Embedded Intelligence for Electrical Network Operation and Control

Victoria M. Catterson, Euan M. Davidson, and Stephen D.J. McArthur,
University of Strathclyde

Integrating multiple types of intelligent, multiagent data analysis within a smart grid can pave the way for flexible, extensible, and robust solutions to power network management.

Within North America, Europe, and other territories, there is political will to adapt and augment existing electrical networks to meet our future energy needs. Such *smart grids* involve changes broadly aimed to increase efficiencies while lowering costs and energy usage. Although one

method is to focus on changing consumer behavior using metering and smart appliances,[1] another considers the network itself and favors a new approach to controlling and operating the existing grid infrastructure (see http://www.smartgrids.eu).

Smarter operation of the existing network naturally suggests the use of intelligent systems techniques to perform or enhance many of the activities currently performed by engineers. Traditionally, engineers view electrical networks as either transmission or distribution networks. Although both types are managed from a centralized control room with the aim of balancing supply and demand, certain features differentiate the two. Namely, transmission networks are characterized by

- high-voltage, meshed networks used for bulk transfer of power;
- large-scale generators (such as coal-fired, gas, and nuclear stations), where the network control engineers are responsible for the dispatch of generation;
- supervisory control and data acquisition (SCADA) systems that give engineers

visibility of voltages and currents around the network and provide remote control of devices such as circuit breakers and telecontrolled switches; and
- engineers reconfiguring the network as needed to handle various contingencies.

In contrast, distribution networks traditionally are characterized by

- lower-voltage, radial networks designed to deliver energy from the transmission network to customers (in homes, businesses, and industry);
- little or no generation connected, making dispatch of generation impossible;
- SCADA systems that offer only visibility of the network, as many devices are manually operated; and
- engineers rarely reconfiguring the network to handle emergency situations.

Various smart-grid visions presuppose a future in which these passive, radial distribution networks with unidirectional power flows evolve into networks with increasing levels of embedded generation, demand response,

monitoring, automation, and control. This radically alters the network's operational requirements, challenging engineers to design active control for response to system events while offering more data for automated analysis of network and asset health.

In this article, we present industrial examples of intelligent systems providing the smart functionality for these tasks. Specifically, we discuss autonomous control at the distribution level, automated analysis of data from highly observable power systems, and the use of multiagent systems technology for distributing and integrating smart functionality. We show how using intelligent systems is key to addressing these problems in ways that are flexible, robust, and understandable to engineers, supporting the delivery of a more efficient grid.

Autonomous Control of a Distribution Network

Connecting small-scale generation to distribution networks leads to a raft of problems collectively termed *congestion*. This includes rises in voltage (caused by the reverse power flow back to transmission), current overloads at network bottlenecks (due to the generators producing currents the network is not rated to handle), or increased load associated with new customers or electric vehicles.

Traditionally, engineers designed or adapted networks to avoid congestion by building extra lines or replacing transformers. However, such network reinforcements can be expensive. An alternative is using an active network management (ANM) strategy to provide automated control to deal with overload conditions. An ANM scheme can include the following control tasks: automatic network reconfiguration to keep as many customers as possible on supply, control of generators to keep power flows within

limits and avoid overloading of lines, coordinated control of voltage regulation measures, and control of the network to minimize losses.

For all these functions, network operators require solutions with the following properties:

- *Flexibility.* Engineers desire solutions that can be deployed in various situations—that is, in different network topologies with different control measures.
- *Extensibility.* Networks change; new generators can be added and old ones removed. Regulations also change, opening up new services and possibilities for network control. ANM systems can be extended without replacing the entire control system.
- *Graceful degradation.* Measurement error and failure of the communication systems underpinning ANM schemes are inevitable. Thus, performance must degrade gracefully when a scheme is presented with bad data or communications failures occur.

Figure 1 illustrates an example utility that desired an ANM solution to manage power flows in the electrical network. The figure shows two 132-kV lines feeding a distribution network through two transformers, rated at 45 megavolt ampere (MVA). The distribution network is at 33 kV, with almost 85 MW of installed generation (generators 1, 2, 6, 7, and 8). The minimum local load in this network is 10 MW, meaning that certain conditions see an export of 75 MW from the 33-kV network through the transformers (85 MW of generation, 10-MW load).

In this scenario, if one of the two 45-MVA transformers is removed from service, the remaining transformer will be overloaded. To handle

such eventualities, bespoke systems for detecting overloads were previously installed on the network. One system monitors for the loss of one transformer to mitigate such a scenario. If a transformer is lost, the current on lines 1 and 3 is monitored, and too high a current results in a trip signal being sent to the hydro generator. Removing the hydro plant from the network reduces the power being exported through the remaining transformer, and thus power flow is kept within limits.

Due to the incremental building of wind farms in this area, the number of bespoke systems in place has increased over time. There are now nine special control schemes in place, each monitoring for a different scenario requiring different control. It is hard for engineers to manage this area of network because the complexity of the automated systems leads to control actions being taken that are not straightforward and understandable.

Engineers wanted a solution that could automatically handle power-flow management in an explainable way and offer the option of manual power-flow management, when desired.[2] It should be possible to apply this technique to multiple distribution networks—reasoning generically on a model of the network, rather than being tuned to suit only the particular installation topology.

Constraint programming technology can help meet these needs.[3] Power-flow management was formulated as a finite discrete domain constraint satisfaction problem (CSP),[2] where controllable generators constitute the problem variables, with discrete domains for the control signals that limit their output. For our example network, we chose the domain of all wind farms, $D = \{1, 0.75, 0.5, 0.25, 0\}$, and the hydrogenerator is $D = \{1, 0\}$.
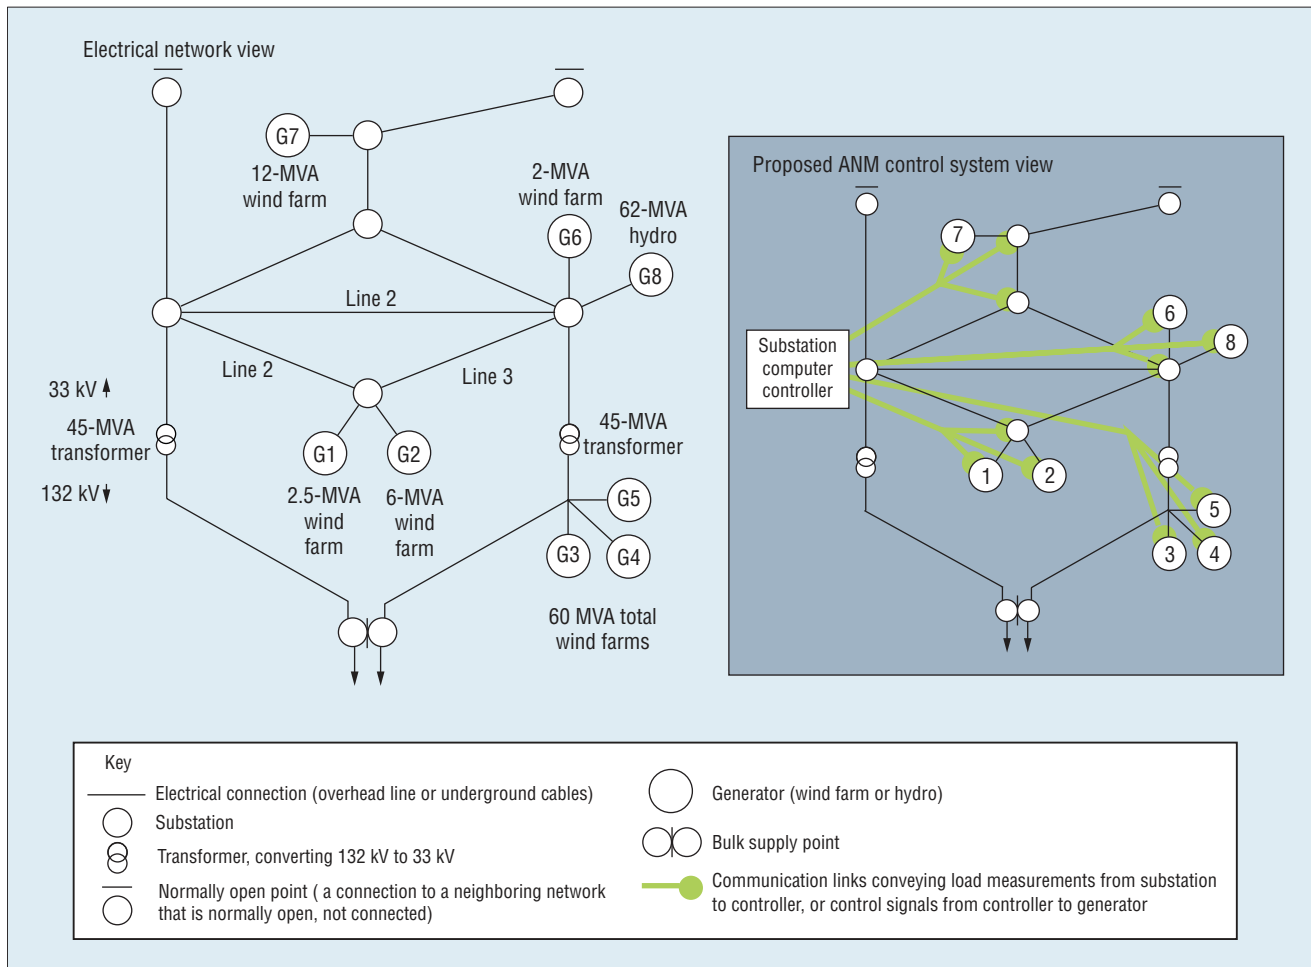
Figure 1. Example distribution network. This network automatically handles power-flow management in an explainable way and offers the option of manual power-flow management. In this case, the off-the-shelf constraint satisfaction problem (CSP) solver running on a substation computer could generate timely solutions for all normal network conditions.

Three distinct types of constraints apply to power-flow management. First, the flow of power around the network must not exceed the thermal limits of any equipment. This technical constraint can be checked using a network model and a load-flow engine, which engineers use to manually check thermal constraints.

Second, contractual constraints determine which generators have priority access to connect. Under current contracts, generators either have a firm connection agreement, which lets them generate whenever and however much they like, or they have a nonfirm agreement, in which they are removed from the network in a last-in-first-off (LIFO) order. To connect more wind farms to the network while avoiding the cost of network reinforcement, many of the generators were offered nonfirm connections. For the network in Figure 1, if an overload occurs, the most recently-built wind farm 5 is always first to be removed.

The LIFO regime was introduced in the interest of fairness, so the generators connecting earlier would not see their access rights reduced if additional generators were to connect in the future. However, because of different limits on items of equipment, in some situations generator 6 could reduce its output by 5 MW to alleviate the congestion, whereas the contractually obligated reduction of generator 5 must be 10 MW. In this situation, it would be in generator 5's best interest to negotiate a bilateral contract with generator 6, paying it to reduce by 5 MW in order to continue outputting (and selling) an extra 10 MW. Such bilateral contracts are not yet allowed, but the increase in wind penetration and subsequent network congestion issues mean regulation is under scrutiny and might change in the coming years. With the constraint programming approach to power-flow management, the contractual constraints can be updated easily with changes in regulation.

The final type of constraint is the preference for maximizing access to the network for distributed generators. Thermal overloads can be removed

by tripping off all generation, but it is preferable to maintain wind generation whenever possible.

With this formulation of the problem, an off-the-shelf CSP solver running on a substation computer (see Figure 1) could generate solutions in a worst-case time of 9.7 seconds and under 2 seconds for most normal network conditions. This is well within the timescales required for a control system to remove a thermal overload and is therefore a suitable approach to power-flow management. It is also extensible, by updating constraints if there is a change in regime, and generic to all networks, by changing the network model for the load flow. Crucially, engineers can understand the control actions this system proposes because the load flows before and after action is taken show the thermal overload being removed. Graceful degradation comes from the CSP solver offering a ranked set of solutions. Should the preferred solution fail due to model error, measurement error, or failure of communications, the controller can try implementing the next solution.

This system is currently in the preliminary stages of field trials with a UK utility. Software is being deployed on a substation computer in a 33-kV/11-kV primary substation for power-flow management of an 11-kV feeder with 2 MW of embedded generation. Initially, an open-loop system will generate and record control actions for removing thermal overloads. After successful trials, closed-loop operation should follow, in which control actions are automatically taken when an overload is detected.

Another facet of ANM is automatically restoring supply after a fault. In contrast to power-flow management, this has been the subject of much research[4] and has become a benchmark problem within AI planning.[5] Although engineers are involved in network control, the protection system is fully automated.

Post-Fault Analysis

The protection system is responsible for detecting faults (such as short circuits) and disconnecting faulty or faulted equipment from the network by opening the appropriate circuit breakers. This must happen within milliseconds up to tens of seconds to limit equipment damage and ensure that the power system remains stable.

After a fault has occurred on the network, engineers are interested in its type and location, how much of the network was affected, whether any areas are still without supply, and whether the protection system operated correctly to clear the fault. These questions must be answered by analyzing the data generated by the fault, but there are different sets of data available. One utility, which arguably operates one of the most heavily monitored networks in Europe, performs post-fault analysis on four data sets:[6]

- SCADA alarms that record automatic actions (such as high currents detected, main protection operating, and circuit-breakers triggering);
- digital fault recorder (DFR) data, which records high-resolution current and voltage traces from a circuit when triggered by an incident;
- traveling-wave fault locator data, which indicates how far along a line the fault is; and
- circuit-breaker trip-coil current traces, indicating when the breaker was triggered and how long it took to operate.

Frequently, engineers manually analyze these data sets. In a typical year, a utility will see 3 to 6 million SCADA alarms and 20,000 fault records, but during storm conditions, when many network incidents occur simultaneously, these figures can reach 20,000 SCADA alarms and 2,000 fault records in a 24-hour period. This rate is overwhelming for timely analysis.

Engineers take a holistic approach, using analysis of one data set to inform the analysis of another. In particular, a first-pass analysis of SCADA alarms can prioritize DFR data for inspection by highlighting the likely faulted line. This approach was the basis for protection engineering diagnostic agents (PEDA), which integrated multiple intelligent systems for different aspects of data analysis.

First, a rule-based expert system groups SCADA alarms into events and then groups events into incidents. Assessment of the protection events within an incident indicates whether it needs further investigation. If so, the fault's location and timing is compared against the capture times of fault records to partition them into those directly related to an incident, those related to the incident, those indirectly related, and miscellaneous records with no associated incident. The directly related records are prioritized for further processing.

The type of fault and fault clearance time are extracted from the prioritized records using a second rule-based expert system. At the same time, fault records are passed through a model-based reasoning engine to identify incorrect operation of components, such as a missing inter-trip signal or failure of a trip relay.

These four system components are integrated using multiagent systems technology (see Figure 2). An incident and event identification agent wraps the SCADA analysis functionality and informs subscribed agents
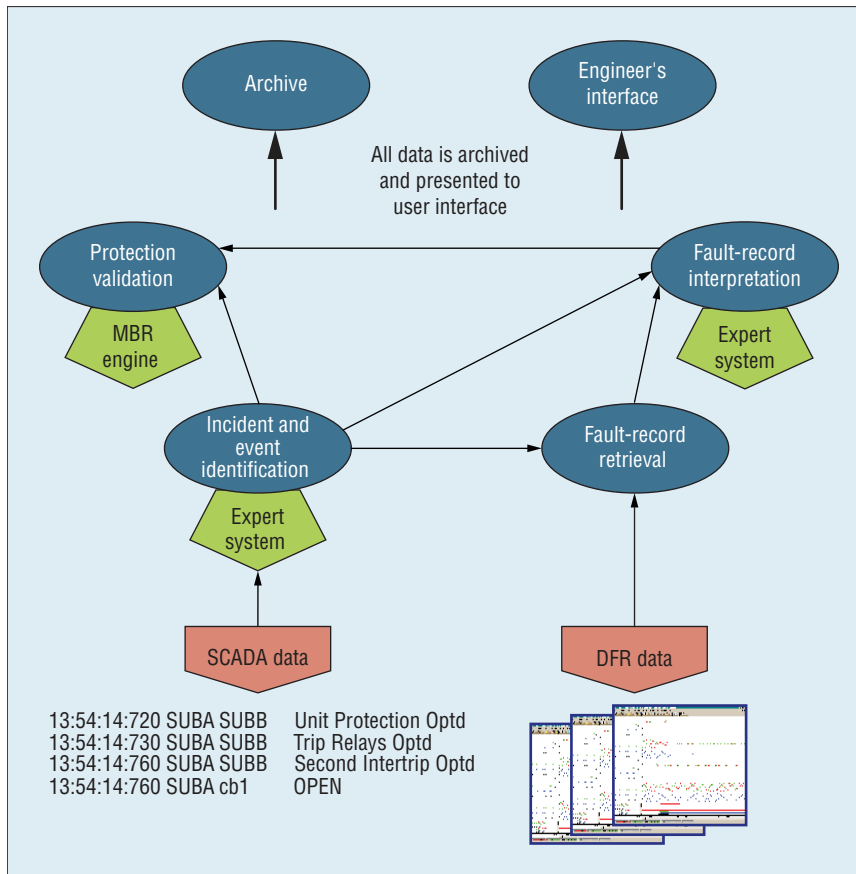
Figure 2. Integrated multiagent systems technology. Protection engineering diagnostic agents (PEDA) incorporate SCADA and DFR data analysis to identify fault type, location, and affected areas of the network and to determine whether the protection system operated correctly.

design life. To prolong their operational lives while minimizing risk of failure in service, online monitoring helps detect behavior changes and diagnose faults.

To this end, we created a *condition-monitoring multiagent system* (COMMAS) that integrates three types of transformer health monitoring.[2] The first is *conventional transformer fault diagnosis*, which involves three techniques for dissolved gas analysis (DGA) that use ratios of gases in transformer oil to identify problems. Second, COMMAS includes a suite of agents for assessing the severity of an insulation deterioration phenomenon called *partial discharge* (PD), where charge travels but only partially bridges the insulation. Earlier work trained three classifiers to diagnose the defect causing PD and used hidden Markov models to detect changes in PD behavior. The third approach is *conditional anomaly detection* (CAD), a technique that meets specific requirements of end-of-life transformer monitoring.

Over a transformer's 40-year life, it develops a particular signature of normal behavior due to low-level defects and shocks that degrade performance while not posing risk of immediate failure or accelerated deterioration. This can be problematic when monitoring health at the end of a transformer's life because automated fault-detection and classification techniques will label its degraded behavior as faulty and needing attention.

Instead of fault diagnosis, engineers are more interested in changes to transformer behavior, either due to fault progression or the development of new problems. At the same time, engineers are sensitive to false positives. The erroneous detection of anomalous behavior on even a few occasions is likely to prejudice

of new incidents when they occur. Fault-record-retrieval agents use incident information to prioritize the collection of DFR data, which is forwarded to a fault-record-interpretation agent for classification and the protection validation and diagnosis agent for model-based analysis. In addition, agents for long-term information storage and an engineer's interface give a complete system for automated post-fault analysis.

From November 2004, many of the PEDA agents were deployed at SP Energy Networks, running online to interpret SCADA and DFR data. After nearly two years of operation, a new SCADA system altered the format and reliability of SCADA data received, requiring nontrivial work to bring it back online. During a

six-month snapshot of its time online, PEDA processed more than 2 million SCADA alarms and 583 fault records, which were correctly reduced to 402 incidents.

Condition Monitoring

Protection systems act to remove immediate threats to equipment integrity, but even when operated within limits, assets experience aging and shocks that degrade their health. Online monitoring can highlight the current condition and aging trend of important assets, letting network operators target maintenance where it will be most beneficial.

Transformers are the most expensive single asset in a transmission network, and within the UK and US, many are coming to the end of their

engineers against anomaly detection techniques as a whole.

To address these challenges, CAD limits false positives by modeling operating conditions and transformer condition.[7] A training set of data is split into environmental and indicator parameters, based on whether they affect the transformer (such as ambient temperature and load current) or indicate its condition (such as oil temperature and oil moisture). Each parameter set is modeled as a mixture of Gaussian distributions, which capture the probability of parameters taking certain values. A third probabilistic model correlates the environment and indicators, mapping Gaussian components of one model to components of the other (see Figure 3).

Anomaly detection consists of interrogating this model with new data and calculating the probability of the transformer condition given the environmental condition. If it exceeds a threshold of unlikeliness, it is deemed a true anomaly and reported to engineers.

The real benefit of this approach is that it will not raise anomalies when the environmental conditions are significantly different from the training examples. If the transformer is operating in a heat wave, its oil temperature could be expected to be anomalously high, but engineering judgment would ascribe the temperature rise to the environmental conditions rather than a transformer problem. Assuming the environmental model does not capture heat-wave conditions, the correlated CAD model reflects its lack of knowledge of how the transformer is expected to behave by returning a higher probability than that of the transformer model alone ($p(ind)$ in Figure 3).

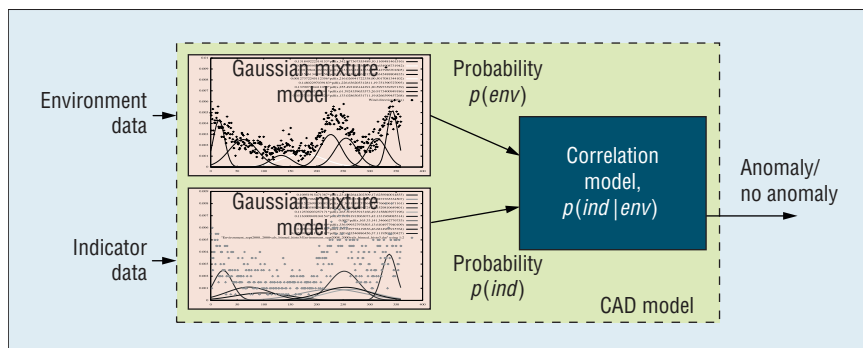CAD was applied to an in-service 275-kV/132-kV, 180-MVA transmission

Figure 3. Conditional anomaly detection. CAD compares the probabilities of transformer indicators and environmental measurements to identify anomalous transformer behavior.

transformer in the UK.[7] This transformer had low levels of fault indicator gasses from historical problems, which made standard diagnostic techniques inaccurate. Thus, anomaly detection was required to model the imperfect normal transformer conditions and alert engineers to any health changes. The CAD model was used to analyze 12 months of online data from the transformer and site, and in that time, it detected 21 anomalies. Further investigation of these anomalies revealed them to be temporary sensor or data-logging faults, where plausible but erroneous values were recorded from the transformer while environmental conditions were normal.

By including CAD within the wider COMMAS agent community, the advantages of anomaly detection are supported by fault diagnosis from the DGA and PD analysis agents. If an anomaly is detected, fault diagnosis can give engineers a familiar categorization of the potential problem to help them plan corrective maintenance. Equally, corroboration between fault diagnoses can help engineers draw conclusions about the transformer's health. That is, if the DGA and PD agents all conclude that PD is occurring, there is a high confidence that this is indeed correct.

Furthermore, the robustness of the agent approach means that if one data source becomes inaccessible, or if an agent otherwise fails to perform its diagnosis, the others in the system can autonomously continue their tasks and present a partial view of the transformer's health. In this way, the combination of appropriate intelligent systems techniques within a multiagent system provides a condition-monitoring approach that meets the requirements of flexibility, extensibility, and graceful degradation.

Agent-Based Smart-Grid Architecture

The potential for intelligent systems in many areas of network management is clear, but these systems must not be left to operate in isolation. The interpretation of one data set can often benefit another, as in the case of post-fault analysis with SCADA and DFR data, but this requires communication between each system. As more systems for data interpretation become capable of interoperation, the scope for higher levels of corroboration grows.

One example is the possibility of making better asset condition assessments with knowledge of network events. Traditionally, incidents and events on the network are not easily available to condition-monitoring engineers, and they assess the health of assets such as transformers purely from condition data. However, a fault can draw a large current that constitutes a shock to the transformer, temporarily increasing temperatures and taxing the insulation system. Without knowledge of network events,

condition-monitoring engineers are speculating about whether these worrying trends are due to external forces or to true degradation of the transformer's condition.

Information about asset health can also benefit control decisions. If it is known that a transformer is experiencing problems, power-flow management can favor the load reduction on the transformer to, say, 50 percent of its rating until maintenance is possible. If a circuit breaker's condition suggests that it might fail, an alternate network configuration could help keep that connection in its current state. This sort of planning is hard to achieve with the manual procedures currently in place, but the integration of multiple types of data analysis within a smart grid paves the way to its realization.

Such an integration of disparate systems requires a standard approach to building and deploying individual components of functionality. Researchers have recognized interoperability as a key facet of smart grid functionality, and within the industry, there is much activity on standardizing data formats, such as from the US National Institute of Standards and Technology[8] and the IEEE P2030 Working Group.[9] However, data formatting is only one component of flexible, open systems that must also adhere to expected levels of robustness, extensibility, and graceful degradation.

Multiagent systems technology is a way of developing loosely coupled autonomous systems, where each segment of the system can independently pursue its goals while using standards-conforming messaging to interact with others.[10] This has several advantages. If one agent fails in its task, it does not necessarily affect or degrade the others in the system. Considering the integration of smart grid capabilities, if a

condition-monitoring agent were to stop responding, it need not adversely affect agents performing control roles. The loss of some functionality might lead to less optimal network management, but it will not prevent other agents from taking reasonable decisions based on less complete information. Particular agent platform deployments can produce a more or less robust system,[2] but in general an agent approach allows for graceful degradation.

A multiagent architecture for the smart grid can build on the standards already developed for interoperability within the power industry. For example, data exchange between utilities is facilitated by the Common Information Model (CIM), and there is a growing use of the IEC 61850 standard for intrasubstation messaging. These standards cover the main domains of communication required by agents performing smart grid roles and hence can form the basis of a smart-grid agent ontology. The IEEE Power and Energy Society Multiagent Systems Working Group aims to promote the openness of agent architectures within the power domain and provide an upper ontology based on CIM on their website (http://ewh. ieee.org/mu/pes-mas).

In addition, we have identified five key agent roles within a smart-grid monitoring system, which offer templates for agent behavior and interactions.[2] By loosely standardizing the roles and interactions, agents can be designed such that behaviors providing social ability for each role can be reused with some parameterization. These roles are not exclusive, with specific provisions for interfacing to clusters of agents (such as PEDA) that do not implement the roles. For those that do, the reuse of social behavior code saves on development time and allows new data sources and services

to be automatically locatable by existing system agents.

The changing nature of the electricity network makes it a more complex, dynamic system to manage. Intelligent systems can provide automation of many of the tasks currently performed by engineers, but in ways that are generically applicable, extensible to changing regimes, and explainable.

The deployment of automated intelligence allows for information exchange across the traditional boundaries of engineering roles, improving decision-making. Realizing this depends on communication between each area of analysis, while robustness relies on the autonomy of each role. We believe the intelligent systems techniques and applications we discuss represent the building blocks of smart-grid architectures, while multiagent systems offer a platform technology for deployment of this functionality. ▬

References

1. A. Vojdani, "Smart Integration," *IEEE Power and Energy*, vol. 6, no. 6, 2008, pp. 71–79.
2. V.M. Catterson, E.M. Davidson, and S.D.J. McArthur, "Agents for Active Network Management and Condition Monitoring in the Smart Grid," *Proc. Agent Technologies for Energy Systems Workshop, 9th Int. Conf. Autonomous Agents and Multiagent Systems*, 2010; http://users.ecs.soton.ac.uk/acr/ates2010/Paper_03.pdf.
3. P. Plener, M. Carlsson, and C. Schulte, "Constraint Programming in Sweden," *IEEE Intelligent Systems*, vol. 24, no. 2, 2009, pp. 87–89.
4. D. Lindenmeyer, H.W. Dommel, and M.M. Adibi, "Power System Restoration: A Bibliographical Survey," *Int'l J. Electrical Power and Energy Systems*, vol. 21, no. 3, 2001, pp. 219–227.

5. S. Thièbaux and M. Cordier, "Supply Restoration in Power Distribution Systems: A Benchmark for Planning Under Uncertainty," *Proc. 6th European Conf. Planning* (ECP-01), IOS Press, 2001, pp. 85–95.

6. E.M. Davidson et al., "Applying Multiagent System Technology in Practice: Automated Management and Analysis of SCADA and Digital Fault Recorder Data," *IEEE Trans. Power Systems*, vol. 21, no. 2, 2006, pp. 559–567.

7. V.M. Catterson, S.D.J. McArthur, and G. Moss, "On-Line Conditional Anomaly Detection in Multivariate Data for Transformer Monitoring," *IEEE Trans. Power Delivery*, vol. 25, no. 4, 2010, pp. 2556–2564.

8. US Nat'l Inst. Standards and Technology, "NIST Framework and Roadmap for Smart Grid Interoperability Standards Release 1.0," 2010; http://www.nist.gov/public_affairs/releases/upload/smartgrid_interoperability_final.pdf.

9. IEEE P2030, "Draft Guide for Smart Grid Interoperability of Energy Technology and Information Technology Operation with the Electric Power System (EPS);" 2010; http://grouper.ieee.org/groups/scc21/2030/2030_index.html.

10. M. Pechoucek et al., "Agents in Industry: The Best from the AAMAS 2005 Industry Track," *IEEE Intelligent Systems*, vol. 21, no. 2, 2006, pp. 1541–1672.

THE AUTHORS

Victoria M. Catterson is a senior research fellow at the Institute for Energy and Environment, University of Strathclyde, Glasgow. Her research interests include intelligent systems for condition monitoring within the power industry. Catterson has a PhD in transformer diagnostics from the University of Strathclyde. She is a member of IEEE and the power-industry association CIGRÉ. Contact her at v.m.catterson@strath.ac.uk.

Euan M. Davidson is a senior research fellow at the Institute for Energy and Environment, University of Strathclyde, Glasgow. His research interests include multiagent systems, model-based reasoning, and applying intelligent system techniques to diagnostic, power-system control, operation, and management problems. Davidson has an MEng in electrical and electronic engineering from the University of Strathclyde. He is a member of IEEE and the Institution of Engineering and Technology (IET). Contact him at edavidson@eee.strath.ac.uk.

Stephen D.J. McArthur is a professor and codirector of the Institute for Energy and Environment, University of Strathclyde, Glasgow. His research interests include intelligent systems applications in power, condition monitoring, fault diagnosis, and multiagent systems. McArthur has a PhD in power system fault diagnostics from the University of Strathclyde. A member of IEEE and IET, he is also the past chair of the IEEE Power and Energy Society Intelligent Systems Subcommittee, and Multiagent Systems Working Group. Contact him at s.mcarthur@eee.strath.ac.uk.

cn *Selected CS articles and columns are also available for free at http://ComputingNow.computer.org.*