Coordinated charging of electric vehicles

A. Albrecht[a]**, P. Pudney** [b]

[a]*Centre for Industrial and Applied Mathematics, University of South Australia, Mawson Lakes, South Australia, 5095*
[b]*Barbara Hardy Institute, University of South Australia, Mawson Lakes, South Australia, 5095*
Email: Peter.Pudney@unisa.edu.au

Abstract: Electric vehicles recharged from renewable energy sources will play an important role in reducing the key adverse impacts of transport: air pollution, depletion of oil resources, and CO_2 emissions. But they will also add to the demand for electricity. Most electric vehicles will be charged at home, overnight. If they all charge at the same time, during the evening peak, they could cause the electricity supply of households or streets to become overloaded. Increases in peak demand ultimately require upgrades to the transmission and distribution infrastructure, which add significantly to the cost of electricity.

Time-of-use tariffs can encourage users to shift their demand to times when electricity is cheaper because demand is usually lower. But to better manage the demand for electricity and control peak demand, we need appliances, including electric vehicles, that respond in real time to changing supply conditions by shifting their demand. Ultimately, having demand that can respond to the available supply will allow greater use of variable, renewable energy sources.

An electric vehicle will typically need to charge for a few hours overnight. Coordinating the charging of electric vehicles distributed over many premises in an area can smooth the total demand in that area. In this paper we describe a simple load control system that uses signals broadcast from street transformers, typically servicing 30–50 premises, to influence the behaviour of smart appliances in the supplied premises to control the demand profile.

Keywords: Electric vehicles, charging, demand management

1 ELECTRIC VEHICLES AND THE GRID

Electric vehicles recharged from renewable energy sources can play an important role in reducing adverse impacts of transport: air pollution, depletion of oil resources, and CO_2 emissions. Widespread use of electric vehicles is feasible. Despite their limited range compared to fossil-fueled vehicles, electric vehicles with a range exceeding 100 km are capable of meeting over 90 per cent of daily trips by passenger cars in Australia. The extra electrical energy that would be required if all current passenger car transport in Australia were electric would increase the demand for electricity by about 11% (Albrecht et al., 2009).

Without careful control of where and when electric vehicles are charged, vehicle charging will increase peak demand, which will decrease the reliability of electricity supply or increase the cost of reinforcing the electricity network. With proper control, however, electric vehicles could be used to flatten the demand for electricity, to help balance demand and supply, to improve the overall efficiency of the grid, and, once appropriate standards have been developed and implemented, provide controllable energy storage.

In this paper we describe a simple load control system that uses signals broadcast from street transformers, typically servicing 30–50 premises, to influence the behaviour of smart appliances in the supplied premises to control the demand profile.

Where will EVs charge?

Recent experience in Europe, Japan and the UK has shown that while public charging infrastructure can encourage the use of electric vehicles, most charging is done at home—cars are generally not used at night and so this is an ideal time to charge, there is a greater availability of electricity overnight, and electric vehicle owners with home charging facilities can be assured of getting a park and a charge.

The UK Government's strategy (OLEV, 2011) is:

- most charging will occur at home, at night, off peak, using cheaper, lower-carbon electricity

- workplaces will provide recharging facilities for their fleets, and for employees who cannot recharge at home

- public infrastructure should target places where it is needed and is commercially viable.

When will EVs charge?

Household travel survey data from Brisbane, Perth, Sydney, Melbourne and Adelaide shows when vehicles arrive home after their last trip each day (Holyoak and Zito, 2012). The results are similar for each region. The median arrival time is between 17:00 and 17:45, with no more than 20% of cars returning home in any hour. More than 90% of cars are home by 22:00.

Figure 1 shows how demand on the National Electricity Market varies with time of day; the ideal time to charge is after 22:00, when demand from other uses declines.
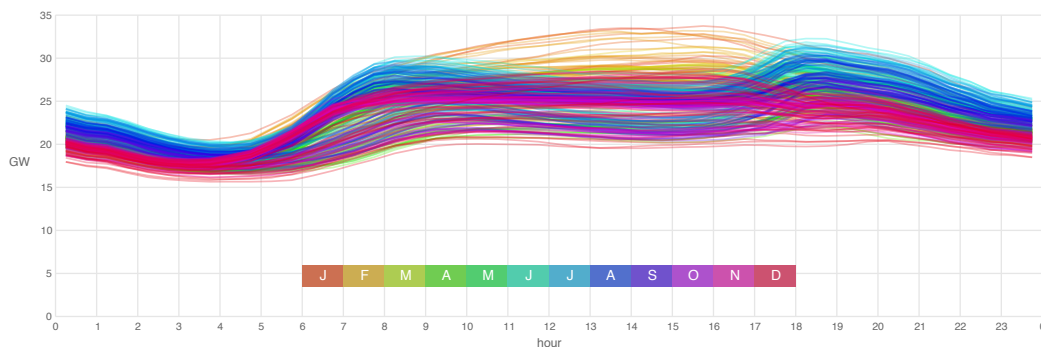

Figure 1. Demand on the National Electricity Market during 2010. Each curve is demand for a day. Demand is still high until late evening.

Charging durations

Charging duration depends on the distance travelled, on the energy consumption of the vehicle, and on the charging power:

$$\text{charging duration (h)} = \frac{\text{energy consumption (Wh/km)} \times \text{distance (km)}}{\text{charging power (W)}}.$$

Cars that travel short distances each day will have relatively low charging durations. The distributions of daily distance travelled by cars in Adelaide, Melbourne, Sydney, the Sydney Greater Metropolitan Area (including Newcastle and Wollongong), Perth and Brisbane are given by Holyoak and Zito (2012). In all regions except the Sydney Greater Metropolitan Area, over 90% of daily travel distances are less than 100 km. The median distance travelled ranges from less than 20 km in Brisbane to 35 km in Sydney.

The energy required to recharge a vehicle, and hence the recharge duration, also depends on the energy consumption of the vehicle—the amount of electrical energy, measured at the charging plug, required to recharge the car per kilometre travelled.

The final factor that determines charging duration is charging power. Electric vehicles in Australia typically charge from the AC supply at 2.3–3.6 kW, though future cars could draw up to 7 kW. High-power DC charging, up to 50 kW, is also possible but requires special charging equipment that uses too much power for home charging.

Coordinating loads

Apart from a few days each year when high demand threatens the reliability of the electricity system, users of electricity are generally unaware of how the system operates and its constraints. Controlling when appliances operate to manage the total demand of an individual premises can prevent overloading and, with Time of Use tariffs, reduce the cost of electricity to the customer. But supply constraints could be at the street level or suburb level, in which case it will be necessary to control the demand of many appliances distributed over many premises.

Ramchurn et al. (2011) describe a model of *decentralised demand side management* where each household has an agent which optimises the scheduling of deferrable loads to minimise energy costs while maintaining the required level of comfort and utility in the home. They propose the use of real-time price signals as an incentive for agents to defer demand, but note that having many agents react simultaneously to a drop in price can cause new demand peaks. They overcome this problem by having the agents react gradually to the predicted prices for the next day. Using a simulation of 5000 homes using real time pricing, they show that they can reduce demand peaks and CO_2 emissions, and that the benefits to households are sufficient to encourage consumers to use the system.

Gerding et al. (2011) describe a system for allocating charging to electric vehicles so that capacity constraints in a local area are not exceeded. Each vehicle owner specifies the value they place on their first hour of charging, second hour of charging, and so on, as well as the time interval during which the vehicle will be available for charging. An online mechanism allocates charging capacity each hour to the vehicles with the greatest value. The pricing mechanism is designed so that there is no incentive for owners to misreport their values or charging interval.

Mets et al. (2010) describe optimisation methods for determining how much energy a plug-in vehicle charger should use in each 15 minute period in order to make the household load as flat as possible, or to make the load from many households as flat as possible. Their simulations indicate that with 30% penetration of plug-in hybrid electric vehicles, uncontrolled charging increased the peak demand of a group of 150 homes by a factor of 1.62. Local energy control of individual chargers reduced the new peak so that it was 1.19 times the original peak; a global control strategy reduced the peak to 1.12 times the original peak.

2 A SIMPLE AUTOMATED LOAD CONTROL SYSTEM

We have developed a simple heuristic load control system that would allow signals from street transformers, typically servicing 30–50 premises, to influence the behaviour of smart appliances, including electric vehicle chargers, in the supplied premises. The design requirements were:

- the load control signal is based only on the instantaneous supply and demand at the street transformer

- load control signals are broadcast to all appliances in the distribution area

- appliances act independently, with no communication between appliances or to the control centre

- the interaction of the control system and the appliances should result in priority being given to appliances with greater urgency.

If there will be enough energy for each appliance to complete its operation then the system should ensure that each appliance finishes. On occasions when there is insufficient capacity, the available energy should be distributed fairly. Possible definitions of 'fair' are discussed later. The key idea is that each appliance determines its priority, which may change in time. The control system broadcasts a signal that indicates the priority required for an appliance to operate. Appliances with a sufficiently high priority can commence or resume operation; appliances with low priority may have to pause operation.

The load control signal comprises two values, $0 \leq u_p < u_r \leq 1$:

- u_p is the *pause level*—operating appliances with priority $0 < p \leq u_p$ should pause operation. If $u_p = 0$ then no appliances need pause.

- u_r is the *resume level*—appliances with priority $u_r < 1 <= 1$ can start or resume operation. Note that appliances with priority $p = 1$ can resume operation only if $u_p < 1$.

These control signal levels are varied by the street transformer in response to the observed total load, in order to drive demand from its current level to some desired level. The control signal is calculated as follows:

- If there is enough spare capacity for one or more additional appliances to operate then the control u_r is gradually reduced until one or more appliances resume and there is no more spare capacity. We do not require any operating appliances to pause, and so set $u_p = 0$.

- If the system is overloaded, the control u_p is increased until enough appliances are paused. We do not want additional paused appliances to start, and so set $u_r = 1$.

- If demand is close to capacity then we set $u_p = 0$ and $u_r = 1$, so that no operating appliances will pause and no paused appliances will resume.

We consider appliances that are required to operate for a specified or estimated duration, called the *operating duration*, within a specified period, called the *enabled duration*. For example, if an electric vehicle charger estimates that it needs to operate for three hours, and it has to be finished within ten hours, then the operating duration is three hours and the enabled duration is ten hours.

An appliance is characterised by:

- the average operating load

- an *enabled timer*, t_e, which indicates the time remaining until the appliance will be disabled

- an *operating timer*, t_o, which indicates the operating time remaining

- the minimum duration that the appliance must be operating before being paused

- the minimum duration that the appliance can be paused.

The *priority* p of an appliance is $p = \min\{1, t_o/t_e\}$. If a paused appliance has priority $p > u_r$ then it is eligible to start or resume operation. Ideally, however, we want only one appliance at a time to respond to the control signal. We can restrict the number of appliances responding at any instant by having each eligible appliance respond at regular intervals with probability less than one, or by having each eligible appliance wait for a random duration before responding.

The behaviour of each appliance is controlled by a finite state machine. The states are:

- **Ready**: the appliance is ready to operate, but not operating. If the priority of the appliance is $p > u_r$ then, with probability p_r, the appliance will start operating and change to the **Operating** 0 state.

- **Operating** t: the appliance resumed operating t seconds ago. If the appliance finishes operating, the state changes to **Finished**. If t is greater than the minimum operating duration and priority $p \leq u_p$ then, with probability p_p, the appliance will pause operation and change to the **Paused** t_p state, where t_p is the minimum pause duration for the appliance.

- **Paused** t: the appliance will be ready to operate in t seconds. The paused timer is decremented every second. If $t = 0$ then the state changes to **Ready**.

- **Finished**: the appliance has finished.

The initial state for each appliance is **Ready**.

Examples

We have developed software that simulates the operation of this control system. Figure 2 shows the operation of 50 uncontrolled electric vehicle chargers, each drawing the same power while operating.

The horizontal axis is time, from 18:00 to 06:00 the next day. In the lower part of the graph, each of the thin grey horizontal lines represents the enabled duration for a charger. The cars are plugged in at random times distributed uniformly between 18:00 and 19:00. The thick coloured lines indicate when each charger is operating. In this case, each charger starts operating as soon as it is plugged in. The operating durations are calculated from the distribution of daily driving distances in Adelaide.

The upper part of the graph shows the total charger load (blue) and the desired load (orange). The desired load has been chosen arbitrarily, but has low spare capacity in the early evening and the highest spare capacity between 03:00 and 04:00.

The middle, shaded part of the graph indicates the priority of each of the tasks (coloured lines). Priority decreases as each charger operates. In this case, the two chargers with the longest required charging durations (green and purple) start with relatively high priority.

In this case, there is no control signal; each charger starts operating as soon as it is enabled. As a result, the total demand peaks early in the evening, exceeding the available capacity.

Implementing our control scheme gives the result shown in the Figure 3. The values of the control signals u_p and u_r are indicated in the shaded middle section of the graph. The resume priority level u_r (green) decreases from one whenever there is spare capacity; the pause priority level u_p (red) increases from zero when the system is overloaded. Appliance priorities increase for paused appliances and decrease for operating appliances. Apart from a few instances where two appliances resumed at the same time, the demand load stays below the desired level. Furthermore, the appliances with the greatest required operating durations tend to start earlier.

There will usually be plenty of energy available overnight for all of the cars in a neighbourhood to charge. Figure 4 shows what happens when there is only just enough energy available.

Figure 5 shows what happens when there is not enough capacity to charge all cars. In this case, the desired demand drops for four hours to a level sufficient to keep just two chargers operating. During this period, the priority of all but the two operating chargers increases. When the capacity increases to the normal level at 04:00, several appliances have priority $p = 1$. These appliances resume operation. The control u_r then reduces and more appliances resume until there is no spare capacity. Each time the decreasing capacity is exceeded, the control u_p is is increased until an appliance pauses. As operating appliances finish, paused appliances resume. At the end of the night, only the appliances requiring the shortest charge durations have missed out.

Fairness

Fairness is particularly important when the total demand for energy exceeds the available capacity, or if power is not available for long enough for some appliances.

Let P_i be the operating power of appliance i, r_i be the initial required operating duration for appliance i and a_i be the final allocated operating duration for appliance i. An allocation $\{a_i\}$ could be considered fair if:

- $a = \max a_i$ is minimised. Appliances requiring an operating duration less than a will finish; the remaining appliances will each have operating duration a.
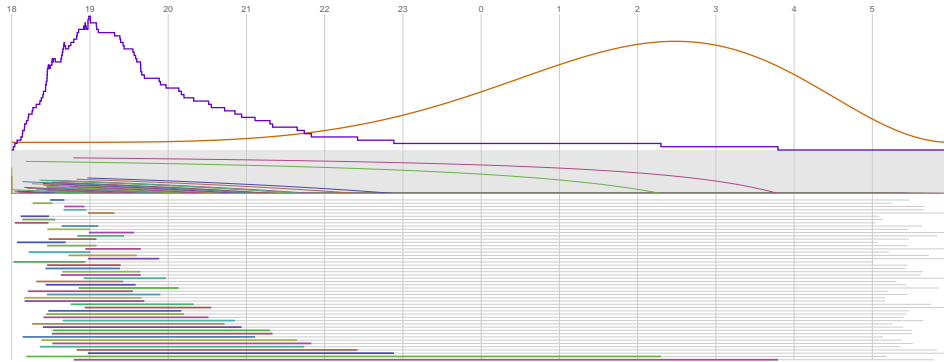
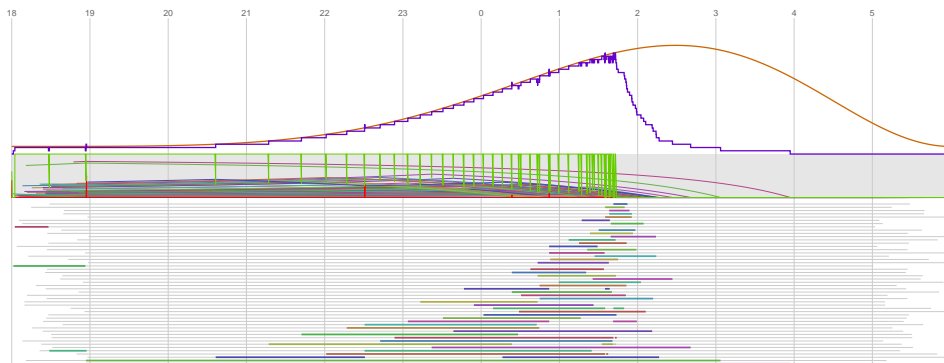**Figure 2**. Uncontrolled charging—see the text for a description of the graph.
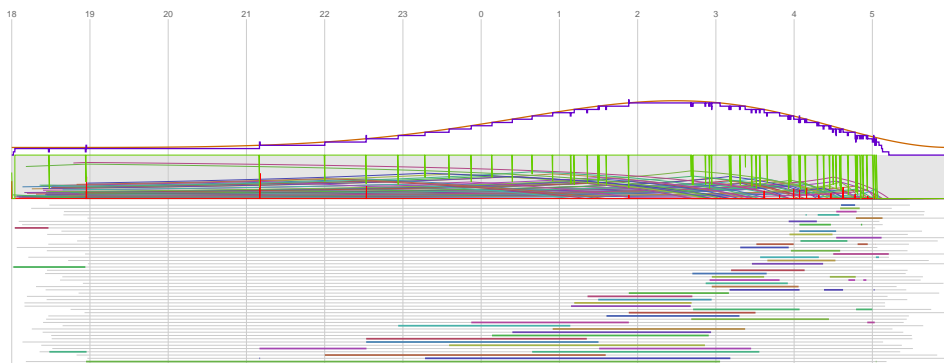

Figure 3. Controlled charging.
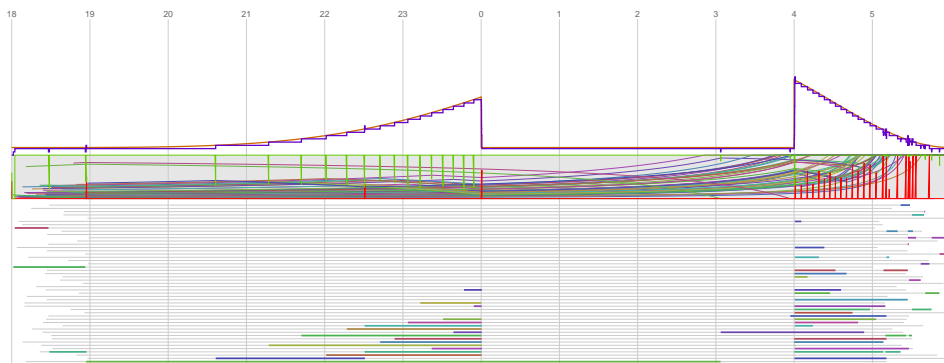

Figure 4. Controlled charging with less capacity.

Figure 5. Controlled charging with a severe power restriction.

- $r = \max(r_i - a_i)$ is minimised. This policy allocates time to the appliances that need it most. At the end of the night, the maximum remaining operating duration r will be as small as possible.

- $E = \max(a_i P_i)$ is minimised. Appliances requiring energy less than E will finish; the remaining appliances will each receive energy E.

- $\rho = \max(a_i/r_i)$ is minimised. Such allocations are fair in that no appliance receives a greater proportion of its required operating duration than any other. Some appliances with large r_i may have $a_i/r_i < \rho$ if spare capacity is available for less than duration r_i.

Other definitions of fairness are also possible. Further work is required to determine the extent to which various fairness outcomes can be achieved using the simple priority-based appliance policy.

A deficiency of our simple control strategy is that operating appliances continue to operate even when paused appliances have higher priority. We can, however, use the same control signals and the same appliance rules to implement more sophisticated control schemes that address these issues. For example:

- If the control signal u_p is periodically increased, an operating appliance with low priority will be paused, the available capacity will increase, and so u_r will decrease until a paused appliance with high priority resumes.

- The control priorities u_p and u_r could be adjusted continuously, forming a priority band. Paused appliances with priority above the band will resume; operating appliances with priorities that fall below the band will be paused. The upper and lower limits of the band can be adjusted to control the number of appliances operating at any instant.

A disadvantage of such schemes is that appliances may pause and resume several times.

Further work is required to develop fair incentives and rewards that will encourage households to participate.

REFERENCES

Albrecht, A., N. Holyoak, R. Raicu, P. Pudney, M. Taylor, R. Zito, and G. Groves (2009). Infrastructure requirements for electric vehicles in Australia. Technical report, CRC for Automotive Technologies.

Gerding, E. H., V. Robu, S. Stein, D. C. Parkes, A. Rogers, and N. R. Jennings (2011). Online mechanism design for electric vehicle charging. In Yolum, Turner, Stone, and Soneberg (Eds.), *Proceedings of the 10th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2011)*.

Holyoak, N. H. and R. Zito (2012). EV applicability: Report on quantification of national daily driving range. Technical report, AutoCRC.

Mets, K., T. Verschueren, W. Haerick, C. Develder, and F. D. Turck (2010). Optimizing smart energy control strategies for plug-in hybrid electric vehicle charging. In *Network Operations and Management Symposium Workshops*, pp. 293–299.

OLEV (2011, June). Making the connection: The plug-in vehicle infrastructure strategy. Technical report, UK Office for Low Emission Vehicles.

Ramchurn, S. D., P. Vytelingum, A. Rogers, and N. R. Jennings (2011, May). Agent-based control for decentralised demand side management in the smart grid. In *The Tenth International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2011)*.