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
Future Cities: Trinity Centre for Smart and Sustainable Cities
Trinity College Dublin
Ireland

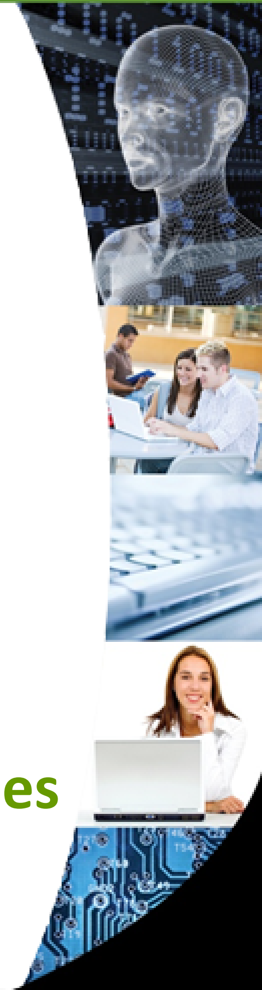

Presentation Outline

- Background
  - Demand response
  - Smart grid as multi-agent system
  - Reinforcement learning
- Demand response using RL
  - Agent design and policies
- Experimental set up
- Results
- Conclusion and Future Work

Demand response (DR)

- Energy usage not distributed evenly during the day
- Morning peak, large evening peak, valley during the night
- Demand response - modification of the consumers' electricity consumption with respect to their expected consumption

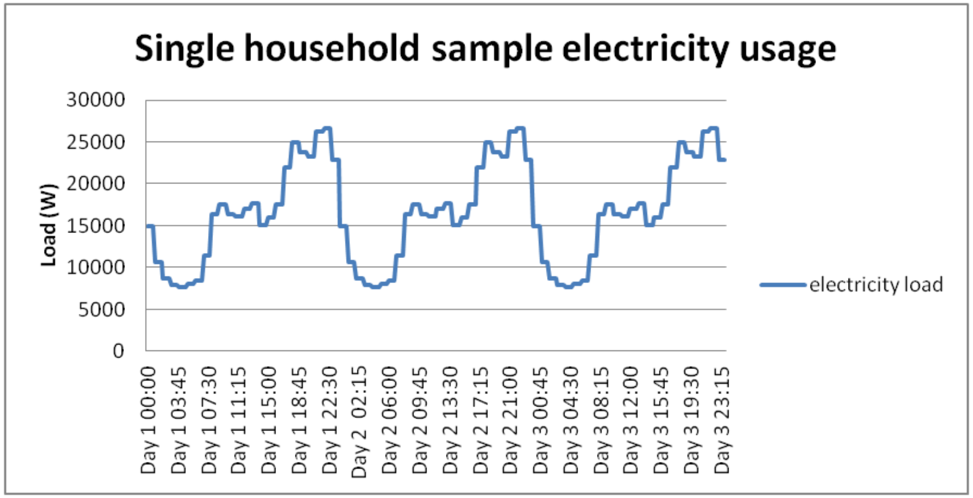

- Demand response techniques – peak clipping, valley filling, load shifting …
- Based on prediction – influence consumers to defer loads that are not essential during the peaks and run them during low demand periods instead
- Can influence use of renewable too – defer loads during the periods of low availability etc

Smart grid as a multi-agent system

- Smart grid is changing:
  - Increasingly dynamic
    - Small and large scale renewable generation
    - Unpredictable weather patterns/renewable generation patterns
  - Increased demand – electric vehicles (EVs), electric heating
  - Changes in demand pattern as more awareness by the consumer (smart meters)
  - Sensor and device usage data available
- Centralized management increasingly unfeasible
- Implement the grid as a multi-agent system
  - Each agent learns its optimal behaviour, cooperates with other agents

Reinforcement learning (RL)

- Use RL to learn agent behaviours
  - Model-free
  - Takes into account long-term effects of agent's actions
- Learn suitable actions through interaction with environment:
  - Receive feedback (reward, reinforcement) from the environment
  - Learn quality of particular actions in particular environment states
  - Stationary environment
- Q-learning (Watkins and Dayan, 1998)
  - Q-value, Q (s,a)
  - Single-agent single-policy
  - Model-free RL technique

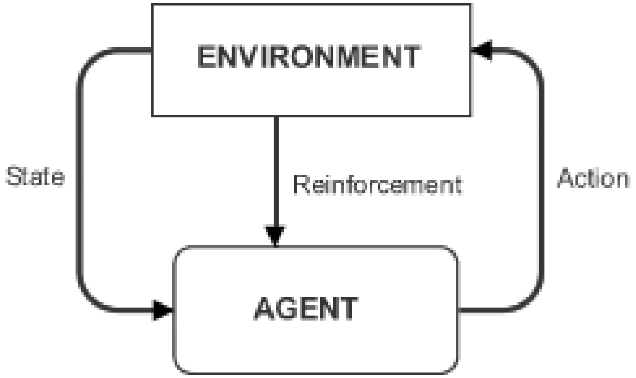

Distributed W-Learning

- W-learning (Humphrys, 1996)
  - Learn dependencies between local policies
- Distributed W-Learning (DWL) (Dusparic and Cahill, 2012)
  - Learn dependencies between neighbouring agents
- Each agent learns how its actions affect its immediate neighbours
  - Implemented as Remote Policies

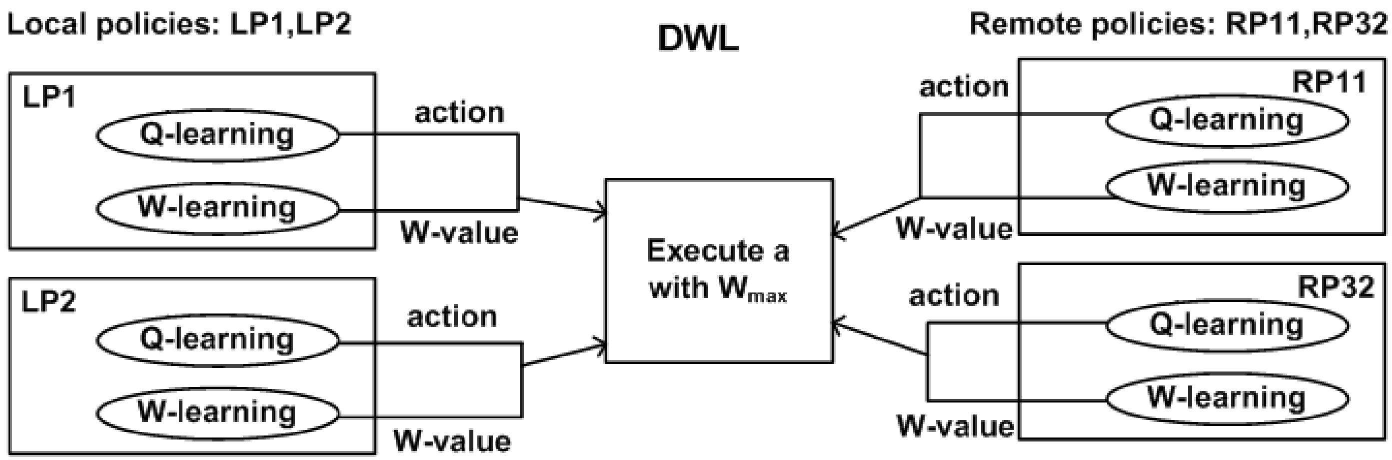

DR using RL

- Each EV is controlled by an RL-agent which is capable of implementing 3 policies:
  - Policy 1: achieve at least the minimum required battery charge
  - Policy 2: charge at the minimum possible price/during the lowest load
  - Policy 3: keep under set transformer limits

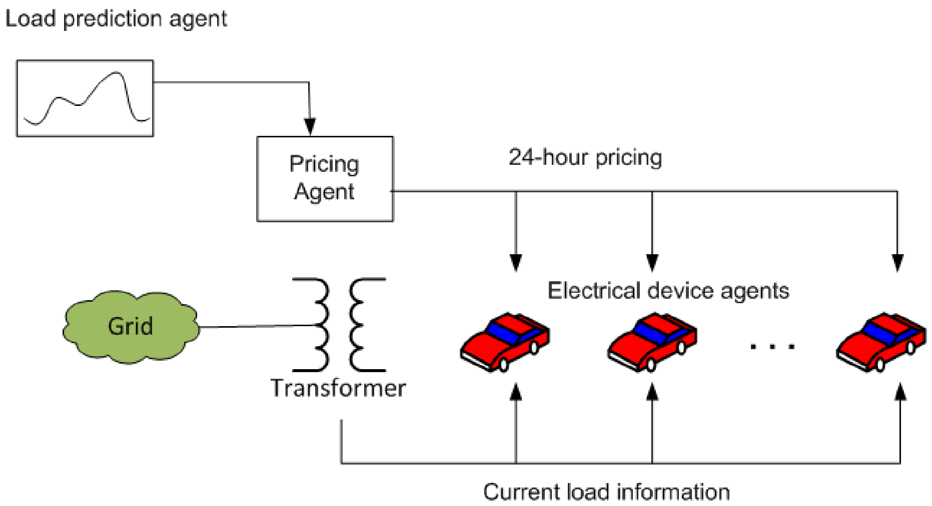

Experimental scenarios

- Baseline:
  - EVs without intelligent control, i.e., charging when they arrive home until fully charged
- Scenario 1:
  - Agents given current load information only
  - Charge and do not go over transform limits
- Scenario 2:
  - Agents given current and predicted load information
  - Charge at minimum cost
- Scenario 3:
  - Agents given current and predicted load information
  - Charge at minimum cost and do not go over transform limits

Experimental setup

- Simulations performed with 9 households (with EV agent + base load each) in Gridlab-D, EV agents implemented using DWL library.
- Vehicles have a battery capacity of 30 kWh
- Vehicles charge at rate of approximately 1.4kW per hour.
- 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).
- Charging decisions made every 15 minutes
- Base load taken from the data recorded in smart meter trial performed in Ireland in 2009.
- Base load in each household ranges from 0.8 kW to 3 kW, based on time of the day

Results – DR based on current load

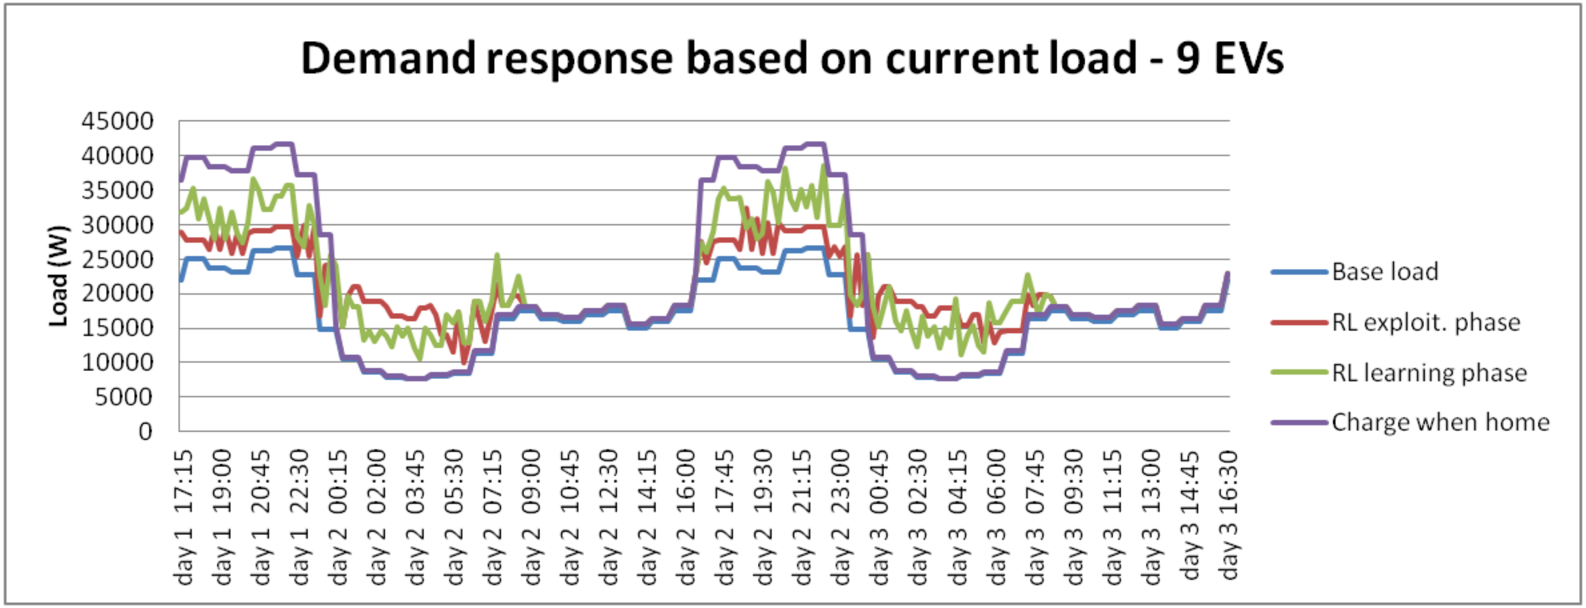

- When EVs charge immediately when arriving home, peak usage increased by additional ~ 60%
- Agents learn (mostly) to postpone charging until off-peak period

Results – DR using load prediction

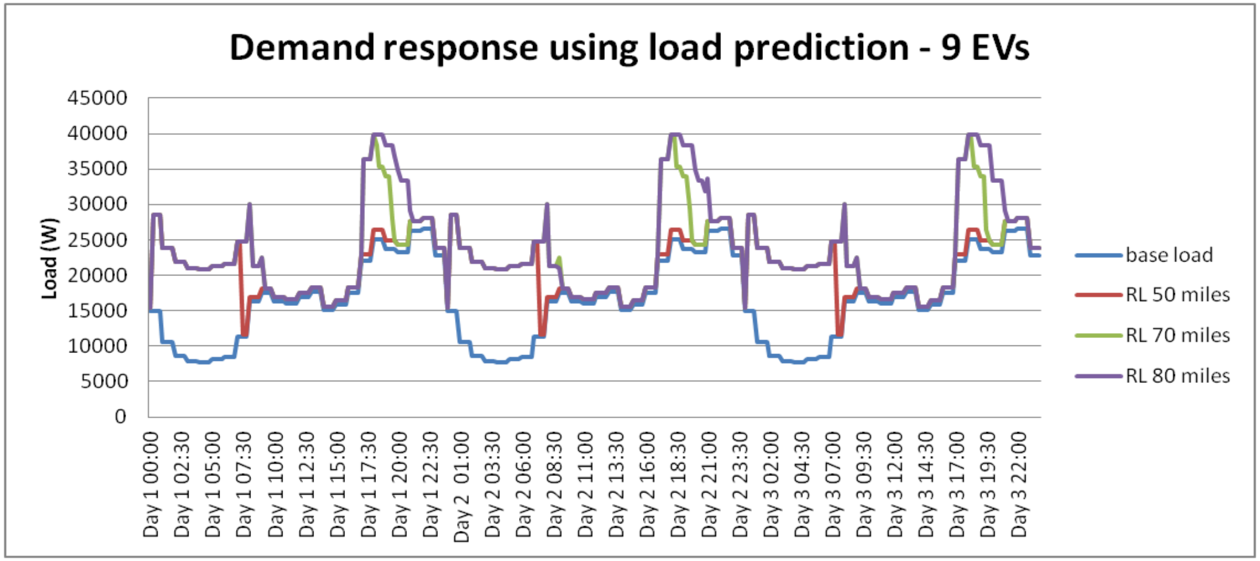

- When mileage is only 50 miles, off-peak period long enough to achieve required charge
- Increase mileage to 70, 80, 100 miles - an agent learn that off-peak not long enough and charge during peak time but only for as long as required to “top up” the battery
- However, when all agents learn that off-peak period not long enough, they all top up the battery at the same peak time
  - Re-introduce current do not overload the transformer policy

THE IRISH SOFTWARE ENGINEERING RESEARCH CENTRE

Results – DR using current and predicted load

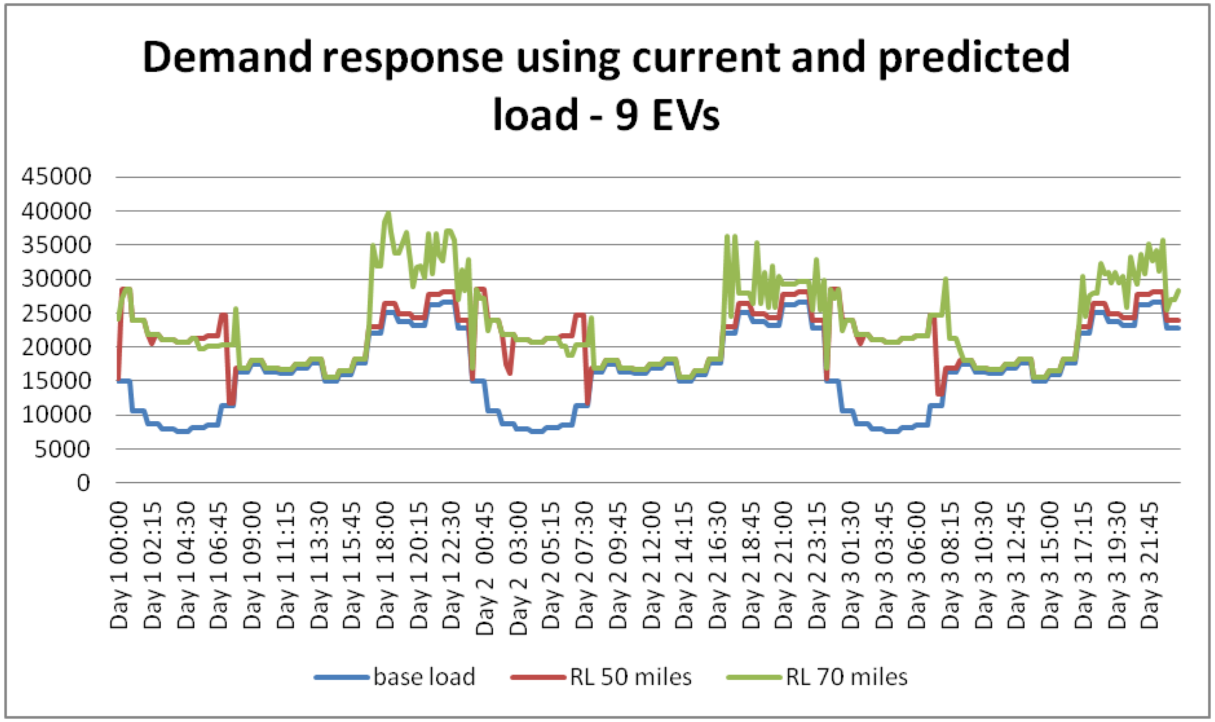

- Agents learn to distribute their top up charging at different times during the peak period
  - Still improvements to be made in smoothing-out the load though
  - Introduce collaboration?

Conclusions and Future Work

- RL suitable for demand response
- Incorporating prediction can improve shifting of loads to off-peak periods
- Still room for improvement
  - Introduce collaboration
- Model other devices apart from EVs
- Introduce supplier/generator agents – introduce renewables
- Address non-stationary properties of the system, e.g.,
  - Current load differs from predicted one
  - Current renewable supply differs from predicted one
  - Unexpected increase in supply/decrease in generation

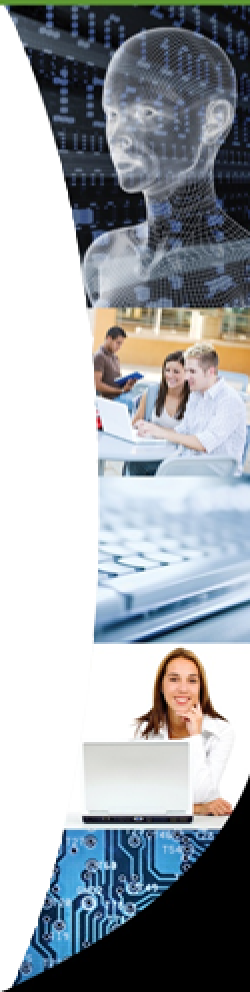

Thank you!

Questions?

Ivana.Dusparic@scss.tcd.ie