Optimizing Residential Energy Resources with an Improved Multi-Objective Genetic Algorithm based on Greedy Mutations

Ivo Gonçalves
INESC Coimbra, DEEC, University of Coimbra
Coimbra, Portugal
icpg@dei.uc.pt

Álvaro Gomes
INESC Coimbra, DEEC, University of Coimbra
Coimbra, Portugal
agomes@deec.uc.pt

Carlos Henggeler Antunes
INESC Coimbra, DEEC, University of Coimbra
Coimbra, Portugal
ch@deec.uc.pt

ABSTRACT

Energy management is increasingly becoming an important issue in face of the penetration of renewable generation and the evolution to smart grids. Home energy management systems are aimed to make the integrated optimization of residential energy resources, taking into account energy prices and end-user's requirements. This paper addresses a residential scenario where energy resources are automatically managed to reduce the overall energy cost while considering a set of user-defined comfort preferences. These energy resources include the grid, shiftable appliances, thermostatic loads, static batteries, electric vehicles, and local energy production. The comfort specifications consist of the time slots where the shiftable appliances are preferred to operate and the temperature ranges desired for the thermostatically controlled loads. The conflicting objectives are addressed by a multi-objective genetic algorithm that aims to minimize the overall energy cost and the user's dissatisfaction. This paper proposes a set of novel operators that result in statistically significant improvements in terms of hypervolume values when compared to a recently proposed genetic algorithm customized to address this same scenario. These novel operators include a different population initialization, a greedy mutation, and two geometric crossovers. The effect of the proposed operators on the resulting allocation of energy resources is analyzed.

CCS CONCEPTS

• **Applied computing → Multi-criterion optimization and decision-making**;

KEYWORDS

Demand response, Energy Management Systems, Genetic Algorithms, Load Management, Smart Grids

ACM Reference Format:
Ivo Gonçalves, Álvaro Gomes, and Carlos Henggeler Antunes. 2018. Optimizing Residential Energy Resources with an Improved Multi-Objective Genetic Algorithm based on Greedy Mutations. In *GECCO '18: Genetic and Evolutionary Computation Conference, July 15–19, 2018, Kyoto, Japan.* ACM, New York, NY, USA, Article 4, 8 pages. https://doi.org/10.1145/3205455.3205616

GECCO '18, July 15–19, 2018, Kyoto, Japan

ACM ISBN 978-1-4503-5618-3/18/07…$15.00
https://doi.org/10.1145/3205455.3205616

1 INTRODUCTION

Renewable energy generation and efficient use of energy are key issues in the pursuit of sustainable development, both strongly benefiting from an adequate integrated management of all energy resources in power grids. In the context of smart digital grids, energy management systems (EMS) endowed with appropriate management algorithms will play an increasingly important role to make the most from renewable sources and controllable energy resources on the demand side [1, 9]. Automated EMS will enable the flexibility in the usage of some end-use loads to have a role in minimizing the electricity bill and also providing services required by the grid. Automated EMS will control energy demand by responding to different possible input signals, such as time-differentiated energy prices, requests from the grid or from an aggregator entity to increase/decrease consumption, direct control actions or according to available local generation (usually based on renewables) or storage [3, 6, 11]. Demand response (DR) has the potential to improve the efficiency of infrastructure capacity, decrease peak load demand, reduce GHG emissions levels, provide ancillary services, and enhance overall grid sustainability. The potential of DR programs is vast and new enabling framework policies are emerging.

From the perspective of a consumer, two objectives are usually at stake: minimize energy costs (or maximize the profits with local energy resources considering selling back to the grid), and minimize the impacts on the quality of the energy services provided by the end-use loads [1, 5, 8]. The quality of the energy services is often measured as a comfort/discomfort indicator, quantifying how much the energy service is far from the end-user's preferences and requirements. For example, the preferred time slots for the operation of certain appliances or the admissible variation in temperature for thermostatically controlled loads are usual preferences set by consumers.

The optimization of this type of decisions involves the resolution of complex models of combinatorial nature, which in general cannot be solved with a reasonable computational effort compatible with (near) real-time application. Metaheuristics have shown to provide good solutions for these optimization problems, being able to deal with multiple objective functions and combinatorial complexity [4, 6, 9]. Furthermore, metaheuristics are effective to deal with technical constraints related with appliance operation, end-user's preferences (utilization of the energy service, thermal comfort or other measures of the quality of the energy service provided), and the capability to tackle unexpected events.

In a smart grid scenario with widespread use of automated EMS, it is critical that solutions be able to comply as much as possible with the diversity of requirements and preferences among end-users.

Automated control will replace the control actions of end-users over the loads whenever the consumer is not at home or is not available/capable to react to different input signals. In the scope of a multi-objective model balancing the economic and quality of service dimensions, the trade-off region depends on the end-user profile (e.g., privileging solutions with lower costs and possibly worse values of comfort or willing to have a higher energy bill by caring mainly about comfort levels). Given the need to deal with different end-user's preferences and requirements, it is important that optimization algorithms are able to identify diverse solutions. In addition, as the automated EMS is supposed to operate 24/7 and at the same time consumers can control some end-use loads, resulting in a variable base load out of the control of the EMS, the ability to re-schedule the utilization of energy resources due to unpredicted usage of some loads is of utmost importance.

This paper addresses a residential scenario where energy resources are automatically managed to reduce the overall energy cost while considering a set of user-defined comfort preferences. A set of novel operators are proposed which include a different population initialization, a greedy mutation, and two geometric crossovers. The resulting algorithmic variants are compared against a recently proposed genetic algorithm customized to address this same scenario. The paper is organized as follows: section 2 presents the residential scenario under study; section 3 describes the baseline variant and the newly proposed variants; section 4 presents and analyzes the experimental results; and section 5 draws some conclusions about the performance of the novel variants.

2 RESIDENTIAL SCENARIO

In the residential scenario addressed, a set of energy resources must be automatically managed to reduce the overall energy cost while considering user-defined comfort preferences. These energy resources include the grid, shiftable appliances, thermostatically controlled loads, static batteries, electric vehicles, and local energy production. The scenario under study is the same used by Soares et al. [9]. The aim in this scenario is to minimize energy costs while also minimizing the discomfort caused to the users. This discomfort is computed based on the time slots where the shiftable appliances are preferred to operate and the temperature ranges desired for the thermostatically controlled loads. For this scenario, the EMS is endowed with a multi-objective genetic algorithm to compute solutions trading-off the cost and discomfort objective functions. This EMS must define the actions to implement over the considered energy resources for a given planning period. In the residential scenario studied, the planning period comprises 2160 minutes (36 hours). As this scenario is rather complex and with a large amount of information required, and as the data used (the same as used by Soares et al. [9]) is not yet openly available, it is provided as supplementary material to this paper in order to permit full reproduction of the results.

The four main types of manageable loads considered are: Shiftable Loads (SLs), Thermostatically Controlled Loads (TCLs), Stationary Storage System (SSS), and Electric Vehicle (EV). The EMS should be able to deal with any number of manageable loads. The residential scenario studied is composed of three SLs, three TCLs, one SSS, and one EV. The three SLs considered are: a dishwasher, a laundry machine, and a tumble dryer. Each SL has a given operation cycle which defines how long the load takes to complete and the power needed at each minute for the load to operate. The EMS must allocate a region for each SL to operate in a given planning period, while ensuring that each SL will be able to complete its operation cycle by the end of the planning period. In the algorithmic implementations considered in this paper, this is ensured by the design of the operators as described in the next section. The three TCLs considered are: an Air Conditioner (AC), an Electric Water Heater (EWH), and a fridge. The EMS must, for each TCL, define the target temperature for all minutes of the planning period. From these target temperatures, the effective temperatures are then computed by a physically-based model (PBM) that considers the physical characteristics of the equipment and its surroundings. Each TCL has a different PBM associated. The PBMs used are the same used by Soares et al. [9]. For the full PBM details the reader is referred to Soares [10]. Regarding the SSS, the EMS must define in which state it is going to operate in each minute of the planning period. There are four possible SSS states encoded as follows: -2, self-consumption and selling excess electricity to the grid; -1, selling electricity to the grid; 0, not in use; 1, charging from the grid. The SSS also has a state of charge associated. The EMS must ensure that the SSS does not go below a given minimum state of charge, a value set by the user. The EMS ensures this by only effectively executing states that extract energy stored (-2 and -1) if they do not result in a violation of the minimum state of charge. In this case, the state is interpreted as neutral (state 0). The minimum state of charge considered here is 20%. Naturally, the state of charge also cannot go above 100%. Similarly, this is ensured by the EMS by only effectively executing the charging state (state 1) if the state of charge is below 100%. With respect to the EV, the EMS must ensure that it is charged by the end of the planning period. In the scenario considered, charging the EV takes 480 minutes (8 hours). Two points in the planning period define the region where the EV is available for charging. In the scenario considered, the EV is available between minutes 200 and 1998 of the planning period. As described in the next section, the algorithmic implementations considered in this paper ensure, by design of the operators, that the EV availability is respected and that it is charged by the end of the planning period. As in Soares et al. [9], the EV is only used in grid-to-vehicle mode, i.e., the battery of the EV is not used for self-consumption or selling electricity to the grid. Because of this, the EMS represents the decisions regarding the EV as states of 0 or 1 for each minute of the planning period. Besides these manageable loads, a non-manageable base load is also considered. As the name entails, this load is not subject to actions of the EMS. However, the power requested by the non-manageable base load must also be taken into account when considering the total energy cost. The only constraint that cannot be directly ensured by design of the operators is the contracted power constraint. The contracted power defines the overall power that can be requested to the grid in a given minute of the planning period.

With respect to the calculation of the total energy cost objective, the following definitions are considered:

T: Number of minutes of the planning period ($t = 1, \ldots, T$)

n: Number of shiftable loads (j = 1, ..., n)
m: Number of thermostatically controlled loads (b = 1, ..., m)
p_{jt}: Power requested by SL j at minute t (W)
p_{bt}: Power requested by TCL b at minute t (W)
v_t: Power requested by the SSS at minute t (W)
e_t: Power requested by the EV at minute t (W)
u_t: Power requested by the non-manageable base load at minute t (W)
a_t: Total power used for self-consumption at minute t (W)
s_t: Total power sold to the grid at minute t (W)
NGP_t: Net grid power requested at minute t (W)
BP_t: Buying price at minute t (€/kWh)
SP_t: Selling price at minute t (€/kWh)

The net grid power requested is defined as:

$$NGP_t = \sum_{j=1}^{n} p_{jt} + \sum_{b=1}^{m} p_{bt} + v_t + e_t + u_t - a_t$$

Then, the total energy cost (€) can be calculated as:

$$\sum_{t=1}^{T} \left(NGP_t \cdot \frac{BP_t}{1000 \cdot 60} - s_t \cdot \frac{SP_t}{1000 \cdot 60} \right)$$

The dissatisfaction objective aggregates three different components: a global interruption risk penalty, time slot penalties for SLs, and temperature variation penalties for TCLs. Regarding the calculation of the dissatisfaction objective, the following definitions are considered:

g_t: Global interruption risk penalty at minute t
C_t: Contracted power at minute t (kW)
h_{jt}: Time slot penalty for SL j at minute t
y_{jt}: Binary variable representing whether SL j is operating at minute t
r_{bt}: Temperature variation penalty for TCL b at minute t
T_{bt}: Temperature of TCL b at minute t determined by the PBM (°C)
HBT_b: Maximum reference temperature for TCL b (°C)
LBT_b: Minimum reference temperature for TCL b (°C)

The global interruption risk penalty is defined as:

$$g_t = \begin{cases} 1, & \text{if } NGP_t > 0.85 \cdot C_t \\ 0, & \text{otherwise} \end{cases}$$

The temperature variation penalty is defined as:

$$r_{bt} = \begin{cases} e^{\frac{T_{bt} - HBT_b}{HBT_b - LBT_b}} - 1, & \text{if } T_{bt} > HBT_b \\ e^{\frac{LBT_b - T_{bt}}{HBT_b - LBT_b}} - 1, & \text{if } T_{bt} < LBT_b \\ 0, & \text{otherwise} \end{cases}$$

Then, the dissatisfaction objective can be calculated as:

$$\sum_{t=1}^{T} \left(g_t + \sum_{j=1}^{n} h_{jt} \cdot y_{jt} + \sum_{b=1}^{m} r_{bt} \right)$$

3 VARIANTS

3.1 Baseline Algorithm and an Alternative SSS Initialization

The baseline algorithm considered for comparison is based on the implementation of Soares et al. [9]. This implementation was designed specifically to address the same residential scenario being studied here. Given the multi-objective nature of the problem considered, the implementation is based on the Non-dominated Sorting Genetic Algorithm II (NSGA-II) [2]. This baseline algorithm defines a set of customized operators that avoid, by design, some issues that could arise under non-customized operators. These customized operators include initializations, crossovers, and mutations. Note that given the nature of the problem, performing an operation on an individual consists of performing the operation on the four types of loads that the individual represents. In other words, an individual-wise operator is essentially a set of four load-wise operators. As an example, an (individual-wise) initialization operator consists of load-wise operators that are able to correspondingly initialize shiftable loads, thermostatically controlled loads, stationary storage systems, and electric vehicles. In the following descriptions, whenever a random selection is said to occur, this is meant to be interpreted as a random selection with uniform probability.

The initialization of a shiftable load works by randomly selecting a starting minute for the operation while ensuring that the operation cycle is fully completed by the end of the planning period. This is an example of an operator that, by design, avoids a possible issue (non-complete operation cycles) from appearing in the solutions. The TCL initialization operator randomly selects a temperature between the maximum and the minimum reference temperatures. This selected temperature is assigned to all minutes of the planning period, i.e., this initialization always results in a constant temperature throughout the planning period. The SSS initialization is performed while having as reference the average energy price of the planning period. If the price of a given minute is above or equal to the average price then a state of -2 or -1 is randomly selected to be assigned to that minute. On the other hand, if the price is below the average price then a state of 0 or 1 is randomly assigned to that given minute. The EV is initialized by randomly selecting a continuous region where the charging is going to occur. In the scenario considered, this region is necessarily composed of 480 minutes (the time required to charge the EV). The individual-wise baseline initialization can now be simply defined as the combination of these four load-wise initialization operators. Given that the baseline SSS initialization has a certain degree of inflexibility (by limiting the state choices based on the average price), an alternative SSS initialization is considered. This initialization divides the planning period into continuous blocks of 10 minutes, and for each block randomly selects any of the 4 states to be assigned. This is meant to confer a greater diversity to the initial population. The alternative individual-wise initialization is composed of the alternative SSS initialization as well as the remaining baseline load-wise initializations.

The crossovers in the baseline algorithm are all swap-based. The shiftable load crossover creates offspring by swapping the starting minute of the two parents. The TCL crossover swaps all the

temperature values of both parents, while preserving each corresponding location in the genotype. Similarly, the SSS and the EV crossovers swap all the states of both parents, while preserving each corresponding location in the genotype. The individual-wise baseline crossover can now be simply defined as the combination of these four load-wise crossovers. In terms of load-wise mutations, the shiftable load mutation is essentially a reinitialization, as it randomly selects a new starting minute for the operation while ensuring that the operation cycle is fully completed by the end of the planning period. The TCL mutation randomly selects a region, and the temperature in this region is mutated by a randomly selected amount. This amount is selected with uniform probability from the range [-deviation bound, deviation bound]. The deviation bound is a TCL specific parameter, as different TCLs can use different deviation bounds. In the baseline algorithm the values used for each deviation bound are the following: fridge and AC 2 °C, EWH 5 °C. Similarly to the TCL mutation, the SSS mutation randomly selects a region, and all the states in this region are changed to a randomly selected state of any of the 4 possible states. As in the shiftable load mutation, the EV mutation is essentially a reinitialization, as it randomly selects a new continuous region where the charging is going to occur. The individual-wise baseline mutation is now defined as the combination of the four load-wise mutations described.

3.2 Greedy Mutations

The idea behind the Greedy Mutation (GM) operator is to select particular characteristics of a solution based on the direct cost or comfort improvement. Hence the name of this operator, as it greedily targets some region of interest, as opposed to performing a random non-directed change. A greediness parameter is also introduced to provide greater flexibility. In this context, greediness refers to the selective pressure placed towards the selection of a given solution characteristic based on the direct improvement. Three load-wise greedy operators are proposed: two are cost-greedy, and the remaining is comfort-greedy. The first cost-greedy operator works over shiftable loads. Here, the goal is to select the region where the corresponding shiftable load is going to operate by greedily considering regions with lower total price. The greediness parameter is defined as a percentage of the total number of admissible regions on which a given shiftable load can operate. This total number of admissible regions depends on the planning period size and the operation cycle duration of the shiftable load considered. In the residential scenario being studied, the total number of admissible regions for each shiftable load is the following: dishwasher 2084, laundry machine 2069, and tumble dryer 2121. A tournament selection is performed to select the region on which the load will be assigned. Regions that enter the tournament are selected randomly while ensuring admissibility. When two regions compete in a tournament round, the region with the lower total price wins. The tournament size is defined by the greediness parameter. Naturally, the bigger the tournament the bigger the probability that a better region is selected. A similar functioning is defined for the remaining load-wise greedy operators. The second cost-greedy operator works in a similar fashion, but in this case for the selection of a single continuous charging region for the EV. The proposed comfort-greedy load-wise operator is also defined to work over shiftable loads. The reasoning is similar to the SL cost-greedy operator, with the exception that instead of looking at the total price of a given region, the operator selects based on the total time slot penalties accumulated. With these three load-wise greedy operators, the individual-wise GM operator can be defined as being composed of the following load-wise operators: SL cost and comfort greedy mutations, TCL baseline mutation, SSS baseline mutation, and EV cost-greedy mutation. Given that two operators exist for the shiftable loads, both are selected with 50% probability of being used. This alternation of operators over shiftable loads may also provide more diversity to the evolutionary process.

3.3 Geometric Crossovers

Another set of operators studied is based on the notion of geometric crossover. A geometric crossover [7] is a type of crossover where, by construction, all offspring are created in the region (or segment) between their parents. Assuming that the parents carry helpful genetic material, then probably the region between them might also be interesting to explore. The operation of a geometric crossover can be seen as an approximation from one parent to the other. The value that defines the approximation is usually generated randomly. Let this value be known as the perturbation (p) applied, and ranging from 0 to 1. An abstract geometric crossover operating over a single gene can be defined as producing an offspring, O, from the parents, $P1$ and $P2$, as follows:

$$O = P1 \cdot p + P2 \cdot (1 - p) \tag{1}$$

Following a similar reasoning, a shiftable load geometric crossover can be defined. Let $x_j(I)$ be the starting minute of shiftable load j in a given individual I. The corresponding geometric crossover follows as:

$$x_j(O) = x_j(P1) \cdot p + x_j(P2) \cdot (1 - p) \tag{2}$$

Note that, by construction, the offspring produced by this geometric crossover always operates in an admissible region, in the sense that the SL is guaranteed to be able to finish its operation cycle by the end of the planning period. This idea could also be extended to implement geometric crossovers over TCLs and the SSS. Here, the most important difference is that these loads are represented by multiple genes, i.e., one value for each minute of the planning period. This creates the possibility of defining two different geometric crossovers. The first one selects a different perturbation for each minute, while the second one uses the same perturbation for all minutes. The first one is simply named geometric, and the second one is name Fixed Perturbation geometric (FP geometric). Let $S_t(I)$ be the state of the SSS at minute t in a given individual I. The corresponding SSS FP geometric crossover follows as:

$$S_t(O) = S_t(P1) \cdot p + S_t(P2) \cdot (1 - p) \tag{3}$$

While the SSS geometric crossover is defined by extending p to represent a different value per minute:

$$S_t(O) = S_t(P1) \cdot p_t + S_t(P2) \cdot (1 - p_t) \tag{4}$$

A similar reasoning is followed to create a geometric and a FP geometric crossover for the TCLs. Naturally, in the case of the TCLs the operations are performed over the target temperature vector. With these five load-wise geometric operators, two individual-wise geometric operators can be defined. The first one is simply named Geometric and is composed of the following load-wise operators: geometric SL crossover, geometric TCL crossover, geometric SSS crossover, and EV baseline crossover. The second one is named FP geometric and is composed of the following load-wise operators: geometric SL crossover, FP geometric TCL crossover, FP geometric SSS crossover, and EV baseline crossover.

4 EXPERIMENTAL STUDY

4.1 Parameters and Methodology

Four main variants are experimentally assessed: baseline, Fixed Perturbation geometric (FP geometric), Geometric, and Greedy Mutation (GM). The baseline variant is based on the implementation of Soares et al. [9] as described in 3.1. The FP geometric variant uses the FP geometric crossover defined in 3.3, as well as the baseline mutation. The Geometric variant uses the Geometric crossover defined in 3.3, and also uses the baseline mutation as its mutation operator. The GM variant uses the greedy mutation defined in 3.2, and the baseline crossover. Different levels of greediness are assessed in the GM variant: 1%, 5%, 15%, 25%, 50%, and 75%. The main parameter values are the same as used by Soares et al. [9]. Each variant performs 30 independent runs. All variants use the same initial population for each equivalent run in order to exclude the possible effect that different populations might have in the outcomes. Experiments are conducted with populations of 50 individuals, being able to evolve for 300 generations. Parent selection is performed with a binary tournament. The individual-wise operators are always applied. The control over the usage of the variation operators is performed at the load-wise level. The load-wise crossover probabilities are the following: SLs 0.5, TCLs 0.5, SSS 0.3, and EV 0.3. The load-wise mutation probabilities are the following: SLs 0.2, TCLs 0.6, SSS 0.3, and EV 0.2.

Claims of statistical significance are based on Mann-Whitney U tests, using a Bonferroni correction, and considering a significance level of $\alpha = 0.05$. A non-parametric test is used because the data are not guaranteed to follow a normal distribution. To compare the different variants, the hypervolume indicator [12] is also considered. The reference point used for the computation of the hypervolume has the following coordinates: 10 for the cost, and 2000 for the dissatisfaction. These values were empirically confirmed to ensure that no solution with higher coordinates was found during the evolution. Preliminary results confirmed that different variations on the reference point coordinates do not influence the relative comparison of the variants assessed.

4.2 Results

Figure 1 shows the median hypervolume evolution over the 30 runs of the considered variants under both initialization scenarios. The first important difference to highlight is that the alternative SSS initialization proposed achieves, with statistical significance, higher hypervolumes than the baseline initialization (p-value: 2.872×10^{-11}). Considering the evolution of the baseline variant under the different initializations, this statistically significant improvement is maintained after conducting 20 generations (p-value: 1.438×10^{-6}). These quick improvements are crucial in a real-world scenario, where a system using these optimizers should ideally reach relatively good solutions as fast as possible when some update is necessary, e.g., if the user updates his comfort preferences. The proposed alternative SSS initialization contributes to this faster convergence to better solutions. At around generation 100 the median values start to become similar. By the end of the evolution, no statistically significant difference is found in terms of the hypervolumes reached.

Focusing on the results under the baseline initialization, GM variants are also able to provide better solutions than the baseline variant in a few generations. Regardless of the greediness used, all GM variants achieve, with statistical significance, higher hypervolumes at generation 20. The resulting p-values are: 1.100×10^{-4} for GM 1%, 2.921×10^{-4} for GM 5%, 9.273×10^{-4} for GM 15%, 1.286×10^{-4} for GM 25%, 1.800×10^{-3} for GM 50%, and 5.434×10^{-5} for GM 75%. Interestingly, this greedy mutation seems robust with respect to the greediness parameter, at least in terms of the median hypervolume evolution. By the end of the evolution, the best GM variant (GM 50%) is also able to achieve significantly higher hypervolumes (p-value: 3.465×10^{-4}). This means that under an appropriate parametrization, the GM variant is also able to provide long-term improvements. Despite the other parametrizations also achieving higher medians than the baseline variant, this is not sufficient to result in statistical significance.

The baseline variant is able to outperform, with statistical significance, the Geometric variant by the end of the evolution (p-value: 3.660×10^{-7}). On the other hand, despite the gap in the final hypervolumes medians, the differences are not statistically significant in the comparison between the baseline variant and the FP geometric variant.

Table 1 presents the best, worst, median, average, and standard deviation (SD) values of each objective for all the variants under the baseline initialization at the end of the evolution. All GM variants are able to achieve lower best costs and dissatisfactions than the baseline. These variants also achieve better median and average costs than the baseline. The FP geometric and Geometric variants also achieve lower best dissatisfactions than the baseline. What might explain the lack of competitiveness of these variants is the lower exploration of solutions that result in considerably higher dissatisfactions. This can be seen from the significantly lower median and average dissatisfactions that these variants achieve in comparison with the remaining variants. In the scenario considered, maximizing the hypervolume entails exploring solutions that result in higher dissatisfactions but lower costs. This is something that the nature of the geometric crossovers do not allow them to achieve, as, by definition, the temperature values selected are always within the range of the temperatures of the parents. Achieving high dissatisfactions can only be achieved by going beyond the defined maximum and minimum reference temperatures.

Figure 2 shows the Pareto front with the overall highest hypervolume of all the runs of all the variants. This Pareto front was evolved by variant GM 5%, and it achieves a hypervolume of approximately 14489. This front is similar to other fronts found in other runs and other variants. The shape of the front is similar to

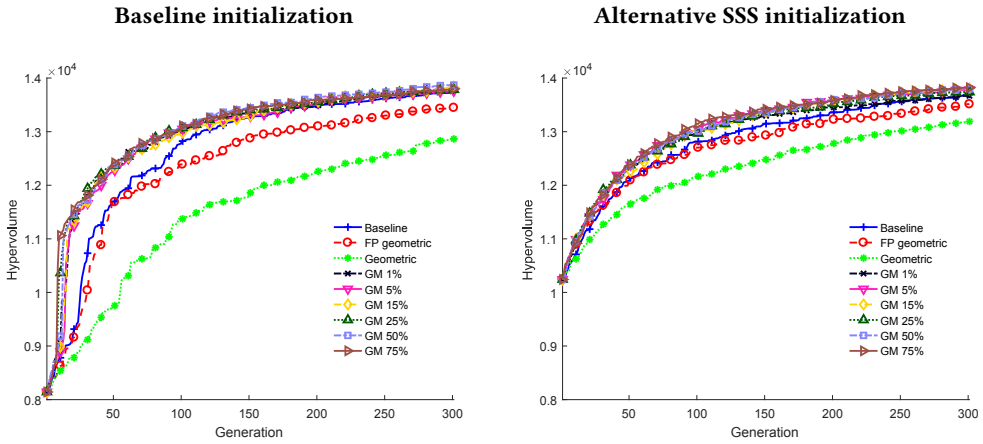

Figure 1: Median hypervolume evolution under the baseline initialization (left), and the alternative SSS initialization (right)

Table 1: Summary of the values achieved in both objectives for each variant under the baseline initialization

Variant	Cost (€)					Dissatisfaction				
	Best	Worst	Median	Average	SD	Best	Worst	Median	Average	SD
Baseline	2.93	5.76	3.57	3.68	0.51	0.457	1313.982	15.316	132.720	254.702
FP geometric	2.97	5.72	3.69	3.78	0.48	0.337	1178.151	10.584	81.472	178.927
Geometric	3.22	5.91	3.99	4.09	0.50	0.360	1127.872	8.231	47.118	125.711
GM 1%	2.89	5.90	3.46	3.59	0.49	0.237	1227.600	17.160	134.181	248.366
GM 5%	2.70	5.70	3.49	3.62	0.52	0.203	1368.942	16.025	138.100	265.160
GM 15%	2.76	5.92	3.41	3.58	0.54	0.268	1365.127	19.228	147.406	269.315
GM 25%	2.76	5.51	3.46	3.60	0.51	0.382	1370.464	15.990	138.106	259.581
GM 50%	2.79	5.80	3.38	3.54	0.51	0.305	1416.478	23.313	164.678	283.367
GM 75%	2.90	5.26	3.48	3.60	0.49	0.274	1288.593	16.908	150.005	269.505

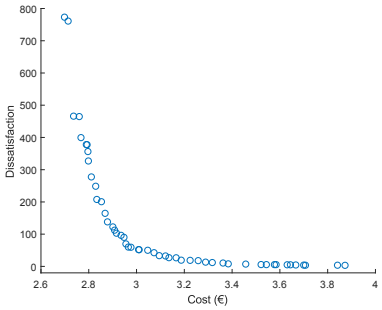

Figure 2: Pareto front with the overall highest hypervolume achieved (approximately 14489) evolved by variant GM 5%

what is commonly found in similar problems. Perhaps the most noticeable feature is the considerable dissatisfaction gap between the two solutions with lower cost and higher dissatisfaction and the remaining solutions. This gap can be explained by the exponentially increasing temperature penalization described in section 2. Solutions might reduce cost but at the expense of a considerable increase in the resulting dissatisfaction.

4.3 Analysis

In the context of the scenario under study, an important analysis aspect is the resulting physical characteristics found in the solutions. One way of analyzing these characteristics is through a power profile plot. A power profile plot visually identifies the power used by each load in a given period. In the power profile plots presented in this section, besides the eight manageable loads, the following information is also presented: the non-manageable base load, the contracted power, the amount of self-consumption, the power sold to the grid, and the price for each minute. The amount of self-consumption and the power sold to the grid are presented in negative values.

Figure 3 shows the power profiles of the solutions with the overall best cost and the overall best dissatisfaction found across all runs and all variants. The solution with the overall best cost is found by variant GM 5%. This solution achieves a cost of approximately 2.70, while resulting in a dissatisfaction of approximately 771.975. The solution with the overall best dissatisfaction is found by variant GM 50% under the alternative SSS initialization. This solution achieves a dissatisfaction of approximately 0.197, while resulting in a cost of approximately 4.65. Figure 4 presents some of the crucial data to complement this analysis: the time slot penalties of the shiftable loads and the hot water consumption.

With respect to the best cost solution, it is clear that this solution shifts consumption away from higher price regions. In particular where it concerns the shiftable loads, most of these resource allocations are clearly done at the expense of the resulting dissatisfaction. The only shiftable load that does not incur in a time slot penalization is the dishwasher. On the other hand, the laundry machine and

tumble dryer do incur in considerable time slot penalizations. The high dissatisfaction achieved (approximately 771.975) also allows to infer that the temperatures of the TCLs are considerably above the preferred range. This can also be assessed from the sporadic power usage from these loads. The fact that the scenario considered does not treat temperatures as hard constraints allows a greater flexibility of the algorithm to substantially reduce the usage of the TCLs. These hard temperature constraints should be considered in a future scenario. With respect to the EV, this solution allocates the charging to a lower price region (between minutes 1441 and 1921). Note that since the charging operation is being considered as a single continuous block and given that it takes 480 minutes to fully charge the EV, it is not possibly to allocate the charging without partially including a higher price region. In the case of this solution, this partial higher price charging occurs between minutes 1862 and 1921. However, this is partially compensated by self-consumption in the same region, which reduces the contribution of this region to the overall cost. The other lower price region alternative for the EV charging would be between minutes 200 and 680 (the EV is only available starting from minute 200). However, this alternative charging would partially include an even longer higher price region (between minutes 422 and 680). Similarly to the EV, the SSS also charges from the grid in the lower price regions. The energy stored is then used to self-consume and sell to the grid, almost always, in higher price regions. The only exception happens around minute 400 where the energy stored is sold to the grid in a lower buying and selling price region. This means that this particular solution could reduce the overall cost even more by, for instance, delaying this selling operation a few minutes. This would result in an admissible solution with the same dissatisfaction and an improved cost. Besides the above-mentioned case where the energy stored was used to compensate the power used by the EV charging, a similar case occurs between minutes 1200 and 1400 where the energy stored is used to counteract some peaks of the non-manageable base load.

Regarding the best dissatisfaction solution, perhaps the most noticeable difference is the behavior of the TCLs. Their more regular power usage is needed to unsure the desired temperature range. In terms of shiftable loads, the laundry machine and the tumble dryer are allocated to different regions than in the best cost solution. The allocation of the tumble dryer results in no time slot penalization. This particular allocation is only possible given the energy being self-consumed from the battery. If this was not the case, the combination of the power required by the EV, the AC, and the tumble dryer between minutes 1700 and 1800 would result in a violation of the contracted power. This long self-consumption in this region is only possible given the long battery charging period that occurs in the beginning of the planning period (up until minute 800). This shows that although the SSS does not directly influence comfort, it can still be indirectly crucial in allowing certain comfort scenarios that would not be possible without the provided self-consumption capability. As in the best cost solution, the EV is mostly in a lower price region. This can be possibly related with the usage of the greedy mutation operator in the variant that achieves this solution.

Figure 5 shows the temperatures of the TCLs along the planning period of two solutions: an intermediate dissatisfaction solution, and the overall best dissatisfaction solution analyzed above. The intermediate dissatisfaction solution is selected from the GM 5%

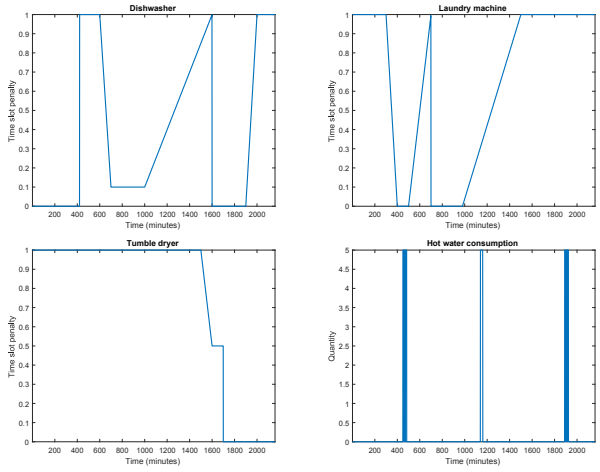

Figure 4: Time slot penalties of the shiftable loads: dishwasher (top left), laundry machine (top right), and tumble dryer (bottom left); and hot water consumption (bottom right)

variant. This solution achieves a dissatisfaction of approximately 3.444, while resulting in a cost of approximately 3.85. Figure 5 also includes, for each TCL, the maximum and minimum reference temperatures from which the temperature penalties are computed. In the fridge and the AC temperature variation, both solutions are mostly able to maintain the temperatures on the preferred range. Naturally, the intermediate dissatisfaction solution is not as precise. With respect to the EWH, both solutions have periods where the temperature falls below the minimum reference temperature, e.g., around minute 1900. The sudden temperature drops happen when a considerable amount of hot water is consumed. In these cases, in order to avoid falling below the minimum reference temperature, the water in the EWH must be set to a higher temperature to counteract the subsequent hot water consumption. This strategy may avoid a temperature penalization. However, the risk of going above the maximum reference temperature also exists. This is precisely what happens in the intermediate dissatisfaction solution before minute 1200. Around this time a considerable hot water consumption occurs. This consumption significantly reduces the EWH temperature. To counteract this, this solution decides to increase the EWH temperature around minute 1000. This increase is able to avoid the temperature from falling below the minimum reference temperature after the hot water consumption, but in turn it entails going above the maximum reference temperature before the same hot water consumption event.

5 CONCLUSIONS

This paper studied a residential scenario where a set of energy resources must be automatically managed to reduce the overall energy cost while considering user-defined comfort preferences. A set of novel operators were proposed and experimentally assessed. The resulting algorithmic variants were compared against a recently proposed genetic algorithm customized to address this same scenario. Results showed that the best-performing novel variants were

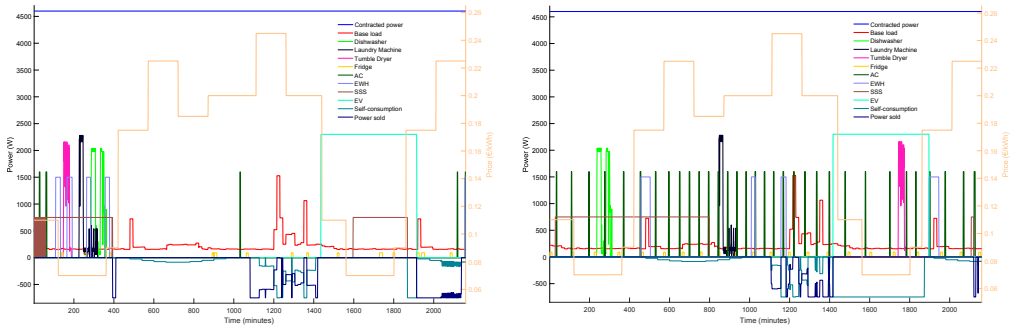

Figure 3: Power profile of the solution with the overall best cost (2.70€, dissatisfaction of 771.975) in the left, and the solution with the overall best dissatisfaction (0.197, cost of 4.65€) in the right

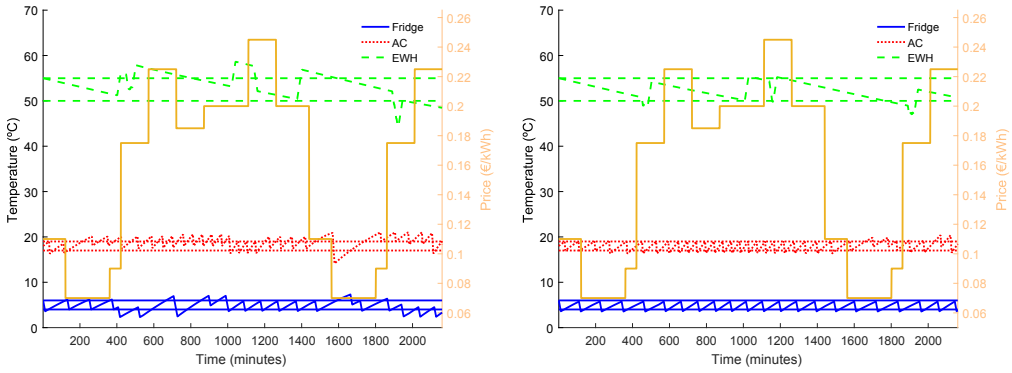

Figure 5: Temperatures of the TCLs of a solution with an intermediate dissatisfaction (3.444, cost of 3.85€) in the left, and the solution with the overall best dissatisfaction (0.197, cost of 4.65€) in the right

able to achieve both short-term and long-term improvements. The short-term improvements are particularly important in real-time scenarios where it is expected that an energy management system should be able to quickly provide interesting solutions when some update is issued.

ACKNOWLEDGMENTS

This work was partially supported by projects UID/MULTI/00308/2013 and by the European Regional Development Fund through the COMPETE 2020 Programme, FCT - Portuguese Foundation for Science and Technology and Regional Operational Program of the Center Region (CENTRO2020) within projects ESGRIDS (POCI-01-0145-FEDER-016434) and SUSPENSE (CENTRO-01-0145-FEDER-000006).

REFERENCES

[1] Benjamin Biegel, Lars Henrik Hansen, Jakob Stoustrup, Palle Andersen, and Silas Harbo. 2014. Value of flexible consumption in the electricity markets. *Energy* 66 (2014), 354–362.

[2] Kalyanmoy Deb, Amrit Pratap, Sameer Agarwal, and TAMT Meyarivan. 2002. A fast and elitist multiobjective genetic algorithm: NSGA-II. *IEEE transactions on evolutionary computation* 6, 2 (2002), 182–197.

[3] Ruilong Deng, Zaiyue Yang, Mo-Yuen Chow, and Jiming Chen. 2015. A survey on demand response in smart grids: Mathematical models and approaches. *IEEE Transactions on Industrial Informatics* 11, 3 (2015), 570–582.

[4] Jawad Khoury, Rita Mbayed, Georges Salloum, and Eric Monmasson. 2016. Predictive demand side management of a residential house under intermittent primary energy source conditions. *Energy and Buildings* 112 (2016), 110–120.

[5] Jinghuan Ma, He Henry Chen, Lingyang Song, and Yonghui Li. 2016. Residential load scheduling in smart grid: A cost efficiency perspective. *IEEE Transactions on Smart Grid* 7, 2 (2016), 771–784.

[6] Mohammad H Moradi, Mohamad Abedini, and S Mahdi Hosseinian. 2016. A combination of evolutionary algorithm and game theory for optimal location and operation of DG from DG owner standpoints. *IEEE Transactions on Smart Grid* 7, 2 (2016), 608–616.

[7] Alberto Moraglio. 2007. *Towards a Geometric Unification of Evolutionary Algorithms.* Ph.D. Dissertation. Department of Computer Science, University of Essex, UK.

[8] Ditiro Setlhaolo, Xiaohua Xia, and Jiangfeng Zhang. 2014. Optimal scheduling of household appliances for demand response. *Electric Power Systems Research* 116 (2014), 24–28.

[9] Ana Soares, Álvaro Gomes, Carlos Henggeler Antunes, and Carlos Oliveira. 2017. A Customized Evolutionary Algorithm for Multiobjective Management of Residential Energy Resources. *IEEE Transactions on Industrial Informatics* 13, 2 (2017), 492–501.

[10] Ana Raquel Gonçalves Soares. 2016. *Integrated management of residential energy resources: Models, algorithms and application.* Ph.D. Dissertation. Universidade de Coimbra, Portugal.

[11] Koen Vanthournout, Benjamin Dupont, Wim Foubert, Catherine Stuckens, and Sven Claessens. 2015. An automated residential demand response pilot experiment, based on day-ahead dynamic pricing. *Applied Energy* 155 (2015), 195–203.

[12] Eckart Zitzler and Lothar Thiele. 1998. Multiobjective optimization using evolutionary algorithms - a comparative case study. In *international conference on parallel problem solving from nature.* Springer, 292–301.