

Putting the “Smarts” into the Smart Grid: A Grand Challenge for Artificial Intelligence

Sarvapali D. Ramchurn, Perukrishnen Vytelingum
Alex Rogers, and Nicholas R. Jennings
University of Southampton
Southampton, SO17 1BJ, UK
{sdr,pv,acr,nrj}@ecs.soton.ac.uk

1. INTRODUCTION

The phenomenal growth in material wealth experienced in developed countries throughout the twentieth century has largely been driven by the availability of cheap energy derived from fossil fuels (originally coal, then oil, and most recently natural gas). However, the continued availability of this cheap energy cannot be taken for granted given the growing concern that increasing demand for these fuels (and particularly, demand for oil) will outstrip our ability to produce them (so called ‘peak oil’) [9]. Many mature oil and gas fields around the world have already peaked and their annual production is now steadily declining. Predictions of when world oil production will peak vary between 0-20 years into the future, but even the most conservative estimates provide little scope for complacency given the significant price increases that peak oil is likely to precipitate [1]. Furthermore, many of the oil and gas reserves that do remain are in environmentally or politically sensitive regions of the world where threats to supply create increased price volatility (as evidenced by the 2010 Deepwater Horizon disaster and 2011 civil unrest in the Middle East). Finally, the growing consensus on the long term impact of carbon emissions from burning fossil fuels suggests that even if peak oil is avoided, and energy security assured, a future based on fossil fuel use will expose regions of the world to damaging climate change that will make the lives of many of the world’s poorest people even harder [15].

Against this background, many governments around the world have begun taking action to transition to a low carbon economy. For example, the United Kingdom has legislated to reduce CO₂ emissions by 80% by 2050 (compared to 1990 levels) [8]. Achieving this aim requires that the direct use of fossil fuels that we are familiar with today is almost entirely eliminated. Thus, the use of electric vehicles and high speed electric trains will have to become widespread in order to reduce our reliance on oil for transportation.¹ Likewise, our homes and offices will have to be heated by efficient ground and air source heat pumps powered by electricity rather than existing natural gas and oil fired boilers [22]. As a result (and given the general growth of the world economy), electricity demand across the world is predicted to increase by 76%, or 4800 gigawatts (GW), by 2030 (compared to 2007 levels) [20]. Crucially, much of the electricity needed to meet this demand will have to be generated from renewable wind, so-

¹Electric motors are inherently more efficient than internal combustion engines, and are ‘future proof’ in that their carbon emissions reduce as the electricity used to supply them become cleaner.

lar, and tidal sources rather than the coal and natural gas power plants that we use today.

It is this increased demand for electricity, and the requirements for its generation, that present perhaps the greatest challenge. In most countries, the electricity grid has changed very little since it was first installed, and all existing grids are predicated on the central idea that electricity is produced by a relatively small number of large fossil fuel burning power stations and is delivered to a much larger number of customers, often some distance from these generators, on-demand. The grid itself relies on ageing infrastructure (e.g., 40-year old transmission lines and transformers, and 20-year old power stations), is plagued by poor information flow (e.g., most domestic electricity meters are read at intervals of several months), and has significant inefficiencies arising from losses within the transmission (on a national level) and distribution (on a local level) networks [12].

The vision of an electricity grid that makes extensive use of renewable generation challenges this current situation. Renewable generation is both intermittent and distributed, with the output of such generators being determined by local environmental conditions (such as wind speeds and cloud cover in the case of wind turbines and photo-voltaic (PV) solar panels, respectively) that can vary significantly over minutes and hours. Thus, it will no longer be possible for supply to continuously follow the vagaries of consumer demand, but rather, the demand-side will have to be managed to ensure that demand for electricity is matched against the available supply. Electric vehicles will play a part in this, since not only do they represent a significant extra load that must be satisfied, but more positively, they also provide a distributed form of energy storage² which may allow the grid to smooth out this variable supply.

Furthermore, meeting the increased demand for renewable generation may require hundreds of thousands, or even millions of such generators, distributed across both the transmission and distribution networks. These generators may need to act together, effectively working as virtual power plants, or may be located on every building across the grid,

²Energy storage in existing grids is typically limited to a small number of pumped storage generators that pump water from a low reservoir to a high one when electricity is plentiful, and recover this potential energy by letting the water flow back through a turbine, when electricity is in short supply.

resulting in a distributed network of *prosumers*³ who both produce and consume electricity depending on their local requirements. Thus, unlike existing grids where electricity generally flows one-way from generators to consumers, this will result in flows of electricity that vary in magnitude and direction continuously. To guarantee the security of the network (i.e., the maintenance of stable voltages and frequencies, and the reliability of supply) and to avoid the cascading failures that plague today’s grid,⁴ new control procedures must be devised. Indeed, the number and variability of generators will require that the grid is able to act autonomously, under human supervision but not necessarily under human control, to diagnose potential problems and self-heal.

Thus, there is a growing consensus that existing grids cannot simply be extended to address these challenges, but rather, a fundamental re-engineering of the grid is required; one that envisages the creation of a ‘smart grid’, described by the US Department of Energy [12] as:

A fully automated power delivery network that monitors and controls every customer and node, ensuring a two-way flow of electricity and information between the power plant and the appliance, and all points in between. Its distributed intelligence, coupled with broadband communications and automated control systems, enables real-time market transactions and seamless interfaces among people, buildings, industrial plants, generation facilities, and the electric network.

What is perhaps most striking about this vision is that not only does it present many challenges in terms of power systems engineering, telecommunications, and cyber-security, but at its core are concepts, such as distributed intelligence, automation, and information exchange, that have long been the focus of research within the computer science and the artificial intelligence (AI) communities. In particular, in this paper we argue that the smart grid provides significant new challenges for research in AI since smart grid technologies will require algorithms and mechanisms that can solve problems involving a large number of highly heterogeneous actors (e.g., consumers with different demand profiles or generators with different volatilities), each with their own aims and objectives, having to operate within significant levels of uncertainty (i.e., where the network conditions and the outcome of actions taken by individual entities on the grid will be more unpredictable or uncontrollable) and dynamism (i.e., where demand and supply at different points in the network will be in a significant state of flux). Hence, in the following sections, we illustrate how such issues arise within the key components of the smart grid — demand-side management, electric vehicles, virtual power plants, the emergence of prosumers, and self-healing networks — and by showing which components and which interactions need to be smart,

³The term ‘prosumer’ was coined by futurologist Alvin Toffler in his book *Future Shock* in 1970 in order to describe the actors in the marketplace who would not just consume but also actively participate in the production of customised goods.

⁴The Northeast Blackout of 2003 that forced the shut-down of over 100 power plants and affected 55M people — the largest black-out in US history — was precipitated by a single overloaded transmission line, in Ohio, sagging and touching overgrown vegetation.

we provide a research agenda for this community for making the smart grid a reality.

2. DEMAND-SIDE MANAGEMENT

A key requirement for a safe and efficient electricity grid is that supply and demand are always in perfect balance. Now, in the day to day running of the today’s electricity grid, this is achieved by varying the supply-side in real-time to match demand (increasing and decreasing the output of generators such that voltage and frequency are maintained across the grid). Hence, the idea that electricity should be available at all times at the flick of a switch has permeated most, if not all, of our daily activities in the modern world.

However, as far back as the 1980s, Schweppe and colleagues highlighted numerous reasons why demand for electricity should be made more adaptive to supply conditions [34]. They noted that doing so would allow peaks in demand to be ‘flattened’, thus allowing generation assets to be reduced; particularly, expensive (and carbon-intensive) *peaking plant* that might only be used for several hours or less each day. This flattening would result in longer term and cheaper production contracts, producing a more efficient grid with lower prices for consumers. Furthermore, it would also provide significant benefits for grid operators. For example, if generation capacity was temporarily restricted due to some unforeseen event (either due to faults or if renewable energy sources are unavailable), then controlling demand would ensure that those generators which were available were not overloaded. In addition, after a power failure has occurred, the ability to synchronise demand with supply as connections are recovered and generators are brought up to speed would significantly accelerate recovery from such failures (a point we will come back to in Section 6).

The need for demand-side management is even more apparent within a grid that makes extensive use of intermittent renewable generation. In this case, there is a high likelihood that there will be periods when there is insufficient generation capacity to meet demand. It is thus imperative that demand can be reduced at these times. Conversely, there may also be times when renewable energy is plentiful, and demand should increase to make the best use of this energy.

To date, approaches to reduce demand have been limited to either directly controlling the devices used by the consumers (e.g., automatically switching off high load devices such as air conditioners at peak times), or to providing customers with tariffs that deter peak time use of electricity. The advent of the smart grid with two way information flows, and smart meters making real-time measurements of consumption, would allow demand-side management to be deployed at scale across the entire grid, providing every home and every commercial and industrial consumer with the ability to automatically reduce load in response to signals from the grid.

However, doing so may be ineffective, or at worst, detrimental, since *such initiatives tend to reduce the natural diversity of consumers’ peak demands and shift all of these peaks to specific periods* [36]. For example, static time-of-use (TOU) pricing where the price of electricity at night is cheaper than during the day, has been observed to create sig-

nificant additional peaks in demand as soon as the off-peak period is reached [30, 36]. Similarly, critical peak pricing (CPP), which is often applied on the west coast of the USA to control air-conditioners at peak times, can often create additional peaks as devices turn back on as soon as the critical period is over. Given this, a number of researchers have suggested that more sophisticated tariffs, such as real-time pricing (RTP) or spot pricing (where the price per kWh of electricity consumed, is different for each half-hour, and is provided to the consumer a day, or a few hours, ahead of time), in conjunction with more sophisticated ‘agents’ that can autonomously respond to these price signals, would avoid this [34]. However, even RTP can create unexpected peaks in demand, when all individuals respond to a signal in the same way, and inadvertently synchronise with others [30].

Thus, it appears that demand-side management technologies that simply rely on reacting to control or price signals will not be enough. Rather, what is necessary are more sophisticated approaches that are truly adaptive to the state of the grid, that are able to learn the correct response given any particular situation, and that can look ahead and predict both supply and demand trends in the near future, in order to prepare for future reductions in available supply, or to make the most effective use of supply when it is available.

The design of such intelligent systems is challenged by the complexity of the domains in which they are deployed. For example, within a home, demand reduction may involve shifting the time of use of a number of electrical appliances, each with their own individual constraints (e.g., lighting cannot be shifted, a washing machine can be shifted by a day or two, while a dishwasher may be shiftable by a few hours [24]). Similarly, both heating (given that this will be likely to be electrified through the use of efficient heat pumps) and cooling loads can be shifted as long as the comfort and temperature preferences of the householders are met. To be effective in this, it may also be necessary for such systems to learn the thermal properties of the home in which they are deployed, as well as the local weather conditions, and the way in which these local conditions impact on the heat loss, or gain, of the home. Crucially, these approaches will have to take into account the fact that each individual householder will have her own preferences, and that these preferences must either be explicitly elicited, or learnt. Since these preferences are likely to exhibit change over time, and depend on the current activities of the householder and local weather conditions, in computational terms this translates into an online learning and scheduling problem under uncertainty.

Similarly, commercial and industrial consumers will be constrained by existing contracts and commercial considerations (e.g., a factory may have to deliver products within certain deadlines, while a data centre has to be available to its customers twenty four hours a day), and must balance demand reduction against these additional factors. Large industrial consumers of electricity with significant heating, cooling, or pumping loads may have considerable flexibility regarding when they actually consume electricity as long as some overarching constraints are satisfied.⁵ However, to do

⁵During the 2000 California electricity crisis, which saw ex-

so in a responsive way, requires that the usage optimisation algorithm that is deployed is able to model and predict both the prices within the grid, and also the industrial processes themselves (similar to the home heating setting above where a thermal model of the home must be learnt). Furthermore, in both settings, it will be essential that the householders and business owners are able to understand the consequences of the automated actions that are taken, and are happy to delegate control to an intelligent device or software agent. In this respect, it will be important to define the *adjustable autonomy* of such systems; to what extent should the agent automatically decide to shift devices to run at certain times, and when should it ask for confirmation from the user [33].

Now, the development of these autonomous technologies raises the prospect that such systems will be widely deployed in possibly millions of homes; each individually reacting to prices and to the preferences of householders. Defining the convergence properties (i.e., how the aggregate demand profile will respond to price signals) of such a complex system will be central to the definition of what constitutes safe and efficient behaviours for the grid. In particular, it will be necessary to ensure that neither significant inefficiencies, nor excessive volatility ensue from these autonomous systems converging to poor equilibria (or not converging at all). Hence, it will be important to design simulation systems that can accurately represent both the grid and the reaction of consumers, in order to predict the emergent properties of the system under a range of different conditions (e.g., weather patterns or social activities) and worst case scenarios (e.g., some generators fail or lines trip).

Against this background, recent work has begun to research the use of autonomous agents, representing individual consumers, that interact through markets [40, 10], and individually learn to optimise their use of electrical loads or storage devices in a number of simplified settings [30, 28]. Simulations of such systems point to the effectiveness of adaptive behaviours (that learn to react to prices) on the grid. In addition, human-computer interaction technologies have also been proposed to improve the reaction of users to the information from smart meters [37, 16]. While promising, we believe that this work represents only the beginnings of the research needed in this area.

Thus, in summary, we believe the key AI challenges in demand-side management are:

- Designing automation technologies for heterogeneous devices that learn to adapt their energy consumption against real-time price signals when faced with uncertainty in predictions of future demand and supply, the individual users’ preferences, and the constraints of the overarching system (domestic, commercial, or industrial) within which it is deployed.
- Developing the means by which the automated decisions of these systems can be effectively communicated to, and controlled by, their human owners, whilst allowing a varying range of autonomous behaviours.

tremely high spot prices, several bauxite smelters realised that there was greater profit to be had in reselling electricity that they had bought in long-term forward contracts, than in using it themselves to produce aluminium [3].

- Developing simulation and prediction tools to allow the system-wide consequences of deploying pricing mechanisms and energy management agents to be assessed by grid operators and suppliers.

3. ELECTRIC VEHICLES

With the advent of commercially viable electric vehicles (EV), such as the Nissan Leaf and the Chevy Volt, the coming years are likely to see the large-scale adoption of electric vehicles that will shift the energy requirements of transport from fossil fuels to renewable electricity from the smart grid [12, 26]. EVs are one of the key mechanisms to deliver significant reductions in carbon emissions as the transport sector is one of the largest contributors in most developed countries (about 20% in the UK and 30% in the US), and the majority of these emissions are the result of private motor vehicles. As millions of EVs are deployed onto the roads, novel mechanisms, building upon the communication infrastructure and distributed intelligence in the smart grid, will be needed to ensure that the batteries of these vehicles are fully charged when their owners need to use them, without overloading the network. In addition, these same batteries will form part of the decentralised demand-side management system used to reduced variations in demand and supply by charging when low-carbon renewable energy is plentiful, and discharging back into the grid when it is in short supply; so called vehicle-to-grid or V2G.

In more detail, electric vehicles place a considerable additional load on the grid due to the high charging rates that are necessary to ensure both a reasonable vehicle range of around 100 miles, and the ability to rapidly charge the battery. While a typical house may use between 20 to 50 kWh of energy per day, an EV battery may be charged with 32 kWh of energy in just a few hours [18]. Thus, the total energy required by these vehicles may be comparable to the total electricity consumption within the domestic sector, but all of this demand is likely to be concentrated over particular periods of the day, and over particular geographical areas; both of which are subject to shifts. For example, if all the EVs in a local neighbourhood are charged at the same time (as is likely to happen as householders return home at the end of the day), the local distribution network, and in particular, the street level transformer (which is typically undersized and allowed to cool over night), may become a significant bottleneck to supply. When the owners of these vehicles drive to work and plug in, the demand will shift in both time and geographic distribution. Similar issues occur when a large number of EVs simultaneously attend large scale social events at sporting arenas or shopping malls [26].

Given these continuously changing demands imposed on the local distribution network by the movement and charging of vehicles within it, and the variable supply of renewable energy, it will be necessary to devise sophisticated approaches to schedule the charging of electric vehicles. This scheduling should make the most effective use of what renewable energy is available, while also ensuring that the vehicles' batteries are fully charged when required by their owners. Furthermore, this must be done in the context of uncertainty regarding both the future availability of renewable energy, and future vehicle use. Building upon this, it will be important to design decentralised control mechanisms that can guide

the charging of EVs to various points in the network, given its dynamic conditions and constraints. In particular, these mechanisms will have to take into account that consumers need to be incentivised (e.g., in terms of charging prices or speeds at specific points) to adapt their behaviour as they may only care about their individual travel needs. The challenge is to ensure such incentives are properly designed to induce charging profiles that stabilise the grid (i.e., ensure flows are secure and transformers are not overloaded) while satisfying the needs and preferences of the highly heterogeneous population of EVs each with their individual battery capacity, charging speeds, and usage pattern.

More positively, EVs will also be a key resource in the demand-side management systems discussed previously. In such systems, the ability to defer demand to times when renewable energy is more plentiful is essential, and currently, this is only possible with subset of electrical loads that are not required to have immediate effect (e.g., washing machines or dishwashers). However, the ability to store energy within large batteries allows any electrical load to be shifted, and we are likely to first see energy from electric vehicle batteries support the shifting of loads within their owners' home (vehicle-to-home or V2H), and then to providing energy back to the grid itself (V2G) [26, 25].⁶ Hence, while the impact of scheduling loads in the home on the user's lifestyle may be minimised through the use of the EV battery, the scheduling of the battery charging and discharging cycles will need to ensure there is sufficient capacity to satisfy the loads in the home, and the travel needs of the vehicle's owner, while minimising the cost of electricity used. Moreover, this schedule will need to be optimised for, and adapt to, the changing needs of the vehicle owner, the (real-time) price paid for feeding back to the grid, as well as the battery capacity and efficiency. Hence, such optimisations will also require learning algorithms to predict the pattern of use of the vehicle, and also the demand of the home.

Addressing these challenges requires intelligent systems that can fully automate the charging and discharging of these vehicles, whilst taking account of the current and future availability of the renewable generation, and being aware of the local constraints of the distribution network. Recent work has begun to address these challenges with online mechanism design being used to elicit users' travel requirements (i.e., the amount of charge required and the time at which the EV is needed) and schedule the charging of their vehicles [17], and suggestions to apply peak and dynamic pricing to shift demand across a city [25]. These mechanisms are likely to work and be of social value (i.e., not impede the daily activities of the vehicle owners) only if they minimise waiting (charging) times for consumers and never leave consumers stranded. As such, these systems will have to draw on diverse sources of information, such as distribution network load information (e.g., load on the lines, number of EVs connected at various positions and prices at different charge/discharge points), traffic information from road cameras, and geolocation services such as Google Latitude (<http://latitude.google.com>) or Facebook Places (<http://www.facebook.com/places>) which contain rich in-

⁶In addition to providing energy, the vehicles may also be able to provide regulation services to the grid to stabilise both the voltage and frequency of electricity[31].

formation that can be mined to predict future movements of consumers to specific locations and, hence, likely bottlenecks on specific lines and transformers in the system. Systems that can optimise the charging cycle of an EV by making sense of such a wide range of heterogeneous information sources are likely to play a key role in ensuring EVs are seamlessly integrated into the smart grid.

Thus, against this background, we identify the key AI challenges in the deployment of EVs in the smart grid as follows:

- Predicting an individual user’s EV charging needs based on data about her daily activities and travel needs.
- Predicting aggregate EV charging demands at different points in the network given the continuous movement of EVs, the available charge in their batteries, and the social activities their users engage in.
- Designing decentralised control mechanisms that coordinate the movement of EVs (each with different battery capacities and charging speeds) to different charge points by providing incentives to consumers to do so. The aim being to maintain secure flows on the grid and ensure that transformers do not trip due to excess demand.
- Designing algorithms to optimise the charging cycles of EVs to satisfy the predicted needs of the user (to shift loads or to travel) while maximising the profits generated from participating in V2G sessions.

4. VIRTUAL POWER PLANTS

As larger numbers of actors (e.g., EVs, homes, or renewable energy providers) in the smart grid communicate and coordinate with each other to control demand at different points in the network (e.g., using demand-side management to ensure that demand is able to follow the supply of renewable energy, and EV discharging to the grid to cope with excess demand), it will be important to harness synergies that exist between them to improve the efficiency of the grid (e.g., EV discharging to satisfy demand at times when demand-side management techniques cannot shift enough usage to later times). To this end, the concept of a virtual power plant (VPP) [2] has been proposed to capture the notion of a number of actors, coming together to sell electricity, as an aggregate.⁷ However, several challenges arise in the formation and management of VPPs that coordinate a number of heterogeneous actors (e.g., EVs or renewable energy providers) to maximise the amount of energy delivered in the system while minimising the costs and uncertainties in doing so. In particular, these individual actors need to be able to come to an agreement in technical (i.e., how they coordinate their consumption or production patterns) and economical (i.e., how they share the profits generated by the VPP) terms in order to maximise the value of the set of energy services (i.e., providing electricity, storing electricity, or shifting demand) they provide as a VPP.

Now, the process of forming VPPs at a technical level means that the individual actors need to synchronise the largely

⁷The term virtual power plant is also used to describe companies, which may not have any generation capacity and that simply buy generation capacity from a generator. We do not deal with such VPPs here.

heterogeneous services they provide within the VPP in an agile fashion so as to meet the requirements of the contracts they make with their customers. In particular, individual actors need to estimate the impact of their individual production (or demand reduction) on the aggregate performance of the VPP, and communicate and optimise the joint actions taken to meet the VPPs’ objectives (i.e., satisfy demand). These technical arrangements may need to be specified on a daily, and even on an hourly basis to maximise the profits of the individual actors. This is because, if some actors can only produce energy at specific times of the day (e.g., PVs generate energy during the day and tidal energy may be available at night), they will want to choose those partners they can complement better at those times (e.g., a PV farm and a tidal generator may generate energy out of phase with each other and hence be highly complementary, while wind energy providers whose turbines are located in the same region will generate energy at the same time and hence be less complementary). In turn, if new actors become better partners due to changes in the environment (e.g., more wind blows at night resulting in higher predicted wind energy production than tidal or more EVs converge to a specific region due to a social event, resulting in more storage being available), then some of them might decide to leave their current VPP and form a new one (e.g., PV owners may be better off storing their excess energy during the day in the EVs to be able to supply at night rather than collaborate with a tidal energy provider). Given the scale and dynamism of this optimisation problem, it will be important to design decentralised coordination algorithms and strategies that allow individual VPP participants to come to the most efficient arrangements within a reasonable time. Moreover, they will need to ensure such arrangements do not overload the local distribution networks, in which they are connected. Given this, and the restrictions imposed by the network operator due to possible network congestion, the VPP may further have to re-optimize individual members’ operations. Typically, such optimisations would have to be done while being confronted with uncertainty about the individual members’ generation and consumption capacity.

The negotiation of technical arrangements needs to take into account that each potential member of a VPP is typically motivated to maximise its own profit, even though, as a group they compete against other actors (individuals, VPPs or large power stations) in the system to maximise the group’s profits. Hence, it is in each actor’s interest to take actions that will cost it the least while maximising its share of the profits obtained by the VPP operations as a whole. This leaves some room for any individual resource to manipulate what it reveals as its predicted capability (i.e., production, demand-response, or storage ability) as opposed to what it actually delivers on the day. For example, given their uncertainty about their production, some resources may prefer to understate their predicted production profile in case they get penalised by the group for under producing. Alternatively, some resources may prefer to overstate their predicted production in the case that penalties for under producing are not significant, and doing so increases their share of the profits. Such strategic considerations highlight the need to capture the provenance of decision made by the VPP, such that it is possible to track and verify the individual actions, reports, and resulting rewards

of each VPP member. The amount of provenance information this will generate will require efficient frameworks and mechanisms to represent, store, audit, and share it. Building upon provenance information it may then be possible to model the trustworthiness of individual VPP members through trust and reputation mechanisms similar to those used in online marketplaces such as eBay or Amazon for example [29]. These mechanisms would, in turn, need to be designed to ensure they are robust to wrong or manipulative reports so that security measures can then be taken to ensure that those actors with low trust do not cause significant disruption to the network in case they do not fulfil their part of the VPPs' operations.

Assuming trust and reputation mechanisms can render VPPs reliable, it is important to ensure that the negotiations that individual energy providers engage in, converge in such a way that the most efficient VPPs (i.e., generating the maximum social welfare) are most effectively formed (i.e., in minimum time and with minimum communication costs) in the system [11]. Here, convergence is achieved when all the members of the VPP are satisfied with their share of the profits generated. The strategic and computational aspects of such negotiation processes are typically studied within multi-agent systems using tools such as cooperative game theory [4] to partition the profits of groups among their members and combinatorial optimisation algorithms to partition actors into the most efficient groupings for the system respectively [27]. However, the VPP formation process presents a number of unique challenges for AI research. In particular, given that all actors are connected in a network where flows are limited on each line, the actions (energy production or consumption) taken by each actor or VPP restricts the actions (to different degrees) of all VPPs in the system. Hence, the formation of each VPP can have significant externalities (e.g., the flows created by one VPP can congest some lines, which, in turn, may prevent other VPPs from using energy sources or providing energy to consumers at the nodes connected to those lines). Moreover, the fact that each VPP compounds the uncertainty in production of each member (e.g., due to uncertainty in the weather forecast or demand-side managed consumption) renders the VPP formation process highly stochastic.

All these issues will require the definition of computationally efficient search algorithms to allocate the payoffs to individual members of VPPs (as defined by game-theoretic solution concepts), while taking into account uncertainty in defining the relative contributions of each member to the aggregate performance (i.e., mainly the profits generated) of the VPP. Moreover, given that different coalitions may be formed over time, an energy provider will choose its membership of coalitions in such a way as to maximise its revenues in the long run. This makes the search for efficient payoff allocations exponentially harder since it extends the search space to include future possible coalitions (and their expected returns) as well as present ones. Initial work in applying multi-agent systems approaches to the VPP formation process include [5] which provides solutions to the formation of VPPs of wind turbines with uncertain production and [13] which provide an agent-based framework for VPP formation. These approaches, however, are still at a preliminary stage.

To advance the state of the art in this domain, the following key AI challenges still need to be addressed:

- Designing agent-based models of different VPP actors and processes in order to capture the complexity of the technical arrangements needed to form and manage VPPs.
- Distributed combinatorial optimisation of the technical arrangements of demand-side management, V2G sessions, and micro-generation, to maximise rewards.
- Designing online mechanisms to form statistically correct trust measures for energy providers and automatically capture, track, and reason about the provenance of information revealed by energy providers to form VPPs.
- Designing search algorithms and negotiation mechanisms for individual actors to agree on which VPP to form at different points in time and how to share the profits, using computationally efficient game-theoretic solution concepts, of a VPP given uncertainty in their performance, trust in their revealed capabilities, and changing weather and demand patterns.

5. ENERGY PROSUMERS

Our discussion, so far, has highlighted the significant heterogeneity of the large numbers of renewable energy resources in the smart grid and the complexity of the interactions between them and consumers. When taken altogether, this will necessitate significant changes in the way energy is bought and sold. In particular, this is set against the current operation of the grid where, in many countries (e.g., the US, UK, and in many parts of the EU), the electricity market is deregulated, such that large generators (located far from the point of use) trade directly with retailers who then sell the electricity on to consumers through fixed contracts and tariffs [19, 35]. In these countries, electricity is traded in forward and futures markets on a long-term ahead basis (weeks, months, seasons and even years) and on day-ahead spot markets through a range of different contracts (e.g., baseload, off-peak or half-hourly contracts). Any real-time excess or shortfall in supply and demand (with respect to contracted volume) is settled in the balancing market (also termed the settlement process) where the price to buy and sell electricity is typically set by the market maker rather than being based on the direct matching between bids and offers in the day-ahead market.

In contrast, in the smart grid, market operations will have to adjust to a much larger number of heterogeneous entities, distributed throughout the network (closer to the point of use of electricity), trading much smaller amounts of energy. Indeed, the widespread adoption of renewable generation at the level of individual homes and businesses will lead to the creation of markets composed of many millions of prosumers who both produce and consume energy [14]. Given this, while some prosumers may try to find an agreement with other prosumers to form VPPs (and resort to cooperative game-theoretic solutions as discussed in Section 4), many will directly trade in the electricity market (where the game-theoretic considerations are purely non-cooperative). Hence, compared to typical consumers who are mainly concerned about optimising their electricity usage and who are

typically agnostic to the real-time conditions on the electricity market, prosumers will need to optimise *both* their production and consumption of energy in order to make trading decisions in real-time, through internet-based interfaces to spot or forward markets, so that they maximise the profits they can make by buying (to consume or store) and selling energy (either energy that they generate, or have stored earlier). By making their own localised trading decisions, prosumers may reduce the inefficiencies (added costs for end users and lower margins for generators) resulting from retailers hedging their energy purchases to minimise their exposure to risk (in the balancing market) and selling fixed long-term contracts to their consumers at high costs.

To do so, however, means that prosumers will need to be endowed with effective trading strategies that can cope with uncertainty in the market. To minimise this uncertainty, they will need to be informed by predictions of their own demand (that may vary according to their needs and social activities) and generation capacity (e.g., using weather forecasts or their EV usage needs), as well as the future price of electricity on the market. Given that these trading decisions may need to be taken in real-time, these predictions will also need to be generated in real-time, and furthermore, to ensure users understand the life-style or operational implications of, and agree to, autonomously chosen trading decisions, human-computer interaction mechanisms will have to be designed to ensure that large numbers of users trust and participate in these markets.

Essentially, as more prosumers populate the market, electricity will become a commodity with similar properties to those traded on stock markets. Given this, prosumers will be able to speculate in markets, buying and selling not simply to consume or supply electricity, but also to profit. However, while speculation may help make the market more efficient, it may also adversely impact on the operation of the grid, if the traded flows do not actually satisfy the physical constraints of the distribution network. Potential solutions point to the application of regulatory measures to reduce speculation and more importantly, to congestion pricing mechanisms [39] within the distribution network, similar to the locational-based pricing that is used within the transmission network in many parts of the US [35]. In such mechanisms, prices vary geographically throughout the network to ensure that the flows of electricity within it do not exceed the limits of any of the transmission lines. To ensure these mechanisms do guarantee an efficient system it will be important to study the equilibrium conditions (e.g., market efficiency, loads on transmission lines) resulting from the application of these congestion prices against significantly heterogeneous populations of prosumers.

In summary, the AI challenges involved in endowing prosumers with the intelligence to trade in electricity markets whilst ensuring safe network flows include:

- Developing computationally efficient learning algorithms that can accurately predict both the prosumers' consumption and generation profiles (instead of only the usage profile for a consumer) as well as the price of electricity in real-time in order to inform profitable trading decisions.

- Developing autonomous trading agents that can use such predictions to maximise their profit in the electricity market, and efficient algorithms to marry congestion management with market operation in distribution networks while guaranteeing good equilibrium conditions in the system.
- Developing human-agent interaction mechanisms, to allow prosumers to guide their agents trading decisions, that take into account the prosumers' daily constraints and preferences to consume or produce energy.

6. SELF-HEALING NETWORKS

So far, we have discussed a number of ways in which the electricity flows are likely to become both more unpredictable and bidirectional in the smart grid. This will result in a greater need for decentralised control strategies given the sheer numbers of active entities embedded in the system. While this renders fault-correction mechanisms in the network even more complex, the intelligence on which these active entities rely to make their consumption or generation decisions, could also be used to naturally distribute (and hence make more robust) the decision making needed to apply self-healing strategies on the network when faults occur. Generally speaking, faults may arise either because lines become overloaded or because of old infrastructure becoming more prone to failure. To prevent such faults and remedy them, network operators already rely on a number of intelligent systems at the transmission network level. Traditionally, this is achieved with the help of automatic voltage regulators and using supervisory control and data acquisition systems [6] with phasor measure units⁸ for situational awareness. Using such systems, active network management [21] techniques can help to automatically reconfigure the network and send control signals to individual generators to increase generation or to pre-contracted loads to reduce their consumption [7]. By endowing individual components on the network with the intelligence to apply these techniques, they can automatically correct faults as and when they occur and therefore let the network self-heal.

Extending these techniques to the management of the distribution network where large numbers of prosumers will operate, will require a much larger number of phasor measure units to be deployed, both because the distribution network contains many more nodes, but also because the heterogeneity of the prosumers within it means that network conditions are likely to vary more rapidly, necessitating accurate and timely monitoring and control. Fully instrumenting such networks is likely to be too expensive, and thus, there is a clear need for the development of state estimation systems that do not need to have every node in the network monitored. More importantly, we will need systems that can, using information gleaned from across the grid, learn correlations between state parameters at different nodes to provide accurate and robust estimates of the system state. The vast amount of data generated from multiple actors and sensors, and the micro-second level measurements being made, will present formidable computational challenges in trying to estimate or predict the future state of the system.

⁸Phasor measurement units measure both magnitudes and phase angles of voltages and currents within the network, and are used to assess the state of a power system in real-time.

Now, if accurate information about the network can be obtained, active network management techniques, supported by distributed intelligence in the network, could help recover from faults faster than previously possible. For example, if voltages tend to drift in some parts of the network, automatic actions on transformers may be taken to re-establish the correct voltage levels, or assistance may be requested from EVs that are currently plugged into the network [38]. Furthermore, if faults are detected in one part of the network, that part of the system could be disconnected, leaving other independent parts running separated (i.e., effectively ‘islanded’) provided they can sustain the balance between supply and demand (e.g., using demand-side management). This could eventually avoid rolling blackouts or even help recover from those blackouts that do happen.

To build such self-healing mechanisms, however, will require that all these actors can communicate their action space (e.g., limits on voltage regulation, generation capacity, demand reduction ability) and agree on joint actions to implement islanding strategies. Given the uncertainty that permeates the actions of some of these entities (e.g., weather patterns that affect generation or social activities that affect the movement of EVs), it will be important to predict the impact of such uncertainty on the joint actions chosen to avoid electing those that may result in cascading failures in the worst case. Moreover, given the individual preferences of all actors involved (e.g., to consume electricity for specific activities or to sell electricity to maximise profits) these joint actions may need to be negotiated rapidly among them to ensure they end up in an agreement all parties commit to [23]. Initial approaches aiming to achieve this level of coordination express the problem as centralised (constrained) optimisation problems that can be solved using (non) linear programming tools [7]. Clearly, centralising active network management involving potentially thousands of different types of actors, each with their own energy generation and production requirements is unlikely to scale very well in both the communication and computation costs it incurs. Hence, more scaleable decentralised planning approaches that rely on short range communication between individual actors (e.g., distribution network nodes, consumers, and EVs) will be needed [38] or [32].

Hence, we summarise the AI challenges of self-healing mechanisms as follows:

- Designing computationally efficient state estimation algorithms that can predict voltage and phase information at different nodes in the (partially observable) distribution network, in real-time, given the prosumers’ current and predicted energy demand and supply.
- Enabling distributed coordination of automatic voltage regulators and energy providers and consumers for voltage control and balancing demand and supply during recovery from faults.
- Automating distributed active network management strategies given the uncertainty (either because they cannot be accurately measured or there is incomplete information about certain nodes) about demand and supply at different points in the network.

7. CONCLUSIONS

There is a significant drive within the developed world to reduce our reliance on fossil fuels and move to a low-carbon economy in order to guarantee energy security and mitigate the impact of energy use on the environment. This transition requires a fundamental re-think and re-engineering of the electricity grid. The ensuing smart grid must be able to make efficient use of intermittent renewable energy sources and supply the additional electricity required by electric vehicles. Doing so, will require extensive use of demand-side management and virtual power plants to balance supply and demand. It will also see large numbers of prosumers, buying and selling electricity in real-time, whilst automated network control algorithms maintain the safe operation of the grid, and allow it to self-heal when something does go wrong.

The automation, information exchange, and distributed intelligence needed to deliver such technologies creates many new challenges for the AI communities investigating machine learning, search, distributed control, and optimisation. In this paper, we have enumerated what we believe to be the main challenges that, if met, will allow the full potential of the smart grid to be realised. Our claims build upon an extensive survey of the state of the art that goes beyond the papers cited and includes a large number of references (spanning technical papers, books, and policy documents relating to the deployment of specific smart grid technologies and evaluations of these) provided in the **online appendix**. In particular, we have highlighted the key issues in learning and predicting demand or supply at various points in the network given the variety of demand control mechanisms (e.g., demand-side management and EV charging) and energy sources, each with different degrees of uncertainty in their production capability (e.g., VPPs or renewable energy sources). Moreover, we showed that the automated decentralised coordination between such entities (to balance demand and supply while ensuring flows on the network are always secure) will need to factor in both the individual properties of all actors (e.g., EVs with different batteries, different types of renewable energy sources, users with their own understandings of trading decisions and their agents’ decisions) involved and the incentives given to them to behave in certain ways (e.g., consumers shifting demand due to real-time pricing, or VPPs sharing profits equitably). Building upon this, we also discussed some initial attempts at solving them within the various sub-areas of the smart grid. Cutting across these various challenges are the issues of human-computer interaction, heterogeneity, dynamism, and uncertainty that are an intrinsic part of decision making and acting in the smart grid. By dealing effectively with these factors, we believe it will be possible for future generations to rely on their energy systems to deliver electricity efficiently, safely, and reliably.

Finally, we note that many of the issues present within the smart grid also arise within other domains such as water distribution, transportation, and telecommunication networks where large numbers of heterogeneous entities act and interact in a similar fashion to those within the grid. Hence, there is potential to transfer technologies across these domains and also address broader issues that affect the sustainability of such systems in a unified manner, such as cyber-security and the ethics of delegating human decision making to intelligent systems.

Acknowledgements

This work was done as part of the iDEaS project.⁹

8. REFERENCES

- [1] K. Aleklett, M. Höök, K. Jakobsson, M. Lardelli, S. Snowden, and B. Söderbergh. The Peak of the Oil Age-Analyzing the world oil production Reference Scenario in World Energy Outlook 2008. *Energy Policy*, 38(3):1398–1414, 2010.
- [2] S. Awerbuch and A. M. Preston. *The virtual utility : accounting, technology and competitive aspects of the emerging industry*. Kluwer, Boston, 1997.
- [3] G. Binczewski. The energy crisis and the aluminum industry: Can we learn from history? *Journal of the Minerals, Metals and Materials Society*, 54(2):23–29, 2002.
- [4] G. Chalkiadakis and C. Boutilier. Sequentially optimal repeated coalition formation under uncertainty. *Autonomous Agents and Multi-Agent Systems*, pages 1–44, 2010.
- [5] G. Chalkiadakis, V. Robu, R. Kota, A. Rogers, and N. R. Jennings. Cooperatives of distributed energy resources for efficient virtual power plants. In *Proc. of the Tenth Intl. Conf. on Autonomous Agents and Multiagent Systems*, pages 787–794, May 2011.
- [6] S. Chowdhury, S. Chowdhury, and P. Crossley. *Microgrids and Active Distribution Networks*. Institution of Engineering and Technology (IET), 2009.
- [7] E. Davidson, S. McArthur, C. Yuen, and M. Larsson. Aura-nms: Towards the delivery of smarter distribution networks through the application of multi-agent systems technology. In *IEEE Power and Energy Society General Meeting*, pages 1–6, 2008.
- [8] DECC. *The Climate Change Act 2008 Impact Assessment*. DECC, 2009.
- [9] K. S. Deffeyes. *Hubbert’s peak: the impending world oil shortage*. Princeton Univ. Press, 2008.
- [10] M. Deindl, C. Block, R. Vahidov, and D. Neumann. Load shifting agents for automated demand side management in micro energy grids. In *Proc. of the Second IEEE Intl. Conf. on Self-Adaptive and Self-Organizing Systems*, pages 487–488, 2008.
- [11] G. Demange and M. Wooders. *Group formation in economics: networks, clubs and coalitions*. Cambridge Univ. Press, 2005.
- [12] U. S. Department-Of-Energy. Grid 2030: A National Vision For Electricity’s Second 100 Years. Tech. report, Department of Energy, 2003.
- [13] A. Dimeas and N. Hatziairgiou. Agent based control of virtual power plants. In *Proc. of the Intl. Conf. on Intelligent Systems Applications to Power Systems*, pages 1–6, 2007.
- [14] EU SmartGrid Technology Platform. Vision and strategy for europe’s electricity networks of the future. Tech. report, European Union, 2006.
- [15] T. Friedman. *Hot, flat, and crowded: Why we need a green revolution—and how it can renew America*. APS, 2008.
- [16] J. Froehlich, L. Findlater, and J. Landay. The design of eco-feedback technology. In *Proc. of the 28th Intl. Conf. on Human Factors in Computing Systems*, pages 1999–2008. ACM, 2010.
- [17] E. Gerding, V. Robu, S. Stein, D. Parkes, A. Rogers, and N. R. Jennings. Online mechanism design for electric vehicle charging. In *Proc. of the Tenth Intl. Joint Conf. on Autonomous Agents and Multi-Agent Systems*, pages 811–818, May 2011.
- [18] R. C. Green, L. Wang, and M. Alam. The impact of plug-in hybrid electric vehicles on distribution networks: A review and outlook. *Renewable and Sustainable Energy Reviews*, 15(1):544–553, 2011.
- [19] C. Harris. *Electricity Markets: Pricing, Structures, and Economics*. Wiley, 2005.
- [20] IEA. World energy outlook 2009 fact sheet. Tech. report, Intl. Energy Agency, Paris, 2009.
- [21] R. MacDonald, G. Ault, and R. Currie. Deployment of active network management technologies in the UK and their impact on the planning and design of distribution networks. *SmartGrids for Distribution*, pages 1–4, 2009.
- [22] D. MacKay. *Sustainable energy: without the hot air*. UIT, Cambridge, 2009.
- [23] J. McDonald. Adaptive intelligent power systems: Active distribution networks. *Energy Policy*, 36(12):4346–4351, 2008. Foresight Sustainable Energy Management and the Built Environment Project.
- [24] W. Mert, J. Suschek-Berger, and W. Tritthart. Consumer acceptance of smart appliances. Tech. report, EIE project—Smart Domestic Appliances in Sustainable Energy Systems (Smart–A), 2008.
- [25] W. Mitchell, C. Borroni-Bird, and L. Burns. *Reinventing the Automobile*. MIT Press, 2010.
- [26] RAE. Electric vehicles: charged with the potential. Tech. report, The Royal Academy of Engineering, 2010.
- [27] T. Rahwan, S. D. Ramchurn, N. R. Jennings, and A. Giovannucci. An anytime algorithm for optimal coalition structure generation. *Journal of Artif. Intel. Research*, 34:521–567, April 2009.
- [28] S. Ramchurn, P. Vytelingum, A. Rogers, and N. Jennings. Agent-based homeostatic control for green energy in the smart grid. *ACM Transactions on Intelligent Systems and Technology*, 2(4), May 2011.
- [29] S. D. Ramchurn, T. Huynh, and N. R. Jennings. Trust in multiagent systems. *The Knowledge Engineering Review*, 19(1):1–25, 2004.
- [30] S. D. Ramchurn, P. Vytelingum, A. Rogers, and N. R. Jennings. Agent-based control for decentralised demand side management in the smart grid. In *Proc. of the Tenth Intl. Conf. on Autonomous Agents and Multiagent Systems*, pages 5–12, May 2011.
- [31] P. Ribeiro, B. Johnson, M. Crow, A. Arsoy, and Y. Liu. Energy storage systems for advanced power applications. *Proc. of the IEEE*, 89(12):1744–1756, 2001.
- [32] A. Rogers, A. Farinelli, R. Stranders, and N. R. Jennings. Bounded approximate decentralised coordination via the max-sum algorithm. *Artif. Intel.*, 175(2):730–759, 2011.
- [33] P. Scerri, D. Pynadath, and M. Tambe. Towards adjustable autonomy for the real world. *Journal of Artif. Intel. Research*, 17(1):171–228, 2002.
- [34] F. Schweppe, B. Daryanian, and R. Tabors. Algorithms for a spot price responding residential load controller. *Power Engineering Review, IEEE*, 9(5):49–50, 1989.
- [35] F. C. Schweppe, M. C. Caramanis, R. O. Tabors, and R. E. Bohn. *Spot Pricing of Electricity*. Kluwer Academic Publishers, 1988.
- [36] G. Strbac. Demand side management: Benefits and challenges. *Energy Policy*, 36(12):4419–4426, 2008.
- [37] V. Sundramoorthy, G. Cooper, N. Linge, and Q. Liu. Domesticating energy-monitoring systems: Challenges and design concerns. *IEEE Pervasive Computing*, 10:20–27, 2011.
- [38] P. Vovos, A. Kiprakis, A. Wallace, and G. Harrison. Centralized and distributed voltage control: Impact on distributed generation penetration. *Power Systems, IEEE Transactions on*, 22(1):476–483, 2007.
- [39] P. Vytelingum, T. D. Voice, S. D. Ramchurn, A. Rogers, and N. R. Jennings. Agent-based micro-storage management for the smart grid. In *Proc. of the Ninth Intl. Conf. on Autonomous Agents And MultiAgent Systems*, pages 39–46, May 2010.
- [40] F. Ygge, J. M. Akkermans, A. Andersson, M. Krejic, and E. Boertjes. The HOMEBOTS System and Field Test: A Multi-Commodity Market for Predictive Power Load Management. In *Proc. of the Fourth Intl. Conf. on the Practical Application of Intelligent Agents and Multi-Agent Technology*, volume 1, pages 363–382, 1999.

⁹<http://www.ideasproject.info>.