Multi-agent residential demand response based on load forecasting

Ivana Dusparic, Colin Harris, Andrei Marinescu, Vinny Cahill, Siobhán Clarke
School of Computer Science and Statistics
Trinity College Dublin

***Abstract*—Improving the efficiency of the smart grid, and in particular efficient integration of energy from renewable sources, is the key to sustainability of electricity provision. In order to optimize energy usage, efficient demand response mechanisms are needed to shift energy usage to periods of low demand, or to periods of high availability of renewable energy. In this paper we propose a multi-agent approach that uses load forecasting for residential demand response. Electrical devices in a household are controlled by reinforcement learning agents which, using the information on current electricity load and load prediction for the next 24 hours, learn how to meet their electricity needs while ensuring that the overall demand stays within the available transformer limits. Simulations are performed in a small neighbourhood consisting of 9 homes each with an agent-controlled electric vehicle. Performance of agents with 24-hour load prediction is compared to the performance of those with current load information only and those which do not have any load information.**

I. Introduction

Utilizing renewable electricity sources, such as wind power and solar energy, is key to the sustainability of the electric grid. As an example, the European Union is committed to meeting a target of 20% of its energy coming from renewable sources by 2020, and Ireland aims to achieve that goal by having 40% of its electricity come from renewable sources. However, renewable sources have much greater uncertainty than traditional ones and their availability can change within hours or even minutes. Storing energy produced by renewable sources is one of the ways to address this issue, however, encouraging greater energy usage at times of greater renewable energy availability could reduce cost and energy losses associated with energy storage. In order to enable this, accurate prediction of users' demand is required, which, coupled with the information on availability of renewable energy, can be used to dynamically set energy prices reflecting the relationship between current supply and demand. To take advantage of such highly dynamic pricing, electrical devices need to be able to autonomously make decisions about when to turn on or off, while meeting their operating goals. In this paper we focus on shifting demand, i.e., demand response (DR), by enabling devices to learn their operating decisions based on prices informed by current load as well as load prediction for the next 24 hours. Availability of renewable sources can easily be integrated into such a model by further decreasing or increasing the current or predicted price to reflect not just overall load but availability of renewable resources as well. Each device is controlled by a reinforcement learning (RL) agent [1], which learns how to meet its goal (e.g., target battery charge for electric vehicles, desired temperature for heating/cooling systems etc), by the target time at a minimum possible price. Load on the transformer is predicted for the day ahead based on historical load data and weather forecast. Current electricity price and price for the next 24 hours is sent to RL agents and is proportional to the predicted load in the system, therefore encouraging devices to use the electricity during predicted off-peak price times. Agents also have the information on the current load on the transformer, to prevent transformer overload.

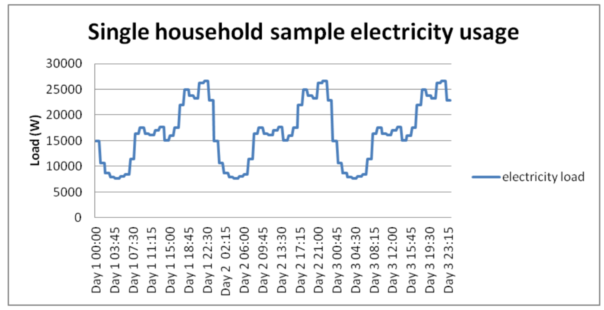

Figure 1. Sample electricity usage of a single household

The rest of this paper is organized as follows. We first briefly summarize existing work in the demand response area in Section II and introduce reinforcement learning which is the basis of our approach. In Section III we present the design of the proposed multi-agent system. Section IV describes evaluation of our approach and presents results while Section V concludes the paper.

II. Related Work

Consumers do not use energy evenly throughout the day. For example, in household consumption, there is a morning peak after inhabitants wake up and are getting ready for their daily activities, and a large evening peak when they arrive home, which is followed by very low electricity usage during the night. Figure 1 shows sample electricity usage for a household over the period of 3 days illustrating this pattern (data is taken from a smart meter trial in Ireland). Demand response is a modification of the consumers' electricity consumption with respect to their expected consumption [2]. For example, if a high usage peak is predicted during the

978-1-4673-4630-6/13/$31.00 ©2013 IEEE

early evening period when residential consumers arrive home, demand response will aim to reduce that peak by encouraging consumers to postpone their non-essential electricity use (so called peak clipping [3]). Similarly, if a low consumption period is predicted, demand response techniques aim to increase the usage at those times, by encouraging consumers to shift their electricity use from earlier or later peak periods to that low-usage period (so called valley-filling [3]). Similarly, in order to encourage greater use of renewable resources, demand response techniques aim to reduce the usage during low renewables availability, and increase the usage during high availability.

A. Current demand response techniques

Effective DR depends critically on price and load forecasting as well as on demand management [4]. Various techniques have been implemented on the consumer side to shift the demand from high demand periods to those of lower demand and have been evaluated on scenarios of differing scales. A lot of techniques focus on implementation within a single household, such as backpropagation neural networks (e.g., [5]), expert systems theory (e.g., [6]), dynamic programming (e.g., [7]), linear programming (e.g., [8]), and RL (e.g., [9], [10]). The issue with single household implementations is that they do not take into account an aggregate effect of all customers shifting the demand to low price/low demand periods, resulting in a sudden increase in demand and overloads on the transformer. DR approaches can also be implemented on a microgrid/community level, e.g., by use of RL in [11], or evolutionary computing in [3]. These approaches are centralized and as such might not be scalable to large communities. In this approach, local devices/households are not in charge of their own schedules but schedules are received from the central scheduling component. DR can also be implemented on the load generation side by scheduling the generation so that it meets the required demand as well as scheduling the consumer devices (e.g., [12]), however such approaches cannot be used with renewable sources as, for example, wind or solar energy, cannot be scheduled. Larger scale multi-agent systems have also been implemented, mainly to coordinate charging of electric vehicles (EVs), such as those in [13] and [14], however, these approaches take into account only current load/price. Our approach will investigate enriching multi-agent learning-based approaches with the information on the predicted 24-hour load (i.e., price which is proportional to predicted load) in order to improve scheduling of devices within the community as well as optimizing the price within each household.

B. Reinforcement learning

In this paper, we implement demand response using RL. RL is a learning technique based on trial and error that has been researched and applied in control theory, machine learning and artificial intelligence problems, as well as non-computer science domains such as psychology. It is considered particularly suitable for implementation of self-organizing

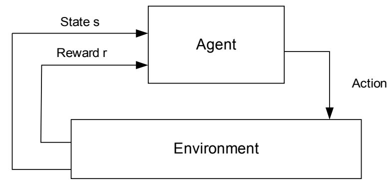

Figure 2. Reinforcement learning process

optimization behaviours in large-scale systems, as it does not require a predefined model of the environment, which, due to the scale and complexity of such systems, is time-consuming and complex to construct [15]. Q-learning [16] is one such model-free RL algorithm. Using Q-learning, an agent learns to associate actions with the expected long-term reward of taking that action in a particular state. An agent tries an action in a particular state and observes the next state and the reward it got for getting to/being in that state. By repeatedly visiting states and trying out actions, it learns which action is the best to take in which state, i.e., has the highest value in terms of long term reward.

Basic RL has been extended both to multi-policy and multi-agent techniques, to address different requirements of application domains. In this paper we use W-learning [17] to address multiple policies on our agents, as the algorithm has proven scalable in other large-scale environments (e.g., urban traffic control [18]). In W-learning each policy is implemented as a separate Q-Learning process with its own state space. Using W-Learning, an agent learns, for each of the states of each of its policies, what happens, in terms of the reward received, if the action nominated by that policy is not obeyed. This is captured in a so-called W-value; the higher the W-value, the more important it is for that policy to have its suggested action executed. An agent then executes the action nominated by the policy which is going to suffer the highest loss if its nominated action is not executed, than any other policy would suffer if their nominated action was not executed, i.e., the one with the highest W-value.

W-learning has also been extended to a multi-agent technique, Distributed W-learning [19], which enables collaboration between heterogeneous agents, and as such could be used in further extensions of our scenario implementations to enable collaboration. We envisage that this work will be a part of a larger smart grid simulation implemented as a large-scale multi-agent system, where devices, households, transformers, generators, and suppliers are all represented by intelligent agents optimizing their own behavior but also cooperating to maintain the optimal performance of the overall system.

III. Agent and Experiment Design

For the purpose of the experiments in this paper we limit the scope of the system to 9 devices, specifically electric vehicles (EVs), operating within a neighbourhood covered by a single transformer, as pictured in Figure 3. At a transformer level, the pricing agent determines the current and future price of the electricity for the neighbourhood based on the current load in

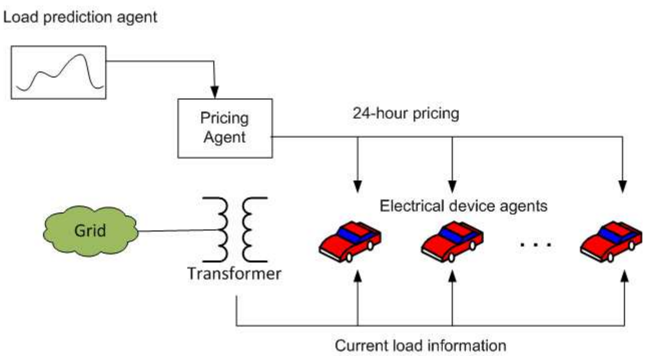

Figure 3. Multi-agent system architecture

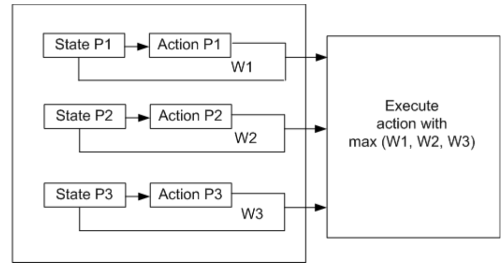

Figure 4. W-learning process

the system as well as load prediction for the next 24 hours as estimated by the prediction agent. The prediction agent combines various forecasting advantages of the methods showcased in [20], where auto-regressive techniques, artificial and wavelet neural networks, and fuzzy logic are all considered in a further hybrid approach to day ahead demand forecasting. The current and predicted price set by the pricing agent is directly proportional to the load in the system.

Device agents must learn to minimize the cost of their operation, while ensuring not to overload the transformer, and ensuring they meet their charging targets. Each EV is capable of implementing 3 policies, which are turned on and off to implement different scenarios:

Policy 1: This policy ensures that a vehicle achieves the desired minimum battery charge. It has information about the minimum required charge and current battery charge, and is rewarded as follows: 500 points for achieving minimum battery charge required to complete the daily journey, 500 points if the battery charge at a given time step is higher than the one at the previous time step, and negative reward of -500 points if current charge is not greater than the one at the previous time step.

Policy 2: This policy ensures that the overall load at the transformer supplying the whole community is kept within assigned limits. The limit has been set so as to discourage vehicles charging during the peak periods of the base load; if multiple EVs charge during the peak base load period this limit will be exceeded. The policy has information only on current load at the transformer (overall for the whole 9 houses community) and is rewarded as follows: 500 points if the load is under the set limit, neutral reward of 0 is the load is very near the limit, and negative reward of -500 points if the load has been exceeded.

Policy 3: This policy ensures that the vehicles are charged during the lowest possible load periods available to them (where lowest load corresponds to lowest price period). It has the information on the current load as well as on the load prediction for the next 24 hours. At each time step, the policy classifies the current load as low, medium, or high, relative to the predicted load for the next 24 hours. (More precisely, the policy does not necessarily look at the whole next 24 hours, but only at the period during which the vehicle can charge. If the vehicle is scheduled to depart in the next 10 hours, it will only look at the next 10-hour window). Vehicles are rewarded 500 points if they charge during the low load period, 250 points if they charge during medium load period, and get a negative reward of -50 points if they charge during high load period. The difference between this policy and Policy 2 is that Policy 2 punishes charging at the high load no matter what, and punishes it with a very high negative reward of -500; Policy 3 accounts for the situation that vehicles sometimes might need to charge during the peak load as well and it punishes that behaviour with only -50 points, while highly rewarding charging at medium and even more so at low load. Therefore, agents will aim to maximize their reward by primarily charging during "low" load periods, but will also aim to avoid the negative reward of -500 by making sure that they do achieve the required charge, even if that required charge has to be achieved during "high" load, which is punished by -50 points. Also, the current load in this policy is expressed as a relative value (low, medium, high) rather than absolute load as in Policy 2. This gives agent information on the predicted load so that, for example, if the current load is "high", the agent knows that high is relative to the remainder of the available charging period, and can therefore wait for the "medium" or "low" load periods in order to charge.

Note that the load information sent to agents in all policies can be true load (or actual current electricity price) or just a pricing signal designed to shift demand to the periods of higher availability of renewable resources. This design enables the same agent implementation to be used regardless of whether devices are connected to the main grid, operate within a local microgrid, or if the price signal is sent from a renewable source within the household, e.g., a photovoltaic panel.

Policies 1, 2, and 3 are activated or deactivated in different combinations to create several evaluation scenarios to investigate influence of different information sets of demand response (described in more detail in the next section). At each time step, each of the active policies on each vehicle suggests an action (whether an EV should charge or not charge at this time step), and using W-learning (as described in Section II) an agent decides which action to execute. This process is pictured in Figure 4.

IV. Evaluation

This section presents results from simulations performed on 9 households, each with a base load and RL-controlled EV. Experiments were run for 55 simulated days, and split in 2 parts: learning or exploration period, which lasts for approximately 80% of the overall experiment duration, and exploitation period, during which agents were only selecting the actions they learnt to be good, which lasts for the remainder of the experiment duration. Results presented are from the exploitation period. Vehicles have a battery capacity of 30 kWh and charge at rate of approximately 1.4kW. The required daily mileage differs in different implementation scenarios and ranges from 50 miles (requiring about 35% of full battery charge) to 80 miles (requiring about 50% of full battery charge). Every 15 minutes each electric vehicle agents makes a decision whether to turn the charging on or off for the next 15 minute period. Base load in each household ranges from 0.8 kW to 3 kW, based on time of the day, and is taken from the data recorded in a smart meter trial performed in Ireland in 2009-2010 [21].

Simulations are performed in GridLab-D [22], and agents are implemented using the DWL library [23]. Scenarios implemented are as follows:

1) 9 EVs without intelligent control, i.e., charging when they arrive home until fully charged. This scenario is used as a baseline to which other approaches are compared.
2) 9 EVs each implementing Policy 1 and Policy 2 simultaneously (i.e., aiming to achieve their desired battery charge while making sure not to go over designated maximum load on the transformer)
3) a single EV implementing Policy 1 and Policy 3 simultaneously (i.e., aiming to achieve its desired battery charge while charging during the lowest available load period, in order to minimize its overall charging cost). In this scenario mileage required is varied from 50 to 80 miles in order to vary the battery charge required and hence the charging duration.
4) 9 EVs each implementing Policy 1 and Policy 3 simultaneously (i.e., aiming to achieve their desired battery charge while charging during the lowest available load period, in order to minimize their overall charging cost). In this scenario, as in Scenario 3, mileage required by all vehicles is varied from 50 to 80 miles in order to vary the battery charge required and hence the charging duration.
5) 9 EVs each implementing Policy 1, Policy 2 and Policy 3 simultaneously (i.e., aiming to achieve their desired battery charge while charging during the lowest available load period, in order to minimize their overall charging cost, but also making sure not to go over the designated maximum load on the transformer).

A. Experimental results

The results presented show our RL approach is suitable for shifting the demand from peak usage periods to off-peak ones. Figure 5 shows the load pattern in a single household implementing Policy 1 and Policy 2 over 55 days (~5000 time-steps). During the exploration period (~4300 time-steps in the graph) the load reaches 4.5 kW and goes as low as less than 1 kW, while in the exploitation stage (time-steps 4300 - 5000), when the agents have learnt to respond to load information, the maximum load is just over 3 kW, and minimum load is 1.5 kW. Therefore, the load pattern has been significantly smoothed out: peak usage was reduced by ~33%, and off-peak usage was increased by ~50%.

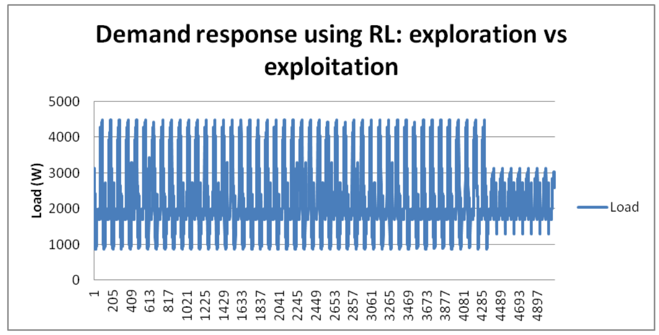

Figure 5. RL-based DR: load fluctuations during the exploration and exploitation stages

Scenario 1: In Scenario 1 EVs are not controlled by agents, and we use this scenario as a baseline for comparison of other approaches.

Scenario 2: Figure 6 compares the performance of Scenario 2 to the baseline Scenario 1. In Scenario 1, where EVs are not controlled by agents, they start charging immediately after arriving home at the peak energy demand time, increasing peak usage to over 40 kW (over 9 households). In Scenario 2, where agents have information on current load, they learn not to charge during high load periods (as they are heavily penalized for doing so by a reward of -500 points), therefore clipping up to ~25% of the overall demand in the neighbourhood from peak times. Maximum load during the peak period in Scenario 2 is 30 kW (over 9 households), while off-peak usage has been increased from 8 kW in Scenario 1 to 15-20 kW in Scenario 2 (so demand has successfully been shifted from the peak to off-peak period).

Scenario 3: Scenario 3 investigates the situations in which the off-peak period isn't long enough for the EVs to achieve their desired charge, so agents will need to charge during medium load and high load periods too. The mileage required to be completed by a vehicle on its daily journey is increased from 50 miles (mileage used in Scenario 2) to 70 and 80 miles. If the agent is learning only not to charge during the high load periods (i.e., it implements Policy 2), it would postpone all its charging until low load periods, achieving the charge required for a part of their journey but not a full charge required for 70 and 80 mile trips, as the off-peak period is not long enough. To enable agents to learn that in order to achieve the desired charge they also need to charge during high load periods (but that they should still fully utilize low

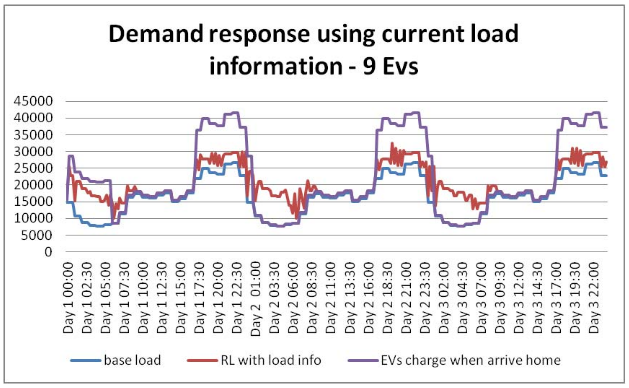

Figure 6. DR based on current load information - 9 EVs

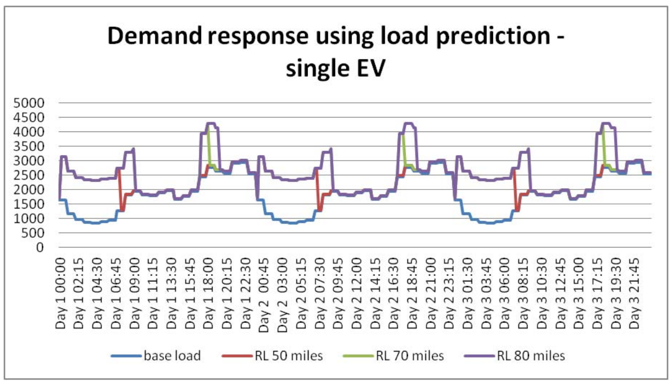

Figure 7. DR based on predicted load information - a single EV

load periods, if such periods are coming), we extended the agents' state space with information about predicted load for the next 24 hours (i.e., enabled them to implement Policy 3). In this scenario we performed the experiments only on a single agent. Results are presented in Figure 7. When a vehicle has only 50 miles to travel to work, the off-peak period is long enough for the vehicle to charge the battery to the desired charge. An agent learns to wait for this off-peak period, as it, based on current load being classified as "high", knows that "low" and "medium" periods are coming. Also, based on experiences from previous days (40+ days of learning during exploration period), it knows that this "low" load period is long enough to achieve the full charge. However, as the mileage required increases, agents learn that the off-peak "low" period is not long enough and learn to charge during the peak "high" periods too. Observe, however, that the off-peak period is still fully utilized, and that peak charging takes place only for as long as it is required to supplement off-peak charge in order to enable a vehicle to complete the required mileage. For example, to achieve the battery charge required for a 80 mile journey, the full duration of the low load period is utilized, but also a few hours during the peak period. For a 70 mile journey, the full duration of the low period is utilized, and only an hour of the high load period. An agent was able to achieve this balance as its punishment (i.e., negative reward) is higher for not achieving required battery charge (- 500 in

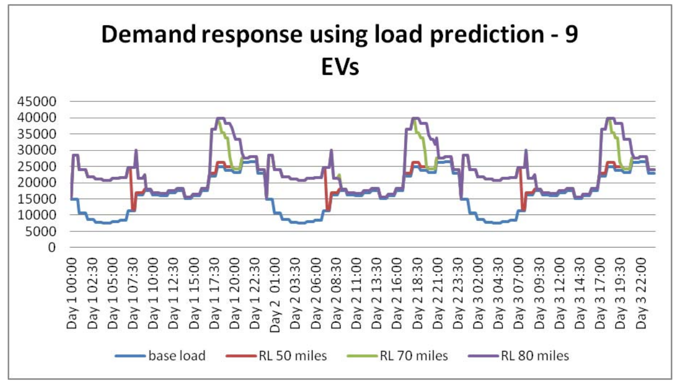

Figure 8. DR based on predicted load information - 9 EVs

Policy 1), than its punishment for charging during high load period (negative reward of -50 in Policy 3).

Scenario 4: Scenario 3 presented results of demand response using load prediction on a single agent. However, the issue arises when all EVs need to complete a large mileage, and all learn that they need to charge during peak times to achieve the desired battery charge (i.e., when numerous EVs implement Policy 3), significantly increasing the load at peak times. Figure 8 presents the results of Scenario 4, where all EVs are provided with predicted load information. All EVs learn to secure their "top-up" charge at the start of the peak period, to make sure that required charge is achieved, thereby increasing the early evening peak by over 50% from base load. As observed in Scenario 3, vehicles need to charge for a certain length of time during the peak period, however, as they only need to charge for about a third or a half of the peak period, rather than the whole duration of it, ideally charging should be spread out. For example, in the case of the 70 mile scenario, all vehicles charge at the start of the peak-period increasing the overall load by 15 kW from the base load for about half of the peak period, following which load drops only to the base load levels for the remainder of the peak period. Agents should be able to learn to spread out their charging, i.e., should also be implementing Policy 2, which puts the maximum limit on the transformer and making sure overall load does not exceed it. We re-introduce Policy 2 in Scenario 5.

Scenario 5: In Scenario 5, we enable EVs to learn when to charge both based on the information on current load (as in Scenario 2), and based on predicted load information (as in Scenario 4). Therefore, in this scenario, agents implement all 3 policies: Policy 1, Policy 2, and Policy 3. Results are graphed in Figure 9. By re-introducing Policy 2, i.e., setting a maximum limit at the transformer and punishing agents for exceeding it, we hoped to enable agents to learn that, even if they have to charge during the peak times, to spread out that charging so that not all of them charge at the same time. This has been achieved to some extent, the maximum load on days 2 and 3 pictured in Figure 9 is 35 kW, comparing to 40 kW in Scenario 4, however, on day 1, maximum load in this scenario reaches 40 kW as well. The reason for this is

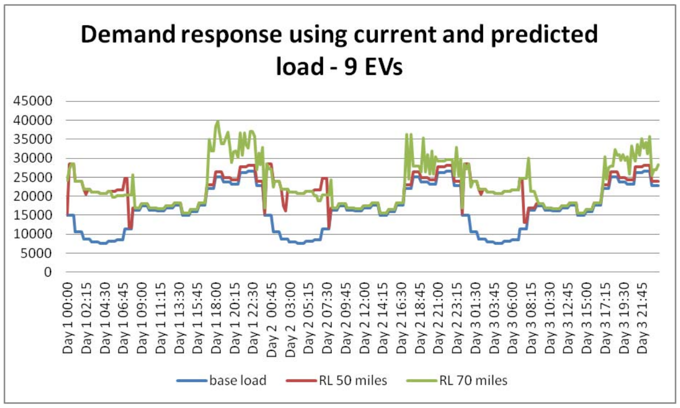

Figure 9. DR based on current and predicted load information - 9 EVs

that there are 9 agents affecting the environment, and each is only aware of its own actions. However, it receives the information on the current load from the transformer, which incorporates the load created by all 9 agents, and gets rewarded or punished based on that overall load on the transformer, i.e., for the actions of all other agents as well. We believe further improvements in decreasing peak load can only be achieved through introduction of the collaboration between agents, which will be the subject of follow up work.

B. Evaluation Summary

Results presented in this paper show that RL is a suitable technique for residential demand response. EVs are controlled by RL agents and given different sets of information to evaluate how they can be influenced to minimize their charging price, achieve desired battery charge and keep transformer load under designated maximum load, by shifting their charging from high load to low load periods. Results show that giving agents information on the current load is able to reduce the peak-charging by ~33%, by punishing them for charging during high load. However, in order to achieve desired battery charge agents sometimes need to charge during the high load as well, and are therefore provided with the information on 24-hour load prediction as well. Agents can then learn to charge only for the minimum required amount of time during the high load, and wait for low load periods for the majority of their charging period. As agents also get rewards for the overall load at the transformer, which is influenced by all agents in the system, we believe that further improvements on the results obtained here can be achieved by introducing agent collaboration.

V. Conclusions and Future Work

This paper proposes a multi-agent RL-based approach to demand response, using information on current load as well as 24-hour load prediction to shift demand to periods of low demand or high availability of renewable resources. Results show agents successfully learning to shift neighbourhood demand to the off-peak periods based on providing them current load information and load prediction for the next 24 hours. We envisage this work to be part of a large-scale multi-agent smart grid simulation, and will therefore be further extended, in the first instance, to include collaboration between EV agents to obtain further improvements in demand shifting. The multi-agent system will also be extended to include scheduling and learning on other types of devices (e.g., heating and cooling systems), to incorporate load priority and introduce supply agents (e.g., those controlling wind turbines). Load prediction will also be further investigated as it will need to account for the changes in demand based on demand response achieved by RL agents.

VI. Acknowledgment

This work was supported, in part, by Science Foundation Ireland grant 10/CE/I1855 to Lero - the Irish Software Engineering Research Centre (www.lero.ie)

References

[1] R. S. Suton and A. G. Barto, *Reinforcement Learning: An Introduction*. Cambridge, Massachusetts: A Bradford Book. The MIT Press, 1998.

[2] P. Luh, L. Michel, P. Friedland, C. Guan, and Y. Wang, “Load forecasting and demand response,” in *Power and Energy Society General Meeting, 2010 IEEE*, july 2010, pp. 1 –3.

[3] T. Logenthiran, D. Srinivasan, and T. Z. Shun, “Demand side management in smart grid using heuristic optimization,” *Smart Grid, IEEE Transactions on*, vol. 3, no. 3, pp. 1244 –1252, sept. 2012.

[4] S. Chan, K. Tsui, H. Wu, Y. Hou, Y.-C. Wu, and F. Wu, “Load/price forecasting and managing demand response for smart grids: Methodologies and challenges,” *Signal Processing Magazine, IEEE*, vol. 29, no. 5, pp. 68 –85, sept. 2012.

[5] H. Jiang and Z. Tan, “Load forecasting in demand response,” in *APPEEC*, march 2012, pp. 1 –4.

[6] Q. Dam, S. Mohagheghi, and J. Stoupis, “Intelligent demand response scheme for customer side load management,” in *Energy 2030 Conference, 2008. ENERGY 2008. IEEE*, nov. 2008, pp. 1 –7.

[7] Y.-Y. Hsu and C.-C. Su, “Dispatch of direct load control using dynamic programming,” *Power Systems, IEEE Trans. on*, vol. 6, no. 3, aug 1991.

[8] K.-H. Ng and G. Sheble, “Direct load control-a profit-based load management using linear programming,” *Power Systems, IEEE Transactions on*, vol. 13, no. 2, pp. 688 –694, may 1998.

[9] E. Galvan, C. Harris, I. Dusparic, S. Clarke, and V. Cahill, “Reducing electricity costs in a dynamic pricing environment,” in *IEEE SmartGridComm*, November 2012.

[10] Z. Wen, H. R. Maei, and D. ONeill, “Optimal demand response using device based reinforcement learning,” Stanford University, Tech. Rep., 2012.

[11] A. Dimeas and N. Hatziargyriou, “Multi-agent reinforcement learning for microgrids,” in *IEEE PES General Meeting*, july 2010.

[12] R. Fazal, J. Solanki, and S. Solanki, “Demand response using multi-agent system,” in *North American Power Symposium (NAPS), 2012*, sept. 2012, pp. 1 –6.

[13] L. Meschiari, C. Harris, and S. Clarke, “Analysis of approaches to coordinated charging of electric vehicles on the distribution grid,” in *2nd International Conference on Smart Grids and Green IT Systems (SMARTGREENS)*, 2013, to appear.

[14] E. L. Karfopoulos and N. D. Hatziargyriou, “A multi-agent system for controlled charging of a large population of electric vehicles,” *Power Systems, IEEE Transactions on*, vol. PP, no. 99, p. 1, 2012.

[15] G. Tesauro, “Reinforcement learning in autonomic computing: A manifesto and case studies,” *IEEE Internet Computing*, vol. 11, no. 1, pp. 22–30, 2007.

[16] C. J. C. H. Watkins and P. Dayan, “Technical note: Q-learning,” *Machine Learning*, vol. 8, no. 3, pp. 279–292, May 1992.

[17] M. Humphrys, “Action selection methods using reinforcement learning,” in *Proceedings of the Fourth International Conference on Simulation of Adaptive Behavior*. MIT Press, 1996, pp. 135–144.

[18] I. Dusparic and V. Cahill, "Using reinforcement learning for multi-policy optimization in decentralized autonomic systems - an experimental evaluation," in *Proceedings of the 6th International Conference on Autonomic and Trusted Computing*, ser. Lecture Notes in Computer Science, W. Reif, G. Wang, and J. Indulska, Eds., vol. 5586, 2009, pp. 105–119.

[19] ——, "Autonomic multi-policy optimization in pervasive systems: Overview and evaluation," *ACM Trans. Auton. Adapt. Syst.*, vol. 7, no. 1, pp. 11:1–11:25, May 2012.

[20] A. Marinescu, C. Harris, I. Dusparic, S. Clarke, and V. Cahill, "Residential electrical demand forecasting in very small scale: An evaluation of forecasting methods," *2nd International Workshop on Software Engineering Challenges for the Smart Grid*, pp. 25–32, May 2013.

[21] Comission for Energy Regulation smart meter trial, http://www.ucd.ie/issda/data/commissionforenergyregulation/, last accessed May 2013.

[22] GridLAB-D. U.S. Department of Energy at Pacific Northwest National Laboratory. http://www.gridlabd.org/.

[23] I. Dusparic and V. Cahill, "Autonomic multi-policy optimization in pervasive systems: Overview and evaluation," *ACM Trans. Auton. Adapt. Syst.*, vol. 7, no. 1, pp. 11:1–11:25, May 2012. [Online]. Available: http://doi.acm.org/10.1145/2168260.2168271