people money, so we will be offering households the opportunity to get rewards for altering the time at which they use their appliances around six months after installation. The more you use energy outside of peak usage times the more rewards you could receive! This will not involve changing your energy supplier as we can monitor your energy usage remotely and reward you as you go.

Why are we offering this?

Smart Homes is a programme of activities and trials within the REPLICATE Project – a five-year initiative that is exploring how we could use new technologies to save energy and money, travel in more environmentally-friendly ways, and tackle local challenges.

We need to test out appliances in real households to see if our ideas for saving energy will work. If energy prices rise in the future it will be increasingly important for households to use energy at times when it's cheapest. Using learning from this project we hope we'll be able to use technology to help people automate this process – so you don't have to lift a finger to save money.

Who can take part?

You will need to live in the **Ashley**, **Easton** or **Lawrence Hill** areas of Bristol . Our project team will help to confirm if you live within the project area.

You do not need to own your property but you will need to have permission from your landlord to install the appliances.

We want to ensure that the whole of your home is as energy efficient as possible so we would like to talk to you about other energy saving measures we could install for free or at a subsided cost.

If you undertake an unplanned house move to another area within Bristol during the project (until January 2021) you'll just need to let us know so that we can set up your smart home equipment in the new property.

Unfortunately if you are moving away from Bristol then we will require the equipment back so another household in Bristol can benefit. We will then provide a like for like (or as close as possible) replacement for your old appliance.

There are some situations where you may not be able to participate:

- If you currently have an A+, A++ or A+++ rated appliance (we won't save you enough energy by switching these).
- If you do not have the permission of the homeowner to install the equipment.
- If you plan to move house within the project lifetime (until January 2021).
- If you do not currently have the appliance which you would like to upgrade e.g. if you don't already have a tumble dryer, we would not be able to install a new one.
- If your home is deemed electrically unsafe by our installers.
- If you do not have a standard dimension appliance (to fit 60cm wide x 90cm high space).

Our friendly team can help talk you through any eligibility questions you may have.

Eligible area

So what's the catch?

In order to participate you would need to:

- Complete an initial survey questionnaire (which may take up to 30 minutes).
- Be open to discuss further energy saving measures in your home such as improved insulation (which we will have funding available to support from mid 2018).
- Take care when handling the equipment to avoid damage (households will not be liable for any maintenance or repair costs of the system unless caused by wilful damage).
- Provide access to your home for installation of the equipment and any repairs (with pre- arranged appointments).
- Keep a constant power supply to the equipment during the trial.
- Donate your old appliances to a local charitable organisation (we will arrange collection).
- Be willing to change the times you use the appliances, for example switching evening washing loads to be automatically scheduled during the day. This might mean asking households to change their behaviour, but there will be rewards for doing so.
- Share your energy usage data (anonymously) to a data sharing platform so that others can learn from our experiences here in Bristol.
- Have permission of your landlord.
- Participate in the **research elements** of the project as detailed below.

Research element

If you chose to take part in the project you will agree to participate in the research element of Smart Homes. This will involve providing your contact details to our project partners and the *University of the West of England (UWE)* and agreeing to be contacted to provide information about your experience of the project (approximately a year after you have had your new appliance installed).

At the sign-up stage you will also be asked some basic information about your house and home energy usage.

Most of this information can be taken from your property's **Energy Performance Certificate** (EPC) that the project team can find online.

We may be able to offer in-home support to help you collect the information needed to participate if you're not sure where to find it. Please get in touch with the project team to discuss this.

We would encourage you to take and record regular meter readings, as UWE will be in contact a year after install to measure the impact that your Smart appliance has had on your electricity consumption and bills. You will also be required to provide UWE with copies of your energy bills from the previous two years.

With prior consent, some households may also be invited to take part in promotional activities.

Ending your participation

The REPLICATE Project finishes in January 2021. After this you can continue to enjoy your new appliance with no commitments to the project or its research.

You are free to leave the trial of smart appliances at any point but if you do so then we will need to take back the appliance so that it can be provided to another household.

In this instance we will provide you with a like for like replacement appliance based on the item you had before the trial.

Please note this will be a used item and may not be the exact model you previously had. In short you won't be losing out if you leave the trial.

Data Protection

Your personal information will be held and used in accordance with the General Data Protection Regulations 2018. Smart Homes will not disclose such information to any unauthorised person or party outside of the REPLICATE project partnership.

Contact details

If you have any further questions please contact our friendly project team:

Email: **replicate@warmupbristol.co.uk**
Phone: **0117 352 1180**

www.connectingbristol.org/replicate

If you would like this information in another language, Braille, audio tape, large print, easy English, BSL video or CD rom or plain text please contact us on 0117 352 1180

Funded by:

This project has received funding from the European Union's Horizon 2020 Research and Innovation programme under grant agreement No. 691735

In partnership with:

BD10140