CASSANDRA - A simulation-based, decision-support tool for energy market stakeholders

Konstantinos N. Vavliakis
Dept. of Electrical and
Computer Engineering
Aristotle University of
Thessaloniki,
Informatics and Telematics
Institute, CERTH
Thessaloniki, Greece
kvavliak@issel.ee.auth.gr

Anthony C. Chrysopoulos
Dept. of Electrical and
Computer Engineering
Aristotle University of
Thessaloniki,
Informatics and Telematics
Institute, CERTH
Thessaloniki, Greece
achryso@issel.ee.auth.gr

Kyriakos C.
Chatzidimitriou
Dept. of Electrical and
Computer Engineering
Aristotle University of
Thessaloniki,
Informatics and Telematics
Institute, CERTH
Thessaloniki, Greece
kyrcha@issel.ee.auth.gr

Andreas L. Symeonidis
Dept. of Electrical and
Computer Engineering
Aristotle University of
Thessaloniki,
Informatics and Telematics
Institute, CERTH
Thessaloniki, Greece
asymeon@eng.auth.gr

Pericles A. Mitkas
Dept. of Electrical and
Computer Engineering
Aristotle University of
Thessaloniki,
Informatics and Telematics
Institute, CERTH
Thessaloniki, Greece
mitkas@auth.gr

ABSTRACT

Energy gives personal comfort to people, and is essential for the generation of commercial and societal wealth. Nevertheless, energy production and consumption place considerable pressures on the environment, such as the emission of greenhouse gases and air pollutants. They contribute to climate change, damage natural ecosystems and the man-made environment, and cause adverse effects to human health. Lately, novel market schemes emerge, such as the formation and operation of customer coalitions aiming to improve their market power through the pursuit of common benefits.

In this paper we present CASSANDRA, an open source[1], expandable software platform for modelling the demand side of power systems, focusing on small scale consumers. The structural elements of the platform are a) the electrical installations (i.e. households, commercial stores, small industries etc.), b) the respective appliances installed, and c) the electrical consumption-related activities of the people residing in the installations.

CASSANDRA serves as a tool for simulation of real demand-side environments providing decision support for energy market stakeholders. The ultimate goal of the CASSANDRA simulation functionality is the identification of good practices that lead to energy efficiency, clustering electric energy consumers according to their consumption patterns, and the studying consumer change behaviour when presented with various demand response programs.

---

[1]https://github.com/cassandra-project

Categories and Subject Descriptors

I.6.7 [**SIMULATION AND MODELING**]: Simulation Support Systems

General Terms

Design, Experimentation

Keywords

Decision support tool, Consumer analysis, Demand-response

1. INTRODUCTION

Energy markets have undergone fundamental changes at the conceptual level over the last years. Sustainability requirements have transformed the traditional power production scheme to a distributed energy resource one. The future points towards a great number of decentralized, small-scale production sites that include renewable energy sources. The Smart Grid paradigm is here to stay, while the deregulation of energy markets has defined new rules and procedures. This context presents substantial opportunities to actors throughout the electrical energy value chain. These, however often remain unexploited for two main reasons: a) Small scale consumers, comprising the vast majority of electrical energy consumers, each have insignificant market power and are individually unable to affect market related policies and

b) there is currently a lack of tools for modelling the consumers with respect to their consumption behaviour, the identification of trends formulated and the evaluation of policies and decisions against specific socio-economic and environmental performance indicators.

To this end, we designed and developed CASSANDRA, a platform for the realistic modelling of the energy market stakeholders, including small-scale consumers. Through CASSANDRA, users are able to test and benchmark working scenarios that can affect the electrical power system operations, pricing, as well as environmental policies at different levels of abstraction, starting from a single small-scale installation (e.g., an apartment) and shifting up to large numbers of consumers (e.g., an entire city). Furthermore, CASSANDRA supports the grouping of small scale consumers into Consumer Social Networks (CSNs) with increased market power that are able to negotiate their energy needs, but can also respond to specific targeted incentives for behavioural change.

The main objectives of CASSANDRA are to:

1. Provide a theoretical model that clearly defines the energy profiles and energy consumption patterns of all relevant stakeholders, as well as their interrelations, in order to support energy efficiency.

2. Promote and validate methodologies that provide increased market power to low level consumers, through consumer coalitions (Consumer Social Networks - CSNs).

3. Develop a modular, extendable and interoperable software platform that implements the CASSANDRA theoretical paradigm to model, simulate and benchmark scenarios and operations, and act as a decision making tool for energy performance, considering consumer behaviour.

4. Promote the proposed consumer behaviour and load modelling approach, and disseminate the energy and environmental benefit/impact.

The paper is organized as follows: In Section 2 we present related work and motivation, while in Section 3 we introduce the CASSANDRA platform, where we discuss the CASSANDRA entities, the architectural overview, and the respective workflow, as well as the CASANDRA simulation process. Section 4 presents the validation and testing of the Platform. We conclude in Section 5, summarizing the benefits CASSANDRA contributions to the energy market and proposing future work directions.

2. RELATED WORK

During the past 30 years, a number of paradigms have been presented in the international scientific community regarding energy efficiency frameworks and modelling of energy consumption of small-scale electricity consumers. These models correspond mainly to residential electricity consumers; however, in most cases, they can also be used for the modelling of commercial or small industry consumers.

Moreover, during the last years, social networks have been in the center of attention for various research communities. Numerous methodologies, such as social network analysis, graph analysis, statistical models, multilabel classifiers and collaborative systems have been proposed, that either combine information from multiple social resources or analyse specific aspects of the formed networks. These methodologies are based on resource attributes and have been applied in different fields, but their use in energy markets remains limited, despite the fact that they could have an important influence for their operation.

2.1 Energy Efficiency Frameworks

The approaches used for energy-consumer modelling may be divided in two categories: a) top-down and b) bottom-up approaches. Top-down approaches require the implementation of econometric and technological data in a macroscopic level. Bottom-up ones correspond to the cognitive approaches (see [21] for a comprehensive review of existing implementations), where methods such as neural networks [12], fuzzy logic [8], conditional demand analysis and regression techniques [3] are used along with past consumption data, in order to predict future load profiles. The problem, in the latter case, is that cognitive systems are heavily dependent on the nature of the past data used for their training procedure. This data usually corresponds to a specific background (e.g. the energy consumption in a single country), and always refers to a specific time frame, thus not being able to follow changes in consumption patterns, due to technological or policy novelties.

With respect to existing frameworks that either optimize or simulate energy performance, two types of approaches exist: a) frameworks that focus on optimally solving very specialized problems (e.g. specifying the market equilibrium in the power stock market), where algorithmic optimization is performed; and b) frameworks that model specific segments of the Energy market, either in the context of efficiency (buildings and houses), or in the context of affecting behaviours (reduction of consumption).

It should be mentioned, with respect to the latter category of frameworks, that no standardization attempt exists, that would allow for the modelling of all entities and their interactions in a uniform manner. As far as tools that focus on optimally solving specialized problems are concerned, employed approaches are further differentiated in two subcategories: (i) tools that focus on the optimization of energy performance, given a set of hard constraints (not allowing versatility/dynamicity and exploration); and (ii) tools that simulate more dynamic scenarios and are investigating for the optimal strategy/solution. In the former case machine learning techniques, such as neural networks [11], genetic algorithms [9], support vector machines and temporal data mining techniques have been employed for solving problems related to load prediction, market equilibrium identification and decision support. In the latter case, agent modeling and machine learning approaches are adopted, in order to ensure dynamicity and uncertainty. One should mention SEPIA (Simulator for Electrical Power Industry Agents), a multi-agent platform capable of running a plethora of computing experiments for many different energy market scenarios. In the same context, EMCAS (Electricity Market Complex Adaptive System) [7] implements an efficient mechanism for handling the Electric Energy Market, while also studying the complex interactions between the physical entities of the market, in order to analyze the participants and their strategies. As far as real-market analysis is concerned, Bagnall [2] presents a simplified simulation model of Great Britain's Electricity Market with producers being agents that participate in a series of auctions-games, wherein the possibil-

ity of cooperation between agents is also observed. There are also a few implementations specifically designed for Energy Markets, which are, though, merely simulations for the regulations. One of the first platforms created was MA-SCEM (Multi-Agent Simulator for Competitive Electricity Markets) [22], which was used to validate regulations and behaviors within the Electricity Markets, using only naive reinforcement learning strategies. Another Agent platform was soon developed in order to test the reliability of the proposals FERC of USA (Federal Energy Regulatory Commission) applied in the Standard Market Design (SDM) [14]. In this implementation, agents were used to model the Producers (Enterprises or Persons), the Consumers and the Distribution Network Operator. More multi-agent systems have been successfully applied to model Power Markets (see [27] by Zhou et. al. for a survey).

In 2006, Sonnenschein et al. [20] described an agent-based simulation tool for examining the impact of real-time pricing methods on the power consumption of domestic users (households). The main entities within the model are electricity suppliers and electricity consumers scheduling their demand according to real-time prices for electricity. In [26], a hybrid econometric and social influence model was implemented for evaluating the influence of pricing and public education policies on residential habit of electricity using in power resources management. Abras et al. [1] implemented a Multi-Agent Home Automation System (MAHAS) dedicated to power management that adapts power consumption to available power resources according to inhabitant comfort and cost criteria. Karnouskos et al. [13] presented the emergent concept of a future Smart City Grid. They designed and built a simulator based on software agents that attempts to capture the behaviour of a smart city, by simulating discrete heterogeneous devices that consume and/or produce energy, able to act autonomously and collaborate. Similarly, in [25], [18], residential consumers are modelled as autonomous agents (SmartHome) which are responsible for shifting their consumption in order to maximize the residential revenue, in addition to several social welfare factors. These studies do not attempt to model consumer behaviour, since their goal is the decentralized control of shiftable loads and the optimal utilization of renewable energy sources.

On the other hand, Power TAC is a multivariate benchmarking framework that models the energy market from a financial/economic perspective [6]. In Power TAC, focus is given on individual agents that act as retail brokers in a local power distribution region, purchasing power from a wholesale market, as well as from local sources, such as homes and businesses with solar panels, and selling power to local customers and into the wholesale market. However, there are issues unaddressed in Power TAC. One of these issues is consumer modelling, which is indispensable for the development of dynamic consumer models, as preferences essentially guide choices and decisions, both in real life and in a dynamic model that endeavours to simulate it [4].

In bottom-up approaches heuristic methods can be used. Various sets of rules are used, in order to develop a number of basic load profiles. Subsequently, probability distributions are utilized, in order to produce greater numbers of load profiles, based on the initial ones. The set of rules used for the determination of the initial load profiles may be arbitrarily complex, and it may take into consideration demographical data, historic consumption data, information regarding all the involved consumption appliances, seasonal variations, etc. [10, 17].

However, even the most advanced heuristic models available in the literature, share some common weaknesses. They are developed within the traditional structure of the respective energy markets, considering consumers mainly as passive elements with little or no response capabilities. This means that the resulting models are very efficient in producing great numbers of small-scale consumer load profiles, but they are not able to model consumer reaction to possible incentives. A novel approach is, therefore, essential that will be able to combine the results of the aforementioned approaches with the ability to model demand response.

2.2 Consumer Social Networks

Consumer social networks have been leveraged for motivating people into reducing CO2 emissions as well. In [16] Mankoff proposed to explore the use of social networking websites in supporting individual reduction in personal energy consumption. They integrated feedback on ecological footprint data into existing social networking and Internet portal sites, which allowed frequent feedback on performance, while enabling the exploration of motivational schemes that leverage group membership. Approaches like the ones followed by Opower[2] and Bidgely[3], attempt to empower Energy Service Companies (ESCOs) with actionable insights for their customers by engaging the latter in a social loop with tips on their energy efficiency and comparative graphs on others' consumption. Consumers became the center of attention in the early 2012, when the "Which? campaigning charity"[4] started the "Big Switch" initiative, a completely new way to buy - and save money on - people's energy. So, using the power of thousands of consumers, it planned to negotiate with energy suppliers in the UK and seek to secure a market-leading energy deal and help people make the switch. Which? had more than 50.000 people sign up to the "Big Switch" within the first 48 hours and more than 250.000 more since.

In the same context [24] attempts to organize consumers into Virtual Electricity Consumer (VEC) entities, which procure electricity as a single customer in order to get a discount on electricity through collective buying. In their approach, they try to exploit real social relations between consumers and build flattened profiles for the VECs. Regarding demand side management, [15] envisages the formation of cooperatives of medium-large consumers and the design of a mechanism for allowing cooperatives to regularly participate in the existing electricity markets by providing electricity demand reduction services to the Grid. The proposed Consumer Demand Side Management (CDSM) mechanism employs agents that proactively place bids in the electricity market, contribute to the flattening of the energy consumption curve for the day ahead and distributes profit among the cooperating agents. [23] focuses on the concept of a Virtual Power Plant (VPP) and attempts to define a mechanism for creating coalitions between wind generators and electric vehicles, where wind generators seek to use Electric Vehicles (EVs) as a storage medium to overcome the vagaries of generation.

---

[2]www.opower.com
[3]www.bidgely.com
[4]www.which.co.uk

However, available literature on coupling demand side management with social networks is limited, and mainly focused in the field of telecommunication networks. This is mainly based on the fact that the burst of cell phone use took place during the 00's, while the existence of different consumption programs/tariffs is still limited in the energy market. Following the social network literature, [19] described the structure of a CSN with two key characteristics: a) strength of a tie, and b) relational density of a personal network. Shi's study provided many useful guidelines for the design of optimal price plans. In another work, [5] estimated the importance of tariff-mediated network effects in the use of mobile telephones and the impact of the structure of social networks on consumer choice.

From all related work, one may say that work by Vinyals et al. [24] is closer to our approach. However, their hypothesis differs from our approach in three ways. First, authors consider the existence of social interaction between consumers, thus links already exist; we define links based on the "proximity" of consumers, and model consumers accordingly. Second, authors do not consider network topology, and how consumers are assigned under different ML (medium-low) voltage transformers; this is probably not important if applied in large-scale (since authors are referring to the market), nevertheless has to be taken into consideration from the power network perspective. What is more important, though, is that we do not focus on the behavioural aspects of the network and how they will act/interact, given specific incentives; we focus on how structural properties of a generated network will influence peak reduction.

2.3 CASSANDRA's contributions

From all the above, it is obvious that work on energy efficiency, i.e. energy efficiency frameworks and energy-consumer modelling, is not mature enough, and that there are not many published studies that approach residential electrical power consumers as individual and autonomous entities. In addition, consumer social networks have been partially leveraged for motivating people into reducing energy consumption. Nevertheless, none of the already developed tools and methodologies provide an integrated solution to energy market stakeholders for testing and benchmarking scenarios where they provide incentives to consumers for energy reduction.

This is the reason we propose CASSANDRA, a platform for the realistic modelling of the energy market stakeholders, including small-scale consumers. The system integrates behavioural and simulation models with user interfaces to support the following functionalities: a) disaggregation of loads at an installation level, to identify the appliance consumption models and to map active and reactive load curves to appliance switching by consumers, b) modelling of consumer activities, supporting also small-scale consumers, c) a range of simulation models, including simulation of large number of installations using demographics/statistics, d) modelling of consumer behaviour/response to monetary incentives, environmental impact, as well as to different levels of awareness of the consumers to these parameters, e) simulation and comparison of different demand response and feedback scenarios, f) computation of different Key Performance Indicators (KPIs) and their modification in the simulated and applied scenarios, g) public pages for consumer awareness where consumers can view the simulated and real consumption of their installations.

The CASSANDRA platform is expected to be used by aggregators, energy efficiency/monitoring devices providers, regulators/distribution network operators/transmission system operators, energy traders/suppliers, energy efficiency software providers, energy and environmental consultants, university/research centres, and large scale consumers.

3. THE CASSANDRA PLATFORM

3.1 The CASSANDRA Entities

Modelling in the CASSANDRA platform is based on a set of assumptions, used to determine the properties of entities, whose interactions result in the representation of a single small-scale installation with respect to its electrical energy consumption. Conceptually, an installation is assumed to comprise a set of electrical appliances, used by a single equivalent person according to his/her consuming habits. Thus, there is an installation entity, which is correlated with a person entity and a number of appliance entities. The appliances are described by respective consumption models, used to simulate their energy consumption during operation. Moreover, the person is correlated with a number of activities (e.g. cooking, cleaning, entertainment), involving the operation of appliances. Each activity is essentially a set of probability density functions, determining the possible triggering, time of use, and daily usage frequency regarding an appliance within a simulation scenario. The CASSANDRA platform functionality is based on a set of models, which interact in order to represent and simulate real life demand side scenarios, ranging from measurement analysis regarding individual installations, to demand response modelling for a consumer community.

3.2 Architectural overview

The proposed CASSANDRA platform modules are depicted in Figure 1; these are:

1. The Agent module: The agent module plays the role of the simulation engine of the CASSANDRA platform. It contains the necessary structures for modelling the behaviour of agents (persons) and object constructs (i.e. installations, appliances) with respect to energy consumption activities, communication, and response to stimuli in the form of demand side management programmes.

2. the CSN module: The CSN module is a tool for performing general network analyses and clustering the population into groups based on (dis)similarities of behaviour or social activities. It is a benchmark tool for experts to experiment and test different use case scenarios, in order to fine-tune hypotheses on consumer response based on CSNs.

3. the Aggregation module: The aggregation module serves as the interface between the agent module (simulation engine), the CSN module (graph/clustering studies) and the web-service module (API). Even though it is depicted as a box in Figure 1, it is actually the wrapper code that implements the interactions between the modules. The aggregate module is a form of mapping I/O of one module into I/O of another module.

4. the Web service module: The web service module is the socket of the CASSANDRA platform to the outside world. Through this interface it is possible to define simulation scenarios, policies (pricing, legislative) and KPIs in order to be simulated by the agent module. The same interface can be used to query for results and for post-simulation reporting and analysis.

5. the End-user module: The end-user module is an implementation of a front end-client that adheres to the web-service module communication protocol. It provides a graphical interface for using the platform. With this module the CASSANDRA platform is a complete and usable tool and not just a simulation software platform. In addition, the end-user module has ready-made reporting and assessment functionalities. We should also mention that due to the incorporation of the web-services it is possible for any developer to implement an end-user module.
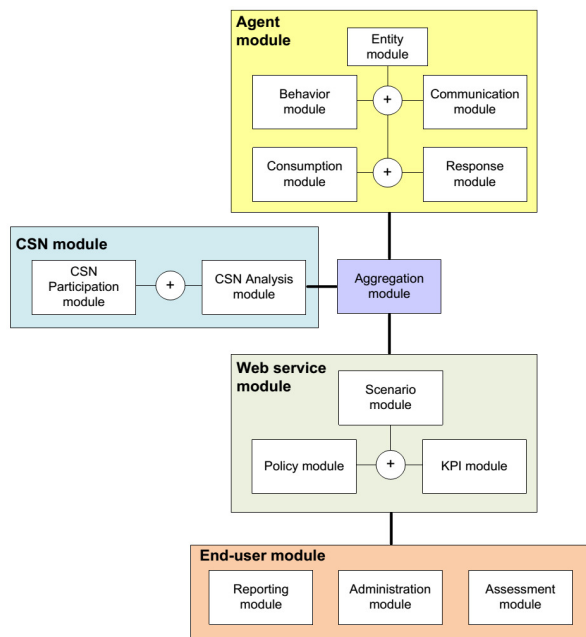
Figure 1: CASSANDRA modules

The overall CASSANDRA platform architecture as projected from the conceptual CASSANDRA platform (Figure 1) is depicted in Figure 2. Besides the aforementioned CASSANDRA modules, several technical layers have been introduced from the conceptual to the implementation design.

The CASSANDRA platform was designed to follow the Software-as-a-Service paradigm, where the software along with its data are hosted on the cloud and access to their computational and data resources will be met, on-demand, by users via thin clients. The decoupling of the CASSANDRA platform and the end-user interface introduced a layer of communication between the server and the client implemented through a RESTful web-service.

Last but not least, a database component was added in order to efficiently store and retrieve data related to the entities of the CASSANDRA platform. There are several types
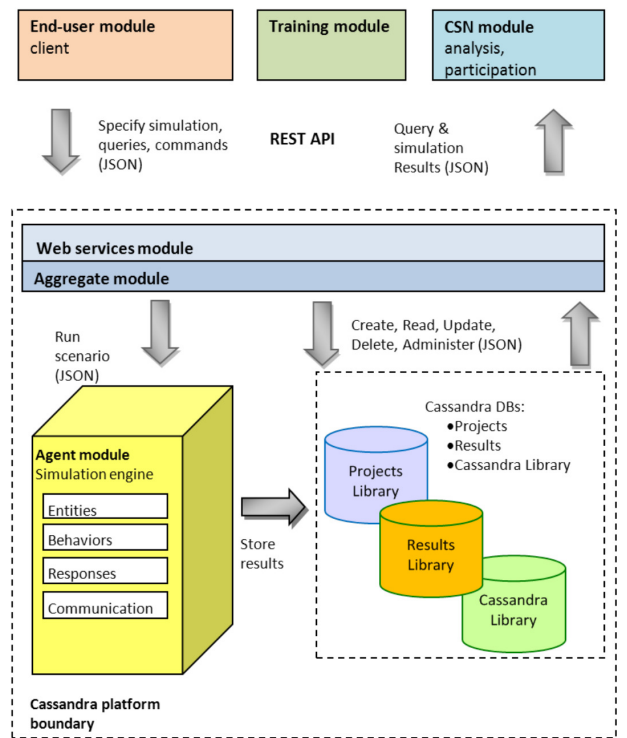

Figure 2: Overall CASSANDRA platform technical architecture

of databases in the platform: (a) the projects database, which holds all the modelling information to run the scenarios defined by the user, (b) the results database, which holds information regarding runs of a scenario along with their definition and (c) the CASSANDRA and user libraries databases, which hold ready-to-use models of CASSANDRA entities so they can be used and re-used in different scenarios. Figure 3 depicts the CASSANDRA platform as a UML deployment diagram

3.3 CASSANDRA Workflow

In order to illustrate the main dataflow and workflow scenario, let's consider an end-user that wants to perform a simulation of the energy demand of a large neighbourhood and find information about the activities that took place during peak demand time (the workflow is depicted in standard fonts and the dataflow in italics):

1. Through the CASSANDRA GUI, the user creates a new project and then a new scenario. The properties of the project and scenario entities are entered in the GUI and stored in the DB in the form of JSON documents.

2. The user defines a number of appliances along with their consumption models and a number of person types along with their behaviour with respect to appliance usage patterns. Also the user defines demographic data of appliances, person types and installation types, number of days and start/end dates of the simulation along with the number of installations to be populated dynamically. The entity definitions are stored in the DB in the form of JSON documents.
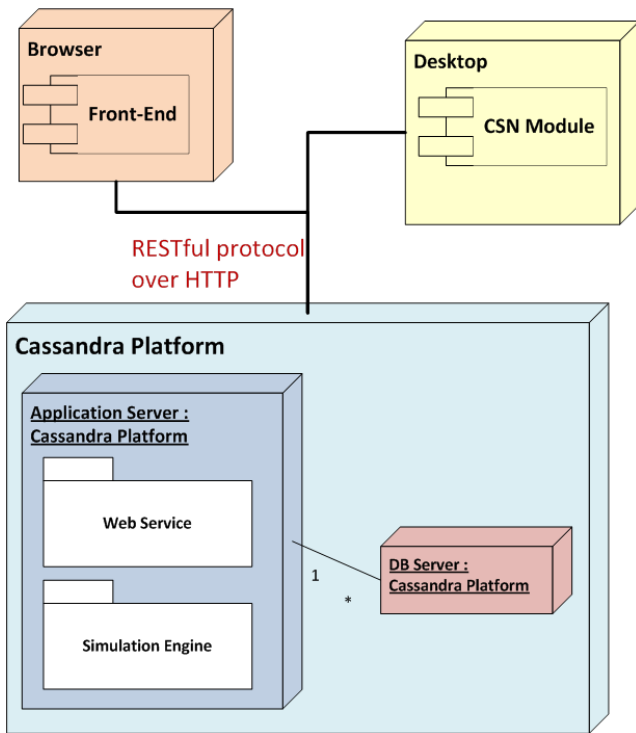
Figure 3: Deployment diagram of the CASSANDRA platform

All the operations are made through the REST Web-service Create Read Update Delete method calls.

3. The user wants to take a person of type "student" from the CASSANDRA library and put it in the collection of person types to be created during the instantiation of the scenario. The JSON documents and their hierarchical structure of an entity of type person (person =¿ activity =¿ activity models =¿ distributions) are retrieved from the CASSANDRA Library and the documents are copied to the projects database for further use.

4. The user requests to execute the simulation scenario with the provided simulation parameters A JSON document is created that contains all the necessary JSON documents (as sub-documents) required to run the simulation. A new database is created in order to store the JSON document defining the scenario and later store the results of the run. The simulation is executed as a background thread.

5. During the simulation, the CASSANDRA entities apply their behavioural patterns to their corresponding environment and every minute the aggregate and per installation power consumption is stored in the database.

6. Upon completion, the user requests the energy consumption time-series of the neighborhood for a specific period of time, say one week. A JSON document is returned containing the power demand for every minute in the simulation.

7. The user spots a peak and requests to find out what types of activities and appliances where on, during

that peak time. A JSON document is returned displaying the types of activities and appliances at that period of time along with their count and percentage.

A more general dataflow of creating and running CASSANDRA scenarios is depicted in Figure 4 as an activity diagram:
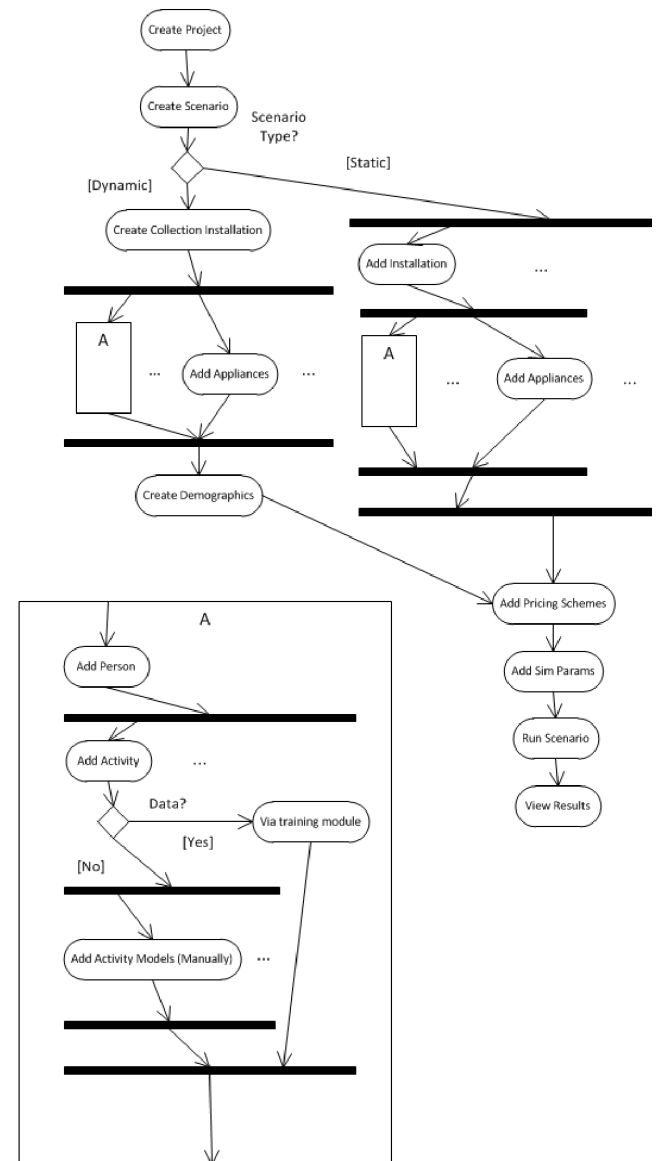

Figure 4: Activity diagram for creating a CASSANDRA platform scenario

3.4 CASSANDRA Simulation

The CASSANDRA platform is a combination of event-based and time-based simulation methods. Algorithm 1 describes the whole simulation procedure.

First of all because of its stochastic nature, each simulation can be executed numerous times in order to get an average estimate of its results and provide more robust conclusions. For each such run, every tick is considered to be

Algorithm 1 Simulation procedure

```
initialize_run();
while(tick < endTick)
 if(beginningOfDay(tick))
  foreach installation
   queue.add(installation.getDailySchedule());
 // search if an event exists and apply it by
 // either switching on or off an appliance
 while(queue.peek() &&
   queue.peek().getTick() == tick)
  event = queue.poll();
  event.apply();
 // For each installation calculate power,
 // energy, costs, kpis
 foreach installation
  installation.update(tick);
```

one minute of time. If the current tick is the beginning of the day, the activity models in each installation create the events of switching on or off appliances and are then inserted into a priority queue. If an event or events exist that are supposed to take place at the current tick, we pop them out of the queue and apply their action. Then, we update the consumption profile and KPIs in the current tick for each installation.

In the end the simulation database contains information about: 1) Active power per minute, 2) Reactive power per minute, 3) Energy consumed, 4) Costs incurred and 5) KPIs (e.g. max active power, average active power, CO2 emissions), for each installation and for the run as a whole.

3.5 CSN Module

CASSANDRA upgrades the role and market power of small-scale electricity consumers by utilizing the concept of consumer social networks (CSN). Within this context, CASSANDRA provides both the essential infrastructure for the development and transparent operation of CSN, as well as the theoretical background concerning the utilization of CSN for the improvement of network reliability and quality of service. The CASSANDRA CSN module focuses on the types of CSN graphs to be identified and developed (real-time interaction social graphs, macro-scale interaction graphs), as well as on the parameters that comprise the proximity and analysis metrics of the graphs (e.g. load consumption instance, peak consumption, consumer type, being some of them). Different metrics and clustering algorithms are supported. The main focus of this module is to model and cluster consumers with respect to mutually exploitable loading levels and interaction with other consumers, rather than merely classifying them on their consumption habits. Consumer profiles provide the basis for a number of operations (personalization, experts finding, etc.). In the CASSANDRA CSN module (Figure 5), we represent consumer networks as an undirected graph: $G = (V, E)$, where each node is denoted as a vertex $v \in V$. Two nodes may be connected or related via common characteristics. This connection is represented as an edge $e \in E$, with $E \subset VV$. We also define a matrix $C = \{c_{ij}\}$, where $c_{ij} = 1$ if and only if nodes c_i and c_j share the characteristic we defined edges to represent. A node/vertex $v \in V$ in the CASSANDRA graph represents an installation, thus a person (since every

installation contains one "equivalent" person). The CASSANDRA CSN module may cluster a set of installations using various clustering methods, such as k-means, hierarchical or graph clustering.
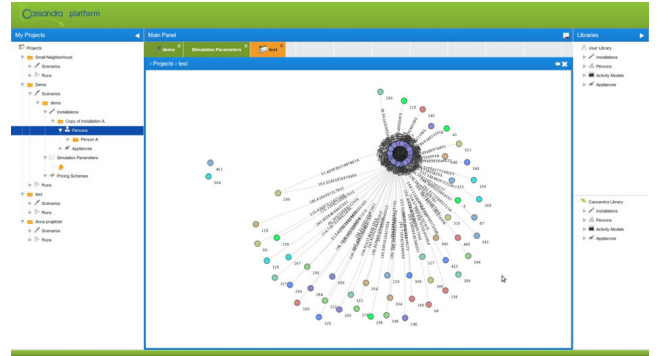

Figure 5: The CASSANDRA CSN Interface

The module can load a set of entities from the CASSANDRA Agent module, define connections among entities based on type, consumption, behaviour, etc., display the resulted network, and cluster them, providing the necessary information about the attributes of each entity/cluster, as well as their consumption. The goal is to group similar (according to selected criteria) consumers and treat them accordingly. For example, if we group consumers based on their consumption regarding cooking, a respective targeted demand response (D/R) program may improve results towards a general Demand Management (D/M) program. We may also achieve improved results when we cluster consumers based on their profile type by identifying groups susceptible to environmental, financial or social incentives.

4. VALIDATION

The first step in the platform evaluation process is the evaluation of the CASSANDRA baseline simulations. In addition, comparison of the consumption distribution throughout the day (using the Jensen-Shannon divergence of power consumption distributions) can be used to compare the real and simulated consumption measurements. The development and verification follows a four-step process:

1. Collection of measurement from real households.

2. CASSANDRA baseline model: The CASSANDRA platform baseline models were created using existing library-based models, through manually generated models and through the disaggregation and training modules.

3. Simulate sites: The models are used to simulate baseline consumption for all three pilot sites

4. Verify simulation against measurement: The CASSANDRA simulated baselines are compared to the pilot baselines and checked for accuracy.

For validation, we used an estimation of the aggregate consumption of a household for a specific day (Expected Power) to compare it with the output of the simulation for one day using several Monte Carlo runs. Figure 6 and Figure

7 display the two results. Figure 6 depicts the expected active power for a group of installations by solving the simulation analytically, while Figure 7 depicts the output produced from the simulation which matches the one from Figure 6. It is obvious that the to results are really close.
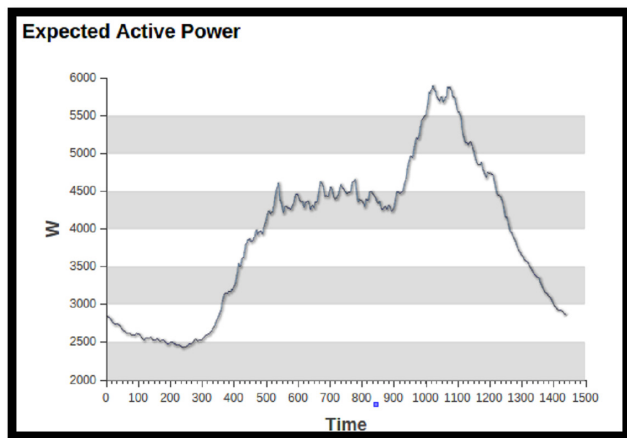

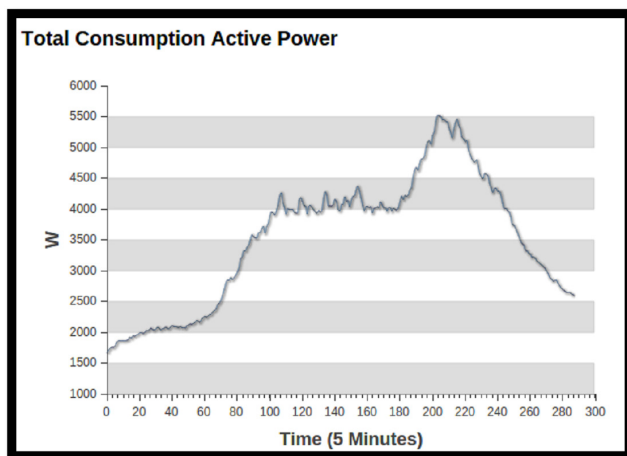Figure 6: Expected active power for a group of installations by solving the simulation analytically

Figure 7: Output produced from the simulation which matches the one from Figure 6

From a different perspective, we have detailed information available for a period of one week. In Figure 8, the measured consumed energy per activity is presented for the same household and time period, whereas in Figure 9 the respective consumed energy per activity is exhibited, as recognized by the CASSANDRA platform. The two results are almost identical.

4.1 CSN Evaluation

In order to test on the usefulness of the CSN approach, we performed the following experiment: we ran a demand-side management scenario, as a means to quantify the outcome of the CSN formation. In this scenario, each consumer agent that has accepted to join the network (his/her willingness to connect) is given an incentive relative to the preferences of
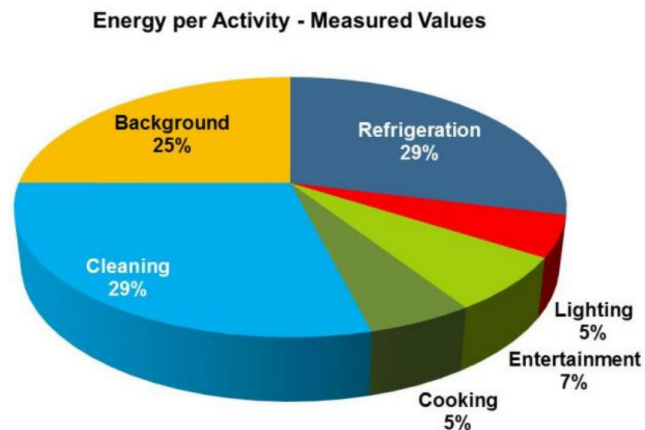

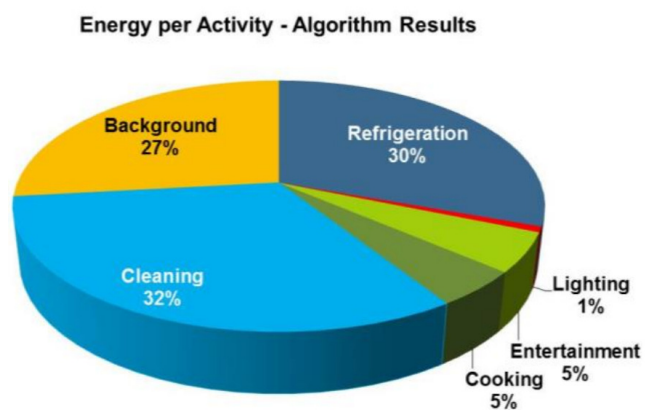Figure 8: Measured energy per activity for a household within a week

Figure 9: Energy per activity for a household within a week, as recognized in the CASSANDRA platform

the group he/she has been assigned, in order to shift his/her load curve by a quarter of the hour earlier/later. Each consumer agent, based on the actions executed by neighbouring agents (links), executes the complementary action so that diversity in actions within the clusters is maintained. Agents that do not participate in the CSN perform as planned initially. Merely by applying this simple incentive scheme, we observe a 6% reduction (Figure 10) in the peak of the load curve (from 2824 MWs to 2652 MWs) plus a reduction in the standard deviation of the power consumption in the 96 quarters of the day (from 573MWs to 549 MWs). This could lead to better forecasting of power demand due to less variance in the resulting curve. In fact, if consumer agents participation in the network increases (from 50% to 75%), peak reduction can reach 14%. Thus, appropriate target incentives to respective consumer groups could obviously lead to numerous benefits related to the energy market and the power system.

5. CONCLUSIONS

In conclusion, CASSANDRA achieved progress in: (i) defining and evaluating theoretical tools that are used for bottom-up modelling of small-scale consumers, including the consumer behaviour modelling methodology, the appliance
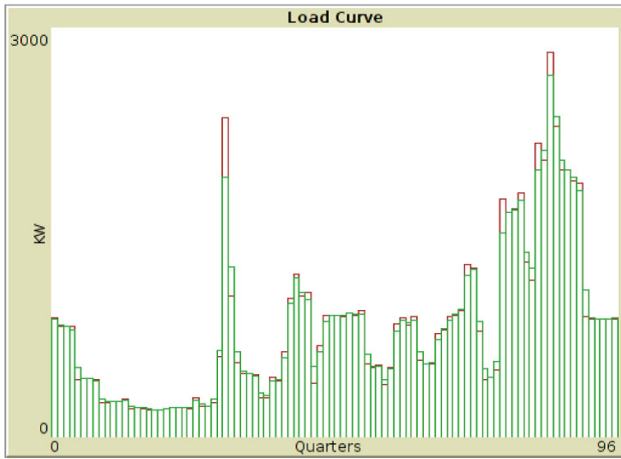
Figure 10: Peak reduction in the load curve. With red bars is the consumption before incentivising the consumer agents and with green after.

consumption models and load disaggregation methodology, (ii) developing the CASSANDRA software platform and (iii) the individual CASSANDRA system components, i.e. the CASSANDRA server, the CSN module or individual software packages/classes.

In addition, CASSANDRA has defined a set of guidelines of proper platform use for different stakeholders based on the lessons learned through the project. CASSANDRA followed an iterative software development process, consisting of four phases: Collection of user requirements, definition of system specifications (including implementation of software tests), and software implementation/integration and software evaluation. This resulted to the development of a data driven, bottom-up demand-side modelling framework that describes consumption at the appliance level and scales up to entire regions of the power system. Combined with models of consumer behaviour and response, CASSANDRA allows the mapping of the observed load curve into consumer activities and the simulation of demand-response scenarios that take into account consume habits and behaviours and their changes in response to monetary or environmental incentives.

CASSANDRA realization concerns the energy markets and offers the necessary tools to develop the representative models for the various involved stakeholders. For this reason, an expandable, easy-to-use architecture was built, in order to model the versatility and dynamic nature of all involved entities in the energy market. CASSANDRA outcomes can support increased participation in the energy market, by providing novel business opportunities for existing stakeholders, and offering incentives to the formerly inactive customers to guarantee their involvement.

CASSANDRA may provide a new ground for aggregators and ESCOs, as it offers them the prospect of initiating and operating CSNs, acting as mediators between them and the TSO. The CASSANDRA platform is expected to boost R&D activities concerning the respective software and hardware. The transparent architecture of CASSANDRA can foster the active participation of small software houses, increasing their efficiency and competitiveness. For SMEs,

CASSANDRA can offer business opportunities to enter the electric power generation market: with a rising number of power generating units, there is a rising demand for maintenance and overhaul. Through planning and installing power generating units, there are several business challenges for service providers to profit.

Consumers have no means of envisioning themselves as part of a larger group. Single households and businesses do not know of the impact they could have joining with their neighbours, or of the benefit that this could bring to them. CASSANDRA can help consumer groups to model and quantify their own options, potential impact and potential benefits. In addition, as of now, most utilities are in the difficult position of having to guesstimate what benefits a particular program could bring to their customers. It makes energy efficiency or a demand response program more difficult to "sell" internally within the utility and externally to consumer groups, regulators, policy makers and the press. CASSANDRA can act as an aid here, providing data and insight.

Future work for CASSANDRA includes further validation of the platform, creating a more user-friendly interface and installation procedure, as well as exploring more activity models based on personalized data.

6. ACKNOWLEDGEMENT

7. REFERENCES

[1] S. Abras, S. Pesty, S. Ploix, and M. Jacomino. An anticipation mechanism for power management in a smart home using multi-agent systems. In *Information and Communication Technologies: From Theory to Applications, 2008. ICTTA 2008. 3rd International Conference on*, pages 1–6, April 2008.

[2] A. Bagnall and G. Smith. Game playing with autonomous adaptive agents in a simplified economic model of the uk market in electricity generation. In *Power System Technology, 2000. Proceedings. PowerCon 2000. International Conference on*, volume 2, pages 891–896 vol.2, 2000.

[3] J. Bentzen and T. Engsted. A revival of the autoregressive distributed lag model in estimating energy demand relationships. *Energy*, 26(1):45 – 55, 2001.

[4] M. Bichler, A. Gupta, and W. Ketter. Research commentarydesigning smart markets. *Information Systems Research*, 21(4):688–699, 2010.

[5] D. Birke and G. P. Swann. Network effects, network structure and consumer interaction in mobile telecommunications in europe and asia. *Journal of Economic Behavior & Organization*, 76(2):153 – 167, 2010.

[6] C. Block, J. Collins, S. Gottwalt, W. Ketter, and C. Weinhardt. Modeling household energy consumption under fixed and variable pricing. In *Proceedings of the Twentieth Workshop on Information Systems and Technology (WITS-10)*, pages 1–6, St. Louis, Missouri, USA, December 2010.

[7] G. Conzelmann, G. Boyd, V. Koritarov, and T. Veselka. Multi-agent power market simulation using emcas. In *Power Engineering Society General Meeting, 2005. IEEE*, pages 2829–2834 Vol. 3, June 2005.

[8] D. Gerbec, S. Gasperic, I. Smon, and F. Gubina. Determining the load profiles of consumers based on fuzzy logic and probability neural networks. *Generation, Transmission and Distribution, IEE Proceedings-*, 151(3):395–400, May 2004.

[9] K. jae Kim and I. Han. Genetic algorithms approach to feature discretization in artificial neural networks for the prediction of stock price index. *Expert Systems with Applications*, 19(2):125 – 132, 2000.

[10] J. Jardini, C. M. Tahan, M. Gouvea, S. U. Ahn, and F. Figueiredo. Daily load profiles for residential, commercial and industrial low voltage consumers. *Power Delivery, IEEE Transactions on*, 15(1):375–380, Jan 2000.

[11] K. Kamijo and T. Tanigawa. Stock price pattern recognition-a recurrent neural network approach. In *Neural Networks, 1990., 1990 IJCNN International Joint Conference on*, pages 215–221 vol.1, June 1990.

[12] S. Karatasou, M. Santamouris, and V. Geros. Modeling and predicting building's energy use with artificial neural networks: Methods and results. *Energy and Buildings*, 38(8):949 – 958, 2006.

[13] S. Karnouskos and T. de Holanda. Simulation of a smart grid city with software agents. In *Computer Modeling and Simulation, 2009. EMS '09. Third UKSim European Symposium on*, pages 424–429, Nov 2009.

[14] D. Koesrindartoto, J. Sun, and L. Tesfatsion. An agent-based computational laboratory for testing the economic reliability of wholesale power market designs. In *Power Engineering Society General Meeting, 2005. IEEE*, pages 2818–2823 Vol. 3, June 2005.

[15] R. Kota, G. Chalkiadakis, V. Robu, A. Rogers, and N. R. Jennings. Cooperatives for demand side management. In *The Seventh Conference on Prestigious Applications of Intelligent Systems (PAIS @ ECAI)*, pages 969–974, August 2012.

[16] J. Mankoff, D. Matthews, S. R. Fussell, and M. Johnson. Leveraging social networks to motivate individuals to reduce their ecological footprints. In *Proceedings of the 40th Annual Hawaii International Conference on System Sciences*, HICSS '07, pages 87–, Washington, DC, USA, 2007. IEEE Computer Society.

[17] J. V. Paatero and P. D. Lund. A model for generating household electricity load profiles. *International Journal of Energy Research*, 30(5):273–290, 2006.

[18] S. D. Ramchurn, P. Vytelingum, A. Rogers, and N. Jennings. Agent-based control for decentralised demand side management in the smart grid. In *The 10th International Conference on Autonomous Agents and Multiagent Systems - Volume 1*, AAMAS '11, pages 5–12, Richland, SC, 2011. International Foundation for Autonomous Agents and Multiagent Systems.

[19] M. Shi. Social network-based discriminatory pricing strategy. *Marketing Letters*, 14(4):239–256, 2003.

[20] M. Sonnenschein, M. Stadler, B. Rapp, J. Bremer, and S. Brunhorn. A modelling and simulation environment for real-time pricing scenarios in energy markets. In A. S. E. Klaus Tochtermann, editor, *Shaker Verlag*, pages 153–161. Shaker Verlag, 2006. ISBN 978-3-8322-5321-9.

[21] L. G. Swan and V. I. Ugursal. Modeling of end-use energy consumption in the residential sector: A review of modeling techniques. *Renewable and Sustainable Energy Reviews*, 13(8):1819 – 1835, 2009.

[22] Z. Vale, T. Pinto, I. Praca, and H. Morais. Mascem: Electricity markets simulation with strategic agents. *Intelligent Systems, IEEE*, 26(2):9–17, March 2011.

[23] M. Vasirani, R. Kota, R. Cavalcante, S. Ossowski, and N. Jennings. An agent-based approach to virtual power plants of wind power generators and electric vehicles. *Smart Grid, IEEE Transactions on*, 4(3):1314–1322, Sept 2013.

[24] M. Vinyals, F. Bistaffa, A. Farinelli, and A. Rogers. Coalitional energy purchasing in the smart grid. In *Energy Conference and Exhibition (ENERGYCON), 2012 IEEE International*, pages 848–853, Sept 2012.

[25] P. Vytelingum, T. D. Voice, S. D. Ramchurn, A. Rogers, and N. R. Jennings. Agent-based micro-storage management for the smart grid. In *Proceedings of the 9th International Conference on Autonomous Agents and Multiagent Systems: Volume 1 - Volume 1*, AAMAS '10, pages 39–46, Richland, SC, 2010. International Foundation for Autonomous Agents and Multiagent Systems.

[26] M. Xu, Z. Hu, J. Wu, and Y. Zhou. A hybrid society model for simulating residential electricity consumption. *International Journal of Electrical Power & Energy Systems*, 30(10):569 – 574, 2008.

[27] Z. Zhou, W. Chan, and J. Chow. Agent-based simulation of electricity markets: a survey of tools. *Artificial Intelligence Review*, 28(4):305–342, 2007.