Smart Home Energy Management

João Sabino
joao.g.sabino@tecnico.ulisboa.pt

Instituto Superior Técnico, Lisboa, Portugal

July 2022

Abstract

In the last century, energy consumption has increased tenfold. Global calls for environmental sustainability along with rising energy costs push companies and individuals to manage energy more efficiently. Demand Response (DR) programs were developed to promote individual and collective contributions to the efficiency of the grid. This dissertation reviews DR programs implemented by energy distribution companies, consumption habits from residential consumers and the incentive dynamic in shared residential complexes. Existing Home Energy Management Systems (HEMS) are analysed regarding approach, energy savings and flexibility. A HEMS focused on peak distribution is proposed. The system is able to schedule the execution of appliances as they are requested, deciding with present information. Scheduling decisions are based on energy consumption and user preferences in device priority and schedulability. Moreover, battery and PV systems are integrated, and different instances of the server application are able to cooperate towards load distribution. Peak reductions of up to 45% versus unmanaged households were achieved in testing simulations, with additional benefits to demand balance in overnight periods.

Keywords: Smart Grid, Energy Management, Demand Response, Appliance Scheduling

1. Introduction

In 2015, the United Nations General Assembly set 17 Sustainable Development Goals, to be achieved globally until 2030. In particular, SDG 7 aims to "ensure access to affordable, reliable, sustainable and modern energy for all". Among other targets, nations are called to double the global rate of improvement in energy efficiency.

Developed countries have complex electrical networks, able to provide electricity to residents. These electrical grids must maintain enough firm capacity to match demand at all times. However, demand is variable: high-demand periods require higher capacity, often met with expensive peaking plants or a higher base capacity. The fluctuation of demand over time - i.e., the demand curve - forces energy companies to maintain a surplus of energy supply, create consumption forecast models, and use additional infrastructure to sustain the grid [7].

Many regions have aging grid infrastructure that doesn't support bidirectional electrical flow [13]. New homes are being built with solar panels and Battery Storage Systems (BSS) to improve energy efficiency and utilize more renewable energy, yet, in a traditional grid, the energy produced locally must be used for self-consumption. Overproduction can't be returned to the network, leading to penalties for end-users and waste of energy. Traditional grids are being modernized with bidirectional flow of energy and communication signals, incorporating sensors and smart meters [3].

From data on household consumption published by the Eurostat in 2019, space heating represents 63.9% of energy end-uses, followed by 14.8% consumed in water heating, 14.1% in lighting and appliances and 6.1% in cooking. Slicing the household consumption by source fuel, electricity represents 24.7% of the energy mix. Odysee shows the share of electricity is increasing rapidly, coming from 21% in 2000 [9]. When accounting exclusively for electricity consumption, lighting and appliances represent most end-uses, as most of the heating in the European Union uses gas and renewables.

The aforementioned factors - drive for energy efficiency, management of self-production, and relevance of appliances in household consumption - expose the need for tools that optimize energy consumption and coordinate demand from a bottom-up approach, compatible with traditional and modern grids. As such, the objective of this work is to develop a Home Energy Management System (HEMS) able to improve demand inefficiencies, both within residential households and local grids, with minimal impact to the consumer. The system is a server application that can control the execution time of electrical appliances to perform at a more adequate period, preventing switchboard trips and local shortages, meant to be installed in a computer within the household. The system is interactive, receiving execution requests as they come and distributing them using present-time information. In shared residential spaces, such as an apartment building or condominium, users are incentivized to distribute their energy loads not only to benefit from a lower contracted power, but also to reduce shared electricity bills, as the electrical installation requires less power.

The solution should account for most types of appliances, BEV, and include support for alternative energy sources such as Photovoltaic Systems and Battery Storage Systems. User preferences must be considered in order to reduce disruption. Energy loads require accurate parameterization for the system to be as effective as possible. The system should be able to simulate the residential environment, including characterization of devices, programs and executions. To complement the program, a simple user interface is to be developed along with test consumption scenarios. Results are produced from these scenarios and evaluated using consumption indicators such as peak, peak-to-average ratio and total delay time, when compared against an unmanaged system.

2. Background

Section 2 reviews data and accomplishments from previous works that is relevant to the solution developed.

2.1. Household consumption

Domestic consumption data is crucial for identifying patterns in demand, as well as paths to optimize distribution. Space heating is used to maintain indoor temperatures at acceptable levels for human thermal comfort. Naturally, more energy is required to do so when outside temperature is lower. Geographic location affects solar irradiance and plays a major role, both in heating necessities and broader lifestyle choices that end up shaping consumption. Countries from southern Europe and the Mediterranean sea, where the climate is more moderate, such as Malta, Portugal and Cyprus, not only consume less energy in space heating, but also rank among the countries with lower overall household energy consumption [9]. Other factors include the number of residents in the household, seasonality, intraday temperature fluctuation and periods of absence.

Concerning the intraday variation on energy consumption, temperature is only one of the many factors that influence expenditure. A significant portion of home owners leave for work in the morning and come back at late afternoon. Cooking is prevailing before launch and dinner time. Lighting is more necessary at night. Additionally, workers commonly use their free time at early night for leisure activities or house chores. These behavior patterns are associated with the use of specific appliances: meal preparation may require the oven, toaster, stove or microwave, house chores include the iron, dishwasher and washing machine. Most of these are large appliances that consume a significant amount of energy. Intraday variation, when plotted in a graph, is called the demand curve.

2.2. Appliance information

Appliances come with information regarding power ratings, energy efficiency and capacity (if any). Rated power is the power needed for execution, even if the consumption is not constant; it represents the maximum power draw of the device. As such, it can be used as an upper bound to represent the consumption of a machine. Some appliances have different consumption profiles, depending on the task. Even when turned off, many appliances enter a ”standby” mode, reducing consumption to a minimum to power digital clocks, smart hub connection, etc. Different modes often have a different rated power.

An interesting categorization from an energy management standpoint is according to the load profile. Resistance-based appliances use resistors to turn electricity into heat. These appliances use constant power: the consumption curve of these devices is steady until interrupted by a timer or the user. As such, the rated power is an adequate estimate of real consumption. Motor-based appliances such as fans also fit this category.

Some appliances have simple ON/OFF self-regulation, like the fridge and some heaters. Self-regulating heaters, despite also being resistance-based, have a thermostat to stop producing heat until temperature falls again below a threshold. The fridge alternates between ON and OFF on a time basis, with additional power used when the door is opened or during defrost cycles. Consumption of these devices can be estimated adequately using equation models that parameterize internal and external factors [2].

Other appliances, such as the washing machine, dishwasher and clothes washing, have more complex self-regulating mechanisms, with the power consumption varying according to the operating cycle [12]. A full cycle of the dishwasher includes washing, rinsing and drying. Stages of self-regulating appliances are hard to identify without a power meter, and vary according to brand and model. Thus, accurate modelling of the consumption curve of these appliances requires specific data from the machine, not accessible to the average consumer.

2.3. Peak demand reduction terminology

The following processes shape the demand curve with focus on reduction of peaks. They are used and often combined in energy management solutions, or encouraged by energy suppliers through variable pricing.

- Load Distribution is a strategy focused on evenly distributing energy loads across time, without affecting the overall energy consumed. The amplitude of the demand curve is reduced, as valleys are filled with loads from peaks. This strategy can be applied by shifting energy loads from peak to off-peak periods. In some literature, it is decomposed into peak shifting and valley filling.

- Peak Shaving is the elimination of consumption peaks whenever they would be formed. It encompasses load denying - rejecting new loads after a certain consumption threshold - and the use of secondary energy sources during peak demand, such as solar panels or batteries. The consumption peaks become plateaus with lower maximum value.

- Load Reduction consists in reducing the overall consumption. The shape of the demand curve remains similar to the original, but with lower values. Utilizing more energy-efficient appliances falls into this category.

2.4. Supply-side techniques

Energy suppliers already implement a range of techniques to shape the demand curve of consumers, both in traditional and smart grids. The most basic control of demand comes from the maximum available capacity chosen by the consumer - the contracted power. Consumers cannot exceed this value; if the instant consumption exceeds the contracted power, the switchboard will trip. The capacity is measured in kVA, and a higher contracted power results in higher energy cost per kWh. Consumers are encouraged to contract power above their average consumption, in order to remain below the threshold during periods of higher activity.

Through financial incentives, contracts may also motivate consumers to adapt their consumption behavior to a more favorable pattern for energy distribution. This is called Demand Response (DR). Consumers can enroll in DR programs and actively contribute to the stability to the grid, through rate models such as Time-of-Use, Real-Time Pricing and Critical Peak Pricing. Other models based on rebates or coupons also reward consumers for reducing consumption during peak hours.

2.5. Active energy management

Home Energy Management Systems provide additional control to consumers, offering remote control over home appliances and automation to varying degrees. Smart plugs are used to control the state of non-smart appliances, where as smart machines can communicate directly with a network. Communication is standardized through protocols such as ZigBee and MQTT [5, 4]. Devices communicate with a coordinator to inform their status or receive new instructions. Depending on the appliance and the level of control over its functionality, the status can range from a mere ON/OFF switch to more fine-grained information such as temperature, intensity and water usage. A web or mobile application can communicate with the network to display the statuses and provide control to the user.

2.6. Scheduling solutions

At their core, the HEMS mentioned above are tools to reduce energy waste. In order to effectively optimize the demand profile of a house, solutions require more control over consumption of appliances. A system that can decide or recommend the time of execution for each appliance is called a scheduling system.

Many academic solutions ([8], [1], [10]) attempt to formulate consumption as a mathematical problem and develop optimization algorithms. These systems schedule the use of appliances in advance, with varying degrees of device categorization. Due to the scheduling happening in advance, these solutions face difficulties adjusting to impulsive or seasonal changes in consumer behavior, as the preferences of the user are disregarded. Additionally, contracted power is a considerable factor in electricity prices regardless of the rate model, and lacks representation in the aforementioned research.

Group optimization has been further explored in the context of shared apartment complexes or building districts. In [11], a HEMS able to manage multiple buildings with heterogeneous consumption, compatible with thermal energy storage and PV systems is developed. Two architecture models - cooperative and coordinated - are compared in performance against uncoordinated rule-based controllers, using Key Performance Indicators (KPI) such as the peak, peak-to-average ratio and self-sufficiency. The results showed reductions of the consumption peak by 9.7% and the peak-to-average ratio by 7.7% in a coordinated approach - a centralized model similar to the cooperative scheduling approach described in this paper.

3. Implementation

The proposed solution is an autonomous HEMS that schedules household appliances with a focus on peak shifting. Peak shifting is accomplished by limiting consumption to a threshold defined by the user. If the threshold is surpassed, energy loads are delayed or shifted to an available period. The system is able to start, shift, delay and interrupt appliances in real time, using a priority system based on user preference to decide. Scheduling interactively using dynamic priorities increases flexibility to the user, enabling sudden consumption changes from unexpected behavior and quickly adapting to them. BSS and PV systems are integrated as peak shaving and valley filling mechanisms. The solution boasts a minimal interface and is able to coordinate loads between residential units in shared dwelling complexes.

3.1. Stakeholders

There are three stakeholders that theoretically stand to gain with such system: residential consumers (or tenants), property managers and energy providers. Consumers benefit from reducing their contracted power and overall consumption. Blackouts are a significant inconvenience for residents, especially in the night. Energy providers benefit from increased stability in demand, reducing the dependency on energy from expensive peaking power plants.

Property managers are responsible for maintaining the shared private infrastructure in condominiums or apartment buildings. The entity may just exist implicitly, if there is infrastructure with expenses split between residents. Depending on the services that are offered, energy costs may represent a significant part of management expenses: the Association of Condominium Owners of Ontario claims that typical condominium corporations spend 35–50% of their annual operating budget on utilities [6]. Managing the consumption within shared spaces to flatten the demand curve may result in a cheaper electricity plan.

3.2. Modes of operation

The scheduling system that is proposed can operate in **single-house** and **multi-house** modes. In single-house mode, loads are distributed locally with focus on reducing delay inconveniences, preventing energy shortages and respecting the consumption limit imposed by the user. This mode directly benefits consumers. Energy companies benefit indirectly from a flatter curve in multiple households, although there is no coordination between instances. The property manager can run its own unit to manage machines in shared spaces, behaving as an individual consumer and reaping the same benefits.

In multi-house mode, consumers provide their aggregate scheduled consumption to a central unit that aggregates demand and provides scheduling recommendations to individual units, based on the loads already scheduled. Large appliances and EV charging are scheduled with awareness of peaks forming in the complex grid, getting distributed more evenly across time. For privacy reasons, aggregate data is stored pseudo-anonymously and temporarily. Moreover, no individual device data is shared outside the local unit. Even so, the central unit could be accessed maliciously and expose consumption of individuals, by gathering enough consumption data to match it with the house. To the consumer, the multi-house mode adds a privacy risk. Executions may also be further delayed to support the grid, although the delay can never surpass the limit set by the user.

However, the benefits to the grid should be much more significant. Evidently, the energy provider is directly benefited by the system, as the aggregate demand curve is softened, and thus less reliant on energy from peaking power plants. In countries where shared energy contracts exist, property manager and tenants gain from the system together, with an overall lower energy expense due to a reduced contracted power. In shared European residences, the additional immediate reward in multi-house mode is the prevention of energy outages due to excess consumption at building/condominium level. Nonetheless, the increased stability provided to the grid has a positive environmental impact to all stakeholders, which could be rewarded by the energy provider or state programs, in a similar manner to existing financial incentives in DR programs.

3.3. System Entities

The system is organized in three main components: the **Coordinator** and **Aggregator** Django applications, and a **Processor** module containing the scheduling logic utilized by the other components. The general purpose language chosen for implementation is Python, using the Django web framework and a database in SQLite.

The Coordinator fully manages the production and consumption of an individual household when the system is disconnected from an Aggregator - operating in single-house mode. All appliances, alternative energy systems, executions and user parameters are stored in this module. The Coordinator receives and handles requests from appliances to start or stop activity. Moreover, it is able to connect to an Aggregator, provide the necessary data for the latter to store an aggregated demand curve and eventually request recommendations for execution times. If BSS exists in the household, the Coordinator always manages the charge and discharge cycles locally, based on individual household consumption.

The Aggregator receives connections from Coordinators that decide to operate in multi-house

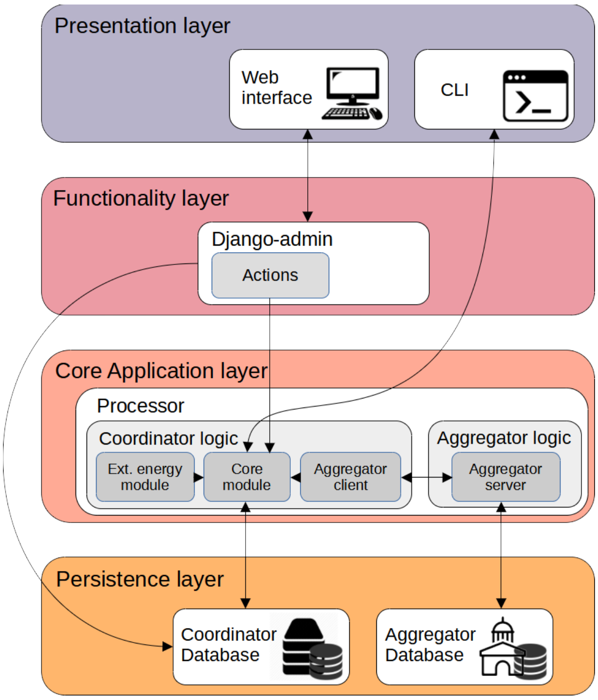

Figure 1: Layered architecture diagram.

mode. It stores pseudo-anonymous consumption data in bulk and responds to scheduling recommendation requests. The Aggregator is to be installed outside of any housing unit, but represents the shared interests of the residential complex; as such, it can be maintained by the property manager. Data regarding individual appliances is not disclosed to the Aggregator. Instead, the Aggregator is only informed of the time and power needed for a new schedule.

When the user intends to activate an appliance, either through an hypothetical interface in the machine or through the Coordinator app, the Coordinator sends the request to the core of the Processor. The Processor contains functions to schedule new requests, activate battery charges and discharges, and an asynchronous thread to update the state of appliances in the background.

3.4. Priority function

Under heavy loads, executions are started or delayed according to a monotonic priority function. The function is used to attribute priority values to executions dynamically. A higher priority will lead to a quicker start, if no energy is available to start immediately without stopping other appliances, and supplemental energy sources are absent or unusable. Part of the priority formula for an execution is based on the priority class defined for the appliance. Users can set a class between **urgent**, **normal** and **low-priority** for each different mode of operation of a machine:

- Urgent requests correspond to time-sensitive appliances, typically small appliances that require human intervention to be used, such as the hair dryer, the coffee machine and the cake mixer. These devices often operate in short spans. This priority class can interrupt appliances with lower priority, to be resumed shortly after.

- Normal priority is given to appliances that provide benefits to the user during operation (e.g., the television) or require some form of active monitoring (e.g., the oven). These devices are still time-sensitive, but can start within a certain delay, without requiring immediate attention.

- Low priority is given to machines with little to no time restrictions, typically large appliances which just need to finish operation before a longer deadline. These devices do not require supervision and generally can operate overnight, such as the washing machine, the dishwasher, and the BEV charger. Low-priority devices are prime candidates for a load distribution scheduling strategy. In single-house mode, they will be allocated to periods with the lowest local consumption, while in multi-house mode, the Aggregator decides on the period with lowest aggregate consumption.

The other factor for the priority formula is the time since request, compared against the maximum acceptable delay defined by the user. As the remaining acceptable delay approaches zero, the priority of the execution increases at a faster rate, potentially gaining higher priority than other executions in upper priority classes. This prevents new requests from further delaying an execution that is already delayed to the limit of acceptableness. The formula is as follows:

$$delay_{\text{request}} = \begin{cases} t_{\text{current}} - t_{\text{request}} & \text{if unscheduled,} \\ t_{\text{start}} - t_{\text{request}} & \text{otherwise.} \end{cases} \tag{1a}$$

$$delay_{\text{remaining}} = delay_{\text{maximum}} - delay_{\text{request}} \tag{1b}$$

$$f = priority_{\text{base}} + floor(\frac{60 \times \alpha}{delay_{\text{remaining}} + 60}) \tag{1c}$$

$$priority = \begin{cases} f & \text{for } f \leq 10, \\ 10 & \text{for } f > 10. \end{cases} \tag{1d}$$

Where $delay_{\text{remaining}}$ is the value, in minutes, of the remaining acceptable delay, $priority_{\text{base}}$ is one of the values in table 1 and α is a constant used to modulate the steepness of the priority curve. By flooring the fraction in equation 1c, priorities are

Priority class	Base value
Urgent	7
Normal	4
Low-priority	1

Table 1: Base priority values.

discretized between 1 and 10. Discretization allows for a more predictable and coherent behavior. The value of $delay_{\text{request}}$ is dynamic if the execution is not yet scheduled, as the current time keeps changing. However, it becomes static once a start time is attributed. Therefore, appliances do not change priority constantly and do not interrupt each other consecutively when trying to regain execution.

When a new request arrives, if there is not enough power available below the threshold using other procedures (elaborated in subsection 3.5), the scheduler will compare the priority value of the new request with the static values of the executions already scheduled. Delayed executions get their priority recalculated when rescheduling.

The values 6, 8 and 10 were experimentally tested for α. With a value of 8, the curves adjusted better to round numbers that the user might input more often as maximum delay, such as twenty minutes and one hour, while keeping an adequate differentiation of priority values for the three categories as the acceptable delay approaches zero. For urgent energy loads, any delay of less than one hundred minutes gets maximum priority (10). Urgent executions generally have short acceptable delay periods, so a new urgent request will typically get maximum priority and start immediately, even without enough energy available, as long as interruptible executions are found. Normal executions have a wider range of priority values, starting from 4. This is the default priority. Appliances with normal priority and longer acceptable delay may not start immediately under heavy load, but reach the higher echelons at one hour remaining, and hit maximum priority at twenty minutes. Finally, low-priority devices are expected to never hit maximum priority. The priority class is intended for devices that will run whenever energy consumption is lower, filling valleys and possibly activating overnight. Figure 2 shows the time intervals at which different priority values are reached, according to the priority class.

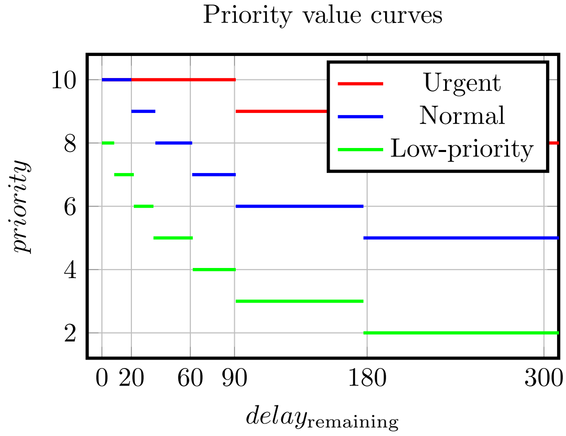

Figure 2: Priority values by class, base priority and acceptable delay remaining, in minutes.

3.5. Scheduling process

When a new scheduling request arrives, the scheduler will preliminarily select three possible start times, according to the operations needed to find enough energy available. There are three possible scheduling procedures:

1. Simple scheduling: there is enough available power to schedule at the suggested time (including PV generation).
2. Battery-enabled scheduling: the BSS can provide complementary energy to enable the execution.
3. Priority scheduling: shiftable loads with lower priority can be interrupted to provide enough power for execution.

The start time obtained in each of these scenarios, as well as the procedure used to satisfy the request, are chosen based on the operation mode of the system and priority class of the appliance. In households without a BSS, the scheduling decisions are simplified to deciding between procedure 1 and 3. Conversely, if the household has a storage system, the start time obtained from procedure 3 will include available energy discharges, acting as a composite of procedures 2 and 3.

For urgent and normal executions, the strategy is focused on **peak shifting**. The goal is to schedule as soon as possible while respecting the threshold. As such, the three scenarios are evaluated to find the simplest procedure that provides the soonest execution time. The suggested time for each scheduling procedure is the soonest where the conditions apply. In low consumption periods, all three will propose the current time: the execution starts immediately, without requiring a battery discharge or shifting appliances. During high consumption, the procedures may find different start times or even fail to find a viable execution period. If requesting a battery discharge anticipates the execution time over the simple scheduling option, or procedure 2 provides a viable execution time where as procedure 1 did not return a start time, the battery will be activated and provide complementary power. If shifting appliances provides an additional advantage in

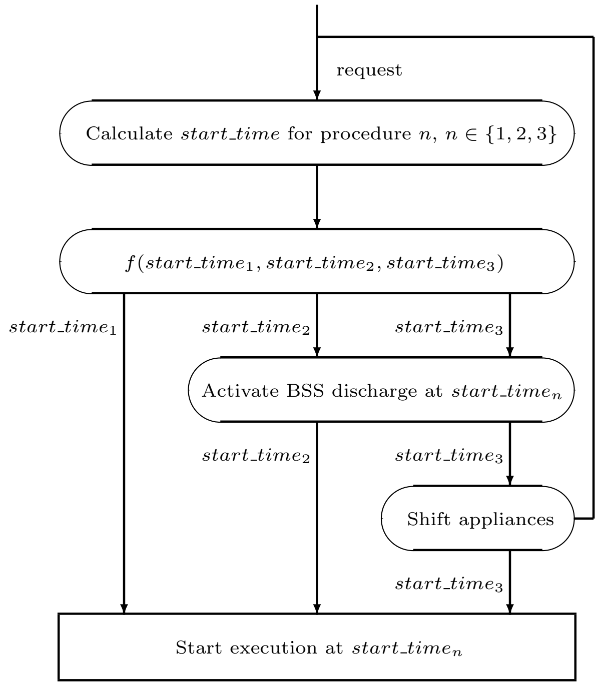

Figure 3: Algorithm for scheduling appliances.

reducing the execution delay over procedure 2, or procedure 2 also failed to find an execution time, the battery will provide as much power as possible and the remaining power needed will be met by shifting appliances.

For low-priority executions, the scheduler aims to schedule to the period of lowest consumption, using local or aggregate values - local if operating in single-house mode, aggregate if in multi-house mode. These strategies are called, respectively, **local load distribution** and **aggregated load distribution**. With both strategies, each procedure returns the start time for the period with the most energy available, if any is found. However, where as in local load distribution this period is found by merely sorting the available periods by available power, in aggregate load distribution all locally available periods are sent to the Aggregator, which returns the start time with less aggregated load. For these requests, delay time is not the decisive factor: instead, the scheduler will choose the first procedure that returns an available execution time. As such, procedures 2 and 3 are only used if there is not enough power available along the entire period of acceptable delay.

PV energy generation is always treated as complementary energy to the grid, effectively raising the threshold during solar hours. Figure 3 represents the scheduling process. For normal and urgent requests, function f returns the minimum value, akin to a $min()$ function. For low-priority requests, f returns the first existing value.

Battery charges and discharges are treated as executions, although with special rules. BSS executions are created by the system when necessary, scheduled using a different logic. Discharges cannot be interrupted and, as they do not require energy from the grid, are never delayed. Charges, on the other hand, are low priority executions, meant for low demand or solar production periods. Charges can only be interrupted if there are no scheduled discharges depending on the energy that would be charged. Interrupted charges are not rescheduled. Instead, new charges are scheduled when the battery is depleted, or periodically every day.

After a regular appliance is scheduled, the system checks if a consumption peak was formed. If consumption is above 70% of the user-defined threshold and the BSS has enough power, a battery discharge with the power difference between consumption and half of the threshold is scheduled for the corresponding period. Discharges can be stacked, still enabling scheduling procedure 2, as long as the battery can deliver enough continuous power. However, the system cannot charge and discharge simultaneously, nor discharge below a maximum depth-of-discharge (depending on the model).

When a discharge finishes, the system checks if the battery is close to its maximum depth-of-discharge. If it is, a battery charge is scheduled up to full power. Additionally, every midnight, the system checks the battery status and schedules a recharge if it is not at full power. Battery charges are broken in executions with duration up to an hour, to improve flexibility and match the hourly solar production data. Scheduling a battery charge also follows one of two procedures:

1. Solar charge: charge during underutilized solar hours to prevent production waste.

2. Low-demand charge: charge during low consumption periods.

If the household has a PV system, the Coordinator will attempt to schedule based on the first procedure. If there is not enough solar production to fully charge the battery, the system will test if the battery could be fully charged using a combination of solar production and grid energy, still during solar hours. This procedure will only fail if solar production if already fully utilized, or there are not enough solar hours to satisfy the request. The second procedure will attempt to recharge the battery whenever consumption is below 30% of the user-defined threshold, creating executions with the power difference between 30% of the threshold and scheduled consumption. The decision of charging without solar production is justified by the usefulness of the BSS as a load balancing and peak shaving mechanism.

3.6. Execution life cycle

An execution is created with information regarding the appliance and energy profile it serves, as well as a timestamp of the request time. By default, an execution is created as part of the scheduling request handling process, immediately followed by the attribution of a start and finish time. In the default scenario, the request time is the moment the execution is created. However, executions can be created with request times in the future, if the user does not want to schedule immediately but rather at a later time. The execution is in the Pending state until it starts.

Executions have four states: **Pending**, **Started**, **Interrupted** and **Finished**. Once the execution starts, the scheduler updates the data object and prompts the appliance to start running. Energy profiles have a maximum execution time, which limits the duration of the execution after it is started. The execution may end in one of three circumstances: the maximum execution time is reached, the user manually terminates the execution, or the execution is stopped by the scheduler to enable higher priority executions. The first two scenarios change the execution state to Finished, while the latter signals the execution as Interrupted. Both state changes stop the appliance.

If the execution is terminated manually, there may be opportunities to anticipate Pending executions: the scheduler will automatically attempt to reschedule Pending executions if a nearer start time is found, starting with the highest priority ones and down to the lowest. If the execution is Interrupted, the scheduler will create a new execution with the same parameters to follow it, and retake the Started state at a later time.

4. Results

Measuring the performance of the proposed solution requires selecting representative consumption data for different household profiles. However, research data on household consumption with daily granularity and segmentation by appliance is incredibly scarce. As such, it is necessary to create demand scenarios based on total consumption data and patterns identified in section 2. Three types of consumption profiles were created to represent households with different energy dependencies, habits and sets of appliances:

- Household 1 represents a small residence with low overall dependency on electricity. Gas is used for water and space heating in addition to the cooker, leaving the dishwasher, oven, microwave, toaster and coffee machine as the only high consumption electrical kitchen appliances. A single resident produces consumption. Additionally, the household has a vacuum cleaner, a small washing machine, a fridge and a hair dryer. The user is absent from the home during the morning and afternoon.
- Household 2 is home to two residents. In addition to the household appliances present in household 1, the second residence includes an air conditioner, a water heater and an induction cooker. Users have separate morning routines and meal preparations using different appliances.
- Household 3 is a larger house, with a diverse set of electrical appliances, belonging to a family of four elements. At least one family member is always in the house. There is a gaming computer, a cake mixer, three televisions and a soundbar system. Moreover, the family has a PHEV that charges in the garage with energy from the household installation.

These scenarios are to be simulated within the system using fabricated daily routines to generate the demand curve over a day. To represent an unmanaged system that will constitute the baseline for comparison, the consumption threshold is set to infinite and all appliances have the same priority class and value. Results are presented for household 1 in single-house mode, and for all residences operating in multi-house mode.

The chosen Key Performance Indicators for evaluating the system are Peak, Peak-to-Average ratio (PAR) and Average Delay-to-Acceptable-Wait ratio (ADAWR). Peak is the maximum instantaneous net consumption reached, in watts. Net consumption corresponds to the power drawn from the grid, obtained by subtracting production from solar panels and BSS discharges from the total consumption. A lower Peak can promote savings to the consumer by reduction of the contracted power. Peak-to-Average ratio is a measurement of load balance, indicating how far the peak was from the average net power consumption. It is obtained by dividing the Peak by the average value. A value of 1 indicates perfect load distribution, although values between 2 and 5 are considered acceptable. ADAWR computes the average delay verified between request and start times of all executions divided by the maximum acceptable delay set for each appliance. It is a value between 0 and 1. Values closer to zero are desired.

4.1. Single-house mode - Household 1

In household 1, the user wakes up at 07:45 AM, takes a shower, prepares breakfast and leaves to work. Comes back at noon for a quick lunch, then leaves again until 05:30 PM. The evening is spent watching television, cleaning the house and preparing dinner. The dishwasher and washing machine are activated at night, before going to bed.

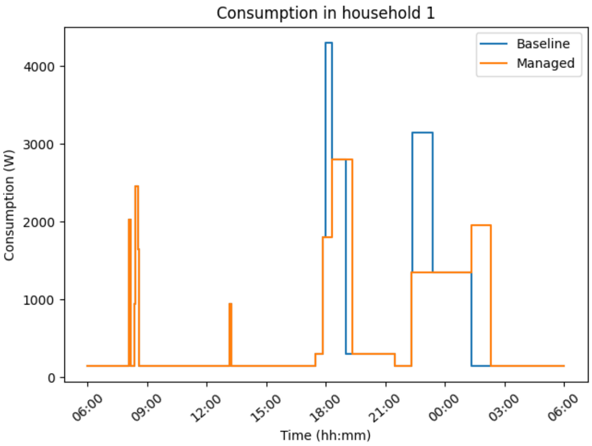

Figure 4: Comparison between unmanaged and system-managed consumption in household 1.

Figure 4 shows the demand curve for the two simulations. In the unmanaged scenario, Peaks are formed in the evening and night. To perform this routine, the user needs an energy contract with a contracted power value of 4.6 kVA. In the managed simulation, the scheduler delays two executions. The oven is delayed by 20 minutes, starting at 06:20 PM, after the vacuum cleaner finishes. This appliance is defined in the system as a normal-priority appliance with an acceptable delay of 40 minutes: as such, the delay proposed by the system comes within the acceptable range. The second delay is applied to the washing machine. The dishwasher and washing machine are low-priority executions that can execute autonomously, assuming as they are loaded with the necessary items. Household 1 would now be able to lower its contracted power to 3.45 kVA.

KPI	Baseline	Managed
Peak	4300 W	2800 W
PAR	7.66	4.99
ADAWR	0.00	0.07

Table 2: Key Performance Indicators for household 1.

Analyzing the KPI in table 2, Peak and PAR are reduced by approximately 35%. It is a very significant reduction. The PAR values are relatively high, in consequence of the data representing a single user: consumption is residual during most of the day, when the house is vacant or the user is asleep. ADAWR is negligible.

4.2. Multi-house mode

To test the multi-house mode in an heterogeneous environment, the three houses represented by the household profiles created in this chapter were connected to the Aggregator and simulated together.

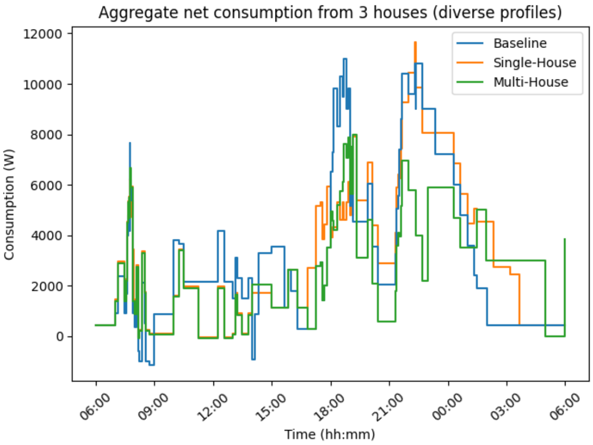

Figure 5: Comparison of aggregate consumption using heterogeneous profiles.

KPI	Baseline	Single	Multi
Peak	11003 W	11651 W	8002 W
PAR	3.60	3.78	3.25
ADAWR	0.00	0.06	0.12

Table 3: Key Performance Indicators for the heterogeneous simulation.

Baseline results compound the unrestricted consumption of each house, leading to a high peak. Surprisingly, the KPI data for the cluster in table 3 shows that the Peak, when operating in single-house mode, is even larger: each house, operating in its best interests without concerns about grid stability, will create a larger consumption peak at 10:20 PM, seen in figure 5. Using the communication with the Aggregator to improve the distribution of low-priority appliances, Peak is reduced to 8002 W. In this scenario, multi-house mode offers a 27% peak reduction in the network shared between the three houses against the baseline, with slightly more disruption to the user than in single-house mode.

KPI	House 1	House 2	House 3
Peak	2800 W	4150 W	4926 W
ADAWR	0.24	0.06	0.07

Table 4: Key Performance Indicators for the individual households, operating in multi-house mode.

Table 4 shows the influence of the operation mode in the user-centered metrics. ADAWR is low across the three houses, although the value in household

1 is larger than in single-house mode. Due to the washing machine of household 3 being scheduled in advance to 06:00 PM by the Aggregator, utilizing its maximum delay, the house is forced to raise the consumption threshold in order to schedule cooking appliances at the same time. This is an undesired behavior resulting from the scheduling happening in real time. However, the increased Peak is still below the recommended threshold for households with a contracted power of 5.75 kVA. Thus, in this situation, the benefits provided to the user were not affected.

5. Conclusions

In this paper, an interactive scheduler focused on peak control is proposed. It is a server application able to start, delay, interrupt and resume appliances based on demand at the time of request, cooperate with other instances, include power generation from alternative sources and utilize battery energy storage systems as peak shaving and load balancing mechanisms.

The results show high effectiveness of the presented solution in most consumption scenarios. In single-house mode, the system is able to reduce consumption peaks by up to 35% by scheduling consumer appliances more efficiently. When used along with a PV system and a BSS, the scheduler is able to shift the energy produced in excess to later periods of the day, and use the battery discharges both as a peak shaving and a load balancing mechanism. Residents in households 1 and 3 were able to reduce their contracted power, receiving an economic incentive. In multi-house mode, the cooperative distribution of low-priority appliances reduced the aggregate Peak by 27% in an heterogeneous environment, at the expense of limiting the ability to shift appliances locally for higher priority requests.

The system could be further iterated with improvements on the algorithm and energy estimation. In a more advanced solution, energy resources could be shared within the microgrid of the shared residential space. Excess PV and BSS energy could be distributed and requested by other households in a local energy market, or the resources could be pooled together to optimize aggregate consumption.

References

[1] X. Chen, T. Wei, and S. Hu. Uncertainty-aware household appliance scheduling considering dynamic electricity pricing in smart home. *IEEE Transactions on Smart Grid*, 4(2):932–941, 2013.

[2] E. H. Et-Tolba, M. Ouassaid, and M. Maaroufi. Smart home appliances modeling and simulation for energy consumption profile development: Application to moroccan real environment case study. In *2016 International Renewable and Sustainable Energy Conference (IRSEC)*, pages 1050–1055, 2016.

[3] X. Fang, S. Misra, G. Xue, and D. Yang. Smart grid — the new and improved power grid: A survey. *IEEE Communications Surveys Tutorials*, 14(4):944–980, 2012.

[4] M. G M, C. Vyjayanthi, and C. Modi. *An Open-Hardware Approach for IoT Enabled Smart Meter Monitoring and Controlling System Using MQTT Protocol*, pages 303–317. Springer, 03 2020.

[5] D. Han and J. Lim. Smart home energy management system using ieee 802.15.4 and zigbee. *IEEE Transactions on Consumer Electronics*, 56(3):1403–1410, 2010.

[6] B. MacLeod. Benefits of modern energy management platforms. *Condominium Manager*, page 49–50, Dec 2019.

[7] G. M. Masters. *Renewable and efficient electric power systems*. John Wiley & Sons, 2013.

[8] A. Mohsenian-Rad, V. W. S. Wong, J. Jatskevich, R. Schober, and A. Leon-Garcia. Autonomous demand-side management based on game-theoretic energy consumption scheduling for the future smart grid. *IEEE Transactions on Smart Grid*, 1(3):320–331, 2010.

[9] Odysee-Mure. Energy efficiency trends for households in the EU, 2021. Last accessed on 27 May, 2022.

[10] Y. Ozturk, D. Senthilkumar, S. Kumar, and G. Lee. An intelligent home energy management system to improve demand response. *IEEE Transactions on Smart Grid*, 4(2):694–701, 2013.

[11] G. Pinto, A. Kathirgamanathan, E. Mangina, D. P. Finn, and A. Capozzoli. Enhancing energy management in grid-interactive buildings: A comparison among cooperative and coordinated architectures. *Applied Energy*, 310:118497, 2022.

[12] M. Pipattanasomporn, M. Kuzlu, S. Rahman, and Y. Teklu. Load profiles of selected major household appliances and their demand response opportunities. *Smart Grid, IEEE Transactions on*, 5:742–750, 03 2014.

[13] T. Stetz, M. Rekinger, and I. Theologitis. Transition from uni-directional to bi-directional distribution grids. *Management Summary of IEA Task*, 14, 2014.