

# Design of an Automatic Demand-Side Management System Based on Evolutionary Algorithms

Edgar Galván-López, Adam Taylor, Siobhán Clarke and Vinny Cahill  
Distributed Systems Group,  
School of Computer Science and Statistics,  
Trinity College Dublin.  
edgar.galvan, tayloral, siobhan.clarke, vinny.cahill @ scss.tcd.ie

## ABSTRACT

Demand-Side Management (DSM) refers to programs that aim to control the energy consumption at the customer side of the meter. Different techniques have been proposed to achieve this. Perhaps the most popular techniques are those based on smart pricing (e.g., critical-peak pricing, real-time pricing). The idea, in a nutshell, is to encourage end users to shift their load consumption based on the price at a particular time (e.g., the higher the price, the less number of electric appliances are expected to be turned on). Motivated by these techniques (e.g., a strong positive correlation between the number of appliances being used and the electricity cost), we propose the use of a stochastic evolutionary-based optimisation technique, Evolutionary Algorithms, to *automatically* generate optimal, or nearly optimal, solutions that represent schedules to charge a number of electric vehicles (EVs) with two goals: (a) that each EV is as fully charged as possible at time of departure, and (b) to avoid charging them at the same time, whenever possible (e.g., load reduction at the transformer level). Instead of using a price signal to shift load consumption, we achieve this by considering what all the EVs might do at a particular time, rather than considering an interaction between an utility company and its user, as normally adopted in DSM programs. We argue that exploiting the interaction of these EVs is crucial at achieving excellent results because it carries the notion of smart pricing (e.g., balance energy usage), which is highly popular in DSM systems. Thus, the main contribution of this work is the notion of load shifting, borrowed from smart pricing methods, implemented in an evolutionary-based algorithm to automatically generate optimal solutions. To test our proposed approach, we used a dynamic scenario, where the state of charge of each EV is different for every day of our 28 days testing period. The results obtained by our proposed approach are highly encouraging in both: EVs being almost fully charged at time of the departure and the transformer load being reduced as a result of avoiding turning on the EVs at the same time.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

SAC'14 March 24-28, 2014, Gyeongju, Korea.

Copyright 2014 ACM 978-1-4503-2469-4/14/03 ...\$15.00.

<http://dx.doi.org/10.1145/2554850.2554959>

## Categories and Subject Descriptors

I.2 [Artificial Intelligence]: Automatic Programming; D.2.8 [Software Engineering]: Metrics—*complexity measures, performance measures*

## General Terms

Smart Grids

## Keywords

Evolutionary Algorithms, Bio-inspired Techniques, Smart Grids, Demand-Side Management Systems

## 1. INTRODUCTION

Demand-Side Management (DSM) is an important research area in the Smart Grid (SG) community as shown by the increasing number of publications over the years (e.g., [5, 6, 9, 10, 11, 12, 13, 15, 18, 19, 20]). DSM is normally considered as a mechanism or program, implemented by utility companies, to control the energy consumption at the customer side [12].

The high importance of the DSM in the SG scenario can be understood by considering the new challenges that are continuously presented to the grid, for example the addition of more and more electric vehicles using the grid. Thus, there is clearly a challenge to keep the grid stable, while at the same time, ensuring that the demand is satisfied by using the already available infrastructure. In other words, the goal of a DSM is to efficiently use the available energy without the necessity of installing new transmission infrastructure.

DSM programs include different approaches (e.g., manual conservation and energy efficiency programs [1], Residential Load Management (RLM) [8, 13]), where RLM programs based on smart pricing are amongst the most popular methods.

The idea behind smart pricing is to encourage users to manage their loads, so that they can reduce electricity prices while, at the same time, the utility companies achieve at reducing peak-to-average ratio (PAR) in load demand by shifting consumption whenever possible [8]. Load-shifting is foreseen to become even more important in the grid as the result of more devices connecting to the grid, where some of them, such as electric vehicles, can double the average household load [13].

Motivated by the smart pricing-based approaches, we are interested in developing a demand-side *automatic intelligent* management system that shifts electricity consumption of

electric vehicles (EVs) by considering what the rest of the EVs do at a particular time, rather than considering a utility - user communication approach, which has been normally adopted in the DSM programs. That is, we believe that we could find optimum, or nearly optimum, solutions by allowing this interaction (i.e., information of all the EVs).

To this end, we use a stochastic optimisation evolutionary algorithm (EAs). The main contribution of this work is the notion of load shifting, borrowed from the popular smart pricing-based methods, by considering what all EVs do at a particular time, rather than considering an scenario where the interaction is done at the consumer - utility company level, as normally adopted in DSM system (Mohsenian-Rad et al. [14] used a similar idea using a game theory approach). The use of all EVs is considered in the solution representation used in our EAs (in Section 2 we give details on this) and we also use this in the evaluation of candidates solutions to *automatically* generate (nearly) optimal solutions.

To test this idea, we considered a dynamic scenario where the ultimate goal of each EV is to being able to complete a journey (e.g., the fuller the battery the better) while at the same time trying to charge them at different time slots (i.e., an action is made every 30 minutes over a period of 14 hours), whenever possible (e.g., reduction of load at the transformer level, reduction of PAR load demand). So, an automatic generated solution via our EA represents an (optimal) scheduling for each EV that indicates when each of them should be turned on to achieve both goals.

This paper is organised as follows. In the following we present the system configuration, the problem used and show in a nutshell how EAs work. In Section 3, we formally introduce the EA solution representation proposed in this work to automatically generate scheduling solutions based on load-shifting. Section 4 describes the experimental setup used in this study. In Section 5 we present and discuss the results obtained by our approach, and finally, Section 6 draws some conclusions and presents some future work.

## 2. SYSTEM CONFIGURATION, PROBLEM DESCRIPTION AND SOLUTION REPRESENTATION

### 2.1 Description of the Problem

We consider a scenario where a single source of energy is used among various household units and where each of them has an EV that charges over a period of time. The demand management of these EVs could happen by turning them “on” or “off” (e.g., [3, 8]) or by modulating the charge rate [16, 17]. In this work, we adopted the former. In our scenario, a decision can be made every 30 minutes. We try to maximise the state of charge of each EV at time of departure, whereas at the same time we try to avoid charging them all at the same time, whenever possible (e.g., reduction of the PAR demand load, reduction of transformer load).

To make this problem challenging, the period of time where each EV can be charged is the same for all them, which results in finding different charging schedules whenever possible (e.g., if two or more EVs needs to be turned on for the entire period of time at home, then it is not possible to find different charging schedules).

### 2.2 Automatic Evolution of Solutions

As indicated previously, we are interested in *automatically* generating (evolving) potential solutions that can accomplish both goals: EVs being as fully charged as possible whereas at the same time trying to avoid charging them at the same time, whenever possible. For this purpose, we use an evolutionary algorithm (EA). These algorithms are strongly influenced by natural selection, where the idea is to “evolve” potential solutions to automatically generate (nearly) optimal solutions.

The evolutionary process includes the initialisation of the population  $P$  at generation  $g = 0$ . The population consists of a number of individuals which represent potential solutions to a specific problem, ( $P(g)$ ). At each iteration or generation ( $g$ ), every individual within the population is evaluated using a *fitness function* that determines their fitness (i.e., how good or bad an individual is). Then a selection takes place to pick the fittest individuals from the population. Some of the selected individuals will be modified by genetic operators and the new population  $P$  at generation  $g + 1$  is created. The process stops when some halting condition is satisfied. Details on how these stochastic optimisation algorithms work can be found in [2, 7].

## 3. EVOLUTIONARY ALGORITHMS SOLUTION REPRESENTATION BASED ON LOAD SHIFTING

Two key elements that are necessary for successfully using an EA are: (a) the representation of individuals, and (b) how they are evaluated via a fitness function, so that automatic solutions can be generated. We now proceed to formally define both of them for the problem we briefly introduced (details of the experimental setup are presented in Section 4).

Let  $N$  denote the number of EVs, where the number of users is denoted by  $U = N$ . For each EV  $n \in N$ , we define the solution representation of  $n$  as

$$ind_n \triangleq [a_{t_i n}, \dots, a_{t_f n}]$$

where  $a_t$  denotes an action taken, at 30 minute granularity, within the range of  $t_i$  (initial time) and  $t_f$  (final time). Thus, the representation of a solution (individual) is given by

$$ind \triangleq ind_n \cdot N \quad (1)$$

Because we want to ensure that all the EVs can complete an entire journey at time of departure (e.g., being as fully charged as possible), while at the same time trying to balance energy usage by shifting energy consumption whenever possible. It is necessary to capture these in the evaluation of the potential solutions via the fitness function.

First, we start given a formal definition on the first target: trying to ensure that all the EVs  $n \in N$  are as fully charged as possible at time of departure. So, we compute the number of time slots that are automatically assigned to each EV ( $ind_n$ )

$$SoC_{ind_n} \triangleq \sum_{ind_n t=t_i}^{t_f} a_t \forall n \in N \quad (2)$$

While Equation 2 calculates the number of slots assigned to an EV ( $ind_n$ ) without considering how much energy it

needs, it is very unlikely that it will find the optimum solution for the first objective (e.g., optimum charging schedule of an EV). Let  $R$  be the number of time slots needed by  $ind_n$  to be as fully charged as possible, we now can compute the *optimal* number of slots needed to accomplish this. Thus,

$$OP\_SoC_{ind_n} \triangleq \begin{cases} \frac{SoC_{ind_n}}{R_{ind_n}}, & \text{if } SoC_{ind_n} < R_{ind_n} \\ \frac{SoC_{ind_n} - R_{ind_n}}{R_{ind_n}}, & \text{if } SoC_{ind_n} > R_{ind_n} \\ k & \text{otherwise} \end{cases} \quad (3)$$

where  $k = 1$  is a constant value which indicates that a perfect match has been generated between the necessary time slots to (fully) charge an EV and the assigned time slots. Equation 3 computes the number of times slots needed to optimally charge an EV. We now proceed to calculate the optimal number of time slots needed for all the EVs to be as fully charged as possible. This is computed as follows

$$OP\_SoC_{ind} \triangleq \sum_{n=1}^N OP\_SoC_{ind_n} \quad (4)$$

Equation 4 captures the first goal (i.e., trying to charge as much as possible all  $N$  EVs) without considering what all of them EVs are doing at time  $t$  which, in consequence, could imply in generating automatic solutions that represent charging all  $N$  at the same time. So, we need to consider the number of occurrences  $O$  at a particular time of the day.

For each customer  $n \in N$ , let  $a_n^t$  denote the action taken at time slot  $t \in T \triangleq \{t_i, \dots, t_f\}$ . The total number of EVs being charged at each time slot  $t \in T$  of the day can be computed as

$$O_t \triangleq \sum_{n \in N} \sum_{ind \in ind_n} a_n^t, \text{ if } a_n^t = 1 \quad (5)$$

Equation 5 computes the occurrences of EVs being charged at a particular time. Because we prioritise the charge of EVs over the number of EVs being charged at a particular time, we need to re-formulate the latter. Thus, we regarded an optimal occurrence at time  $t$  as

$$OP\_O_t \triangleq \begin{cases} \frac{k}{N}, & \text{if } O_t = 1 \\ -k, & \text{if } O_t = 0 \end{cases} \quad (6)$$

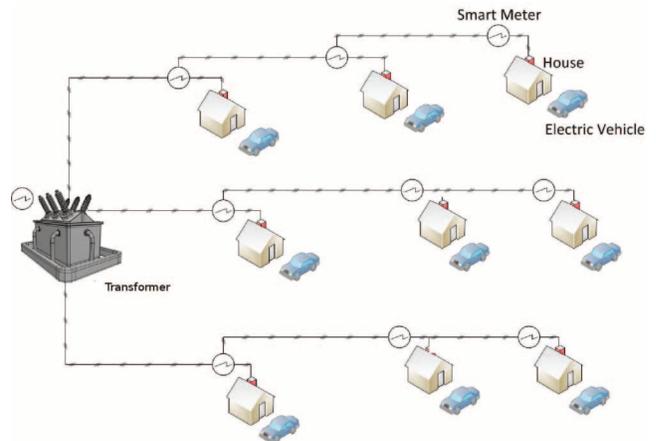
We automatically evaluate potential solutions (individuals) by combining our first and second objective, formally defined in Equations 4 and 6, respectively. Thus, we have that the fitness function is given by

$$f(ind) \triangleq \text{maximise } OP\_SoC_{ind} + OP\_O_t \quad (7)$$

## 4. EXPERIMENTAL SETUP

### 4.1 Smart Grid Scenario

As mentioned previously, the goals are: (a) to guarantee that each EV can complete a journey (e.g., each EV can be charged as much as possible at time of departure), and (b) to avoid charging them at the same time, whenever possible. In Section 5, we show how both are met by our proposed



**Figure 1: An schematic view of the problem used in this work with two main goals: (a) that each EV is as fully charges as possible at time of departure and (b) reduction of load transformer achieved by avoiding charging more than one EV at a particular time, whenever possible.**

approach by showing (a) the state of charge at time of departure, and (b) by showing the averaged transformer load for the period of time where the EVs can be charged and the peak-of-ratio (PAR) in load demand, respectively. The PAR in load demand is calculated as the highest load across all users for a period of time (e.g., time available to charge EVs) over the average load for the same period of time. Details of this calculation can be found in [13]. Thus, a reduction of PAR is generally preferred.

In our considered benchmark smart grid system there are  $U = 9$  users, each with an EV ( $N = U$ ). Each EV  $n \in N$  can only be charged at home. As explained in Section 2, all  $N$  EVs could be charged in the same period of time (i.e.,  $t_i = 18:00$ ,  $t_f = 7:30$ ). An schematic view of the problem used in this problem can be seen in Figure 1

We simulated a dynamic scenario, where every day, the initial state of charge ( $SoC$ ) of each EV at time of arrival ( $t_i$ ) at home varies (recall that EVs can only be charged once they are at home). To be able to easily visualise how good an automatic solution generated by our stochastic optimisation EA algorithm is, we defined the  $SoC$  within the range of  $SoC = [85, 95]$ . By using these values, an optimal solution would be that an EV can be fully charged at time of departure. This would be fairly easy to solve if it was not by the presence of our second goal: avoiding charging more than one EV at the same time whenever possible. Table 1 summarises the parameters used to simulate our grid scenario. We ran our simulations for one month of simulated time, using the well-developed GridLab-D [4] (version 2.3). This is an electrical, open source, grid simulator developed by the US Department of Energy.

### 4.2 Evolutionary-Based Algorithm

The generation of automatic solutions were obtained using an evolutionary-based optimisation stochastic method with fitness proportionate selection and two-crossover point, and bit-flip mutation, run for 300 generations. To obtain meaningful results, we performed 28 independent runs (i.e., one

**Table 1: Summary of Parameters Used for our Smart Grid System**

Parameter	Value
Number of EVs ( $N$ )	9
Initial time and latest time to charge	$t_i = 18:00, t_f = 7:30$
Frequency of making a decision	30 minutes
Number of times slots $T$	28
State of Charge ( $SoC$ ) at $t_i$	[85, 95]
Type of Charging	Low

**Table 2: Summary of Parameters for our Evolutionary-Based System**

Parameter	Value
Length of the individual	$T \cdot N$ (see Table 1)
Population size	300
Generations	300
Mutation rate (per bit)	0.001
Two-point crossover rate	0.7
Independent runs	28

run for each day of the month). Runs were stopped when the maximum number of generations was reached. The parameters we have used are summarised in Table 2.

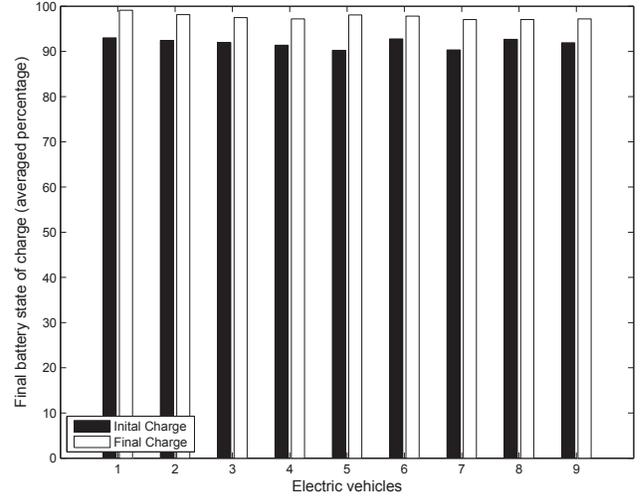
To compare the results achieved by our proposed approach, we used a deterministic approach (default implemented in GridLab-D), denominated in this work as “greedy”, where all the EVs start charging as soon as they reach home. To make a fair comparison of both approaches, each EV in these two methods, used exactly the same randomly generated values for each day of the month (i.e., initial  $SoC$  at time of arrival  $t_i$  at home).

## 5. RESULTS AND ANALYSIS

### 5.1 Final Battery State of Charge

Let us start by analysing the results on the first objective: trying to maximise the  $SoC$  of all the EVs at time of departure. Figure 2 shows the average  $SoC$  over a period of 28 days, at time of arrival at home (initial charge) and their corresponding  $SoC$  at time of departure (final charge). As indicated in Section 4, we randomly initialised the initial  $SoC$  between  $SoC = [85, 95]$ , so that it could be easily identified whenever an optimal solution, for this first objective, (i.e., an EV being fully charged – 100%) was automatically generated by the EA via the proposed fitness function (left-hand side of Equation 7). Our proposed approach behaves fairly well finding, almost, an optimal solution for all the EVs (i.e.,  $SoC > 98\%$ ). It should be noticed that the final  $SoC$  achieved by the greedy approach (i.e., EVs charging as soon as they reach home) is not reported because all the EVs are able to being fully charged at time of departure

We now take a more fine-grained view by analysing how the charge happens for an EV for the 28 days used in this study and using the proposed approach, explained in Section 3. Due to space limitations, we only report the results on  $EV_1$  shown in Figure 3. This shows that regardless the initial state of charge, our proposed approach is able to significantly charge the EV. For example, for days 13, 18 and 28 and the initial  $SoC$  being set no greater than 87%, the



**Figure 2: State of charge ( $SoC$ ), averaged over a period of one month, at time of arrival for each EV (indicated in black bars) and their corresponding  $SoC$  at time of departure (indicated in white bars) achieved by our stochastic bio-inspired optimisation process.**

EA is able to fully charge the EV. We can also see that for those cases where the EV was not being able to be fully charged, the state of charge at the time of departure was set at  $SoC > 94.5\%$  (four days – days 2, 9, 10, 21) and  $SoC > 98\%$  (three days – days 3, 7 and 8). In summary, Figure 3 shows that for most of the days, the proposed approach is able to find an optimal solution, for the first objective (i.e.,  $SoC$  being fully charged at time of departure), and for those non-optimal solutions, the  $SoC$  is at least 94.5% at time of departure.

### 5.2 Transformer Load and Peak-To-Average Ratio

To fully appreciate the results achieved by our proposed approach, we need to analyse the impact of this automatic scheduling on the transformer load. Figure 4 shows the transformer load when the EVs are being charged. Due to the nature of the greedy approach (i.e., EVs starting charging as soon as they arrive at home), it is natural to imagine how this approach will achieve the highest peak in electricity at the initial time of charging ( $t_i = 18:00$ ) and decreases linearly thereafter.

In contrast, our proposed approach, is able to balance energy usage during the period of time where the EVs can be charged ( $t_i = 18:00, t_f = 7:30$ ), as expressed in the ‘x-axis’ of Figure 4) while at the same time all the EVs are being charged as much as possible as discussed in the previous paragraphs (see Figure 2). An interesting trend that can be observed by our approach is how there is a slightly more energy consumption around the initial time of charging (e.g.,  $t < 20:00$ ) compared to the remaining time slots that can be used to charge the EVs. This is expected since the evaluation function (fitness function) expressed in Equation 7 gives a higher importance (left-hand side of the equation) to EVs being as fully charged as possible compared to the energy usage balance (Equation 6).

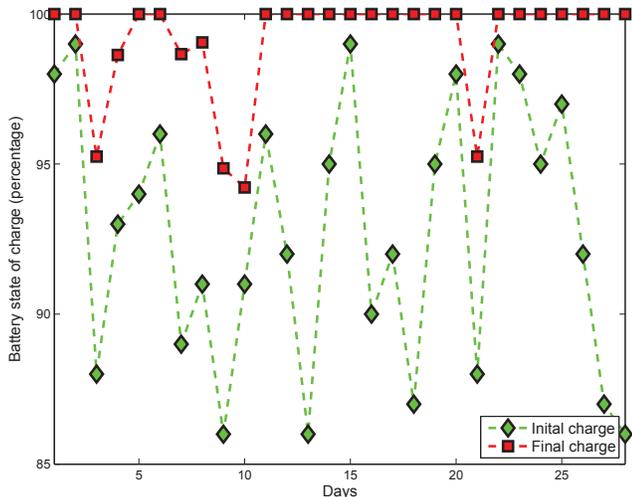


Figure 3: Battery state of charge ( $SoC$ ) represented in terms of percentage for  $EV_1$ . The initial  $SoC$ , indicated in green diamond marks, for each of the 28 days, is randomly initialised within the range  $[85, 95]$  (see Table 1 for details). The proposed approach is able to find the optimal solution, for the first objective (i.e., final  $SoC$ , indicated in red square marks,  $SoC = 100$ ) most of time.

Finally, let us analyse the impact of our approach *vs.* the greedy approach in the PAR load demand, shown in Figure 5. As indicated in Section 4, the PAR is calculated by the maximum load demand for a period of time over the average load demand. Thus, a reduction in PAR is preferred. We can see how our proposed approach (indicated in white bars in Figure 5) is able to reduce the PAR load demand compared to the greedy approach (indicated in black bars in Figure 5) most of the time (i.e., 25 out of 28 days). This is to be expected. That is, we know that the PAR load demand is calculated by considering the maximum load demand of a period of time. We know, from Figure 4, that the maximum load is considerable reduced by our proposed approach hence we should expect to see a reduction of PAR, in general.

## 6. CONCLUSIONS AND FUTURE WORK

Demand-Side Management (DSM) systems play an important role in Smart Grids. Their importance can be understood by considering the new challenges that are continuously presented to the grid, for example, electric appliances that could double the average household (e.g., electric vehicles). The correct design of a DSM results in efficiently use the available energy without the necessity of installing new electricity infrastructure.

In the specialised literature, there are techniques adopted by DSM programs. Perhaps, the most popular techniques are those inspired on smart pricing, such as critical-peak pricing, real-time pricing. Briefly, the idea is to incentivise end-consumers to shift energy consumption to hours when the electricity price is low to reduce both electricity costs and energy-load consumption.

Motivated by these techniques we considered, instead of

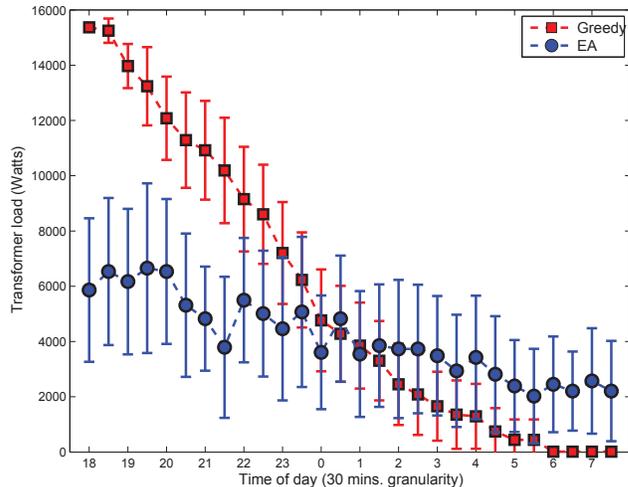


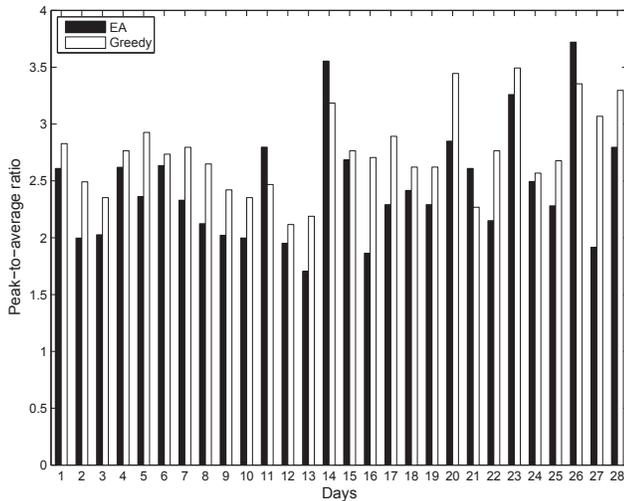
Figure 4: Transformer load, averaged for 28 days, for 9 EVs with different initial state of charge  $SoC$ . All 9 EVs could be charged from 18:00 until 7:30. Red squares shows the transformer load for the greedy approach (i.e., EVs start charging as soon as the reach home) whereas blue circles show the load using our proposed approach.

using a price signal to shift load consumption, what all the EVs do at a particular time to achieve both: (a) for an EV to be as charged as much as possible at time of departure, and (b) to avoid turning them on at the same time whenever possible. We then adopted an evolutionary-based stochastic optimisation algorithm to automatically generate scheduling solutions for each of the EVs for the period of 28-days used in this study.

The results achieved by our proposed approach are highly encouraging. That is, we showed how it is possible to almost fully charged the EVs at the time of departure while at the same time we are able to reduce the load consumption at the transformer level which is reflected in the reduction of peak-to-average ratio of load demand, which is a desired feature in smart grids.

We are planning to extend this work considerable. For example, we decided to use a discrete-valued representation in our EA that indicates the (nearly) optimal scheduling that each EV should follow for a period of time to be as much charged as possible while reducing the transformer load. This works well for a “small” scenario as the one we used in our study (i.e., 9 end-users each with an EV). One natural extension is to consider having a real-valued presentation that could indicate, for instance, how much energy is needed by each EV at a particular time. This is particular important to consider when one deals with a large scenario (e.g., dozens/hundreds of EVs).

One also needs to consider the adoption of a more “natural” form of multi-objective (MO) optimisation. The specialised MO literature is full of approaches that can nicely be integrated with our approach and we will also incorporate it in our EA.



**Figure 5: Peak-to-average ratio (PAR) load demand achieved by both our proposed approach (black bars) vs. the greedy approach (white bars). Our approach is able to reduce the PAR in 25 out of 28 days. A lower PAR is preferred.**

## Acknowledgments

This research was supported by Science Foundation Ireland (SFI) under the Principal Investigator research program 10/IN.1/I2980 “Self-organizing Architectures for Autonomic Management of Smart Cities” and by SFI grant 10/CE/I1855 to Lero - the Irish Software Engineering Research Centre (www.lero.ie).

## 7. REFERENCES

- [1] *Pacific Northwest GridWise Testbed Demonstration Projects, Part I. Olympic Peninsula Project*, October 2007.
- [2] T. Bäck, D. B. Fogel, and Z. Michalewicz, editors. *Evolutionary Computation 1: Basic Algorithms and Operators*. IOP Publishing Ltd., Bristol, UK, 1999.
- [3] A. Brooks, E. Lu, D. Reicher, C. Spirakis, and B. Wehl. Demand dispatch: Using real-time control of demand to help balance generation and load. *IEEE Power & Energy Magazine*, 8:20 – 29, 2010.
- [4] D. Chassin, K. Schneider, and C. Gerkensmeyer. Gridlab-d: An open-source power systems modeling and simulation environment. In *Transmission and Distribution Conference and Exposition, 2008. IEEE/PES*, pages 1–5, 2008.
- [5] A. Collin, J. Acosta, I. Hernando-Gil, and S. Djokic. An 11 kv steady state residential aggregate load model. part 2: Microgeneration and demand-side management. In *PowerTech, 2011 IEEE Trondheim*, pages 1–8, 2011.
- [6] A. Conejo, J. Morales, and L. Baringo. Real-time demand response model. *Smart Grid, IEEE Transactions on*, 1(3):236 –242, dec. 2010.
- [7] A. E. Eiben and J. E. Smith. *Introduction to Evolutionary Computing*. Springer Verlag, 2003.
- [8] E. Galvan, C. Harris, I. Dusparic, S. Clarke, and V. Cahill. Reducing electricity costs in a dynamic pricing environment. In *Proc. Third IEEE International Conference on Smart Grid Communications (SmartGridComm)*, pages 169 – 174, Tainan, Taiwan, november 2012. IEEE Press.
- [9] C. Gellings. The concept of demand-side management for electric utilities. *Proceedings of the IEEE*, 73(10):1468–1470, 1985.
- [10] C. Gellings and W. M. Smith. Integrating demand-side management into utility planning. *Proceedings of the IEEE*, 77(6):908–918, 1989.
- [11] C. Harris, I. Dusparic, E. Galván-López, A. Marinescu, V. Cahill, and S. Clarke. Set point control for charging of electric vehicles on the distribution network. In *2014 IEEE Power & Energy Society Innovative Smart Grid Technologies Conference (ISGT)*, Washington, D.C., USA, February 2014. IEEE Press.
- [12] G. M. Masters. *Renewable and Efficient Electric Power Systems*. Wiley-Interscience, 2004.
- [13] A. Mohsenian-Rad, V. 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. *Smart Grid, IEEE Transactions on*, 1(3):320 –331, dec. 2010.
- [14] A.-H. Mohsenian-Rad and A. Leon-Garcia. Optimal residential load control with price prediction in real-time electricity pricing environments. *Smart Grid, IEEE Transactions on*, 1(2):120 –133, sept. 2010.
- [15] P. Palensky and D. Dietrich. Demand side management: Demand response, intelligent energy systems, and smart loads. *Industrial Informatics, IEEE Transactions on*, 7(3):381–388, Aug.
- [16] P. Richardson, D. Flynn, and A. Keane. Optimal charging of electric vehicles in low-voltage distribution systems. *Power Systems, IEEE Transactions on*, 27(1):268 –279, feb. 2012.
- [17] S. Studli, E. Crisostomi, R. Middleton, and R. Shorten. A flexible distributed framework for realising electric and plug-in hybrid vehicle charging policies. *International Journal of Control*, 85(8):1130–1145, 2012.
- [18] A. Taylor, I. Dusparic, E. Galván-López, S. Clarke, and V. Cahill. Transfer learning in multi-agent systems through parallel transfer. In *Theoretically Grounded Transfer Learning at the 30th International Conference on Machine Learning (ICML)*. Omnipress, 2013.
- [19] A. Taylor, E. Galván-López, S. Clarke, and V. Cahill. Management and control of energy usage and price using participatory sensing data. In *The Third International Workshop on Agent Technologies for Energy Systems at Eleventh International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, pages 111 – 119, 2012.
- [20] D. Westermann and A. John. Demand matching wind power generation with wide-area measurement and demand-side management. *Energy Conversion, IEEE Transactions on*, 22(1):145–149, 2007.