



Catterson, V.M. and Rudd, S.E. and McArthur, S.D.J. and Moss, G. (2009) On-line transformer condition monitoring through diagnostics and anomaly detection. In: IEEE International Conference on Intelligent Systems Application to Power Systems 2009 (ISAP 2009), 2009-11-08 - 2009-11-12, Curitiba, Brazil. ,

This version is available at <https://strathprints.strath.ac.uk/26475/>

Strathprints is designed to allow users to access the research output of the University of Strathclyde. Unless otherwise explicitly stated on the manuscript, Copyright © and Moral Rights for the papers on this site are retained by the individual authors and/or other copyright owners. Please check the manuscript for details of any other licences that may have been applied. You may not engage in further distribution of the material for any profitmaking activities or any commercial gain. You may freely distribute both the url (<https://strathprints.strath.ac.uk/>) and the content of this paper for research or private study, educational, or not-for-profit purposes without prior permission or charge.

Any correspondence concerning this service should be sent to the Strathprints administrator: strathprints@strath.ac.uk

Catterson, V. M. and Rudd, S. E. and McArthur, S. D. J. and Moss, G. (2009) On-line transformer condition monitoring through diagnostics and anomaly detection. In: IEEE International Conference on Intelligent Systems Application to Power Systems 2009 (ISAP 2009), 8-12 Nov 2009, Curitiba, Brazil.

<http://strathprints.strath.ac.uk/26475/>

Strathprints is designed to allow users to access the research output of the University of Strathclyde. Copyright © and Moral Rights for the papers on this site are retained by the individual authors and/or other copyright owners. You may not engage in further distribution of the material for any profitmaking activities or any commercial gain. You may freely distribute both the url (<http://strathprints.strath.ac.uk>) and the content of this paper for research or study, educational, or not-for-profit purposes without prior permission or charge. You may freely distribute the url (<http://strathprints.strath.ac.uk>) of the Strathprints website.

Any correspondence concerning this service should be sent to The Strathprints Administrator: eprints@cis.strath.ac.uk

On-line Transformer Condition Monitoring through Diagnostics and Anomaly Detection

V. M. Catterson*, S. E. Rudd*, S. D. J. McArthur*, and G. Moss[†]

*Institute for Energy and Environment, University of Strathclyde, Glasgow, United Kingdom

Email: v.m.catterson@strath.ac.uk, srudd@eee.strath.ac.uk, s.mcarthur@eee.strath.ac.uk

[†]National Grid, Warwick, United Kingdom

Email: graham.moss@uk.ngrid.com

Abstract—This paper describes the end-to-end components of an on-line system for diagnostics and anomaly detection. The system provides condition monitoring capabilities for two in-service transmission transformers in the UK. These transformers are nearing the end of their design life, and it is hoped that intensive monitoring will enable them to stay in service for longer. The paper discusses the requirements on a system for interpreting data from the sensors installed on site, as well as describing the operation of specific diagnostic and anomaly detection techniques employed. The system is deployed on a substation computer, collecting and interpreting site data on-line.

I. INTRODUCTION

As transformers reach the end of their design lives, utilities wishing to minimize replacement costs must handle the conflicting requirements of retaining plant items in service for as long as possible, whilst avoiding failures in service. The application of various on-line monitoring and diagnostic techniques can aid in the identification of faults within the transformer [1], highlighting the need for maintenance or removal from service. Increasing this level of monitoring, by augmenting various monitoring techniques, can provide a fuller picture of plant health, increasing the certainty when a fault is identified and thus helping to anticipate imminent failure.

However, transformers nearing the end of their lives will have known, minor problems that are being managed and that do not pose an immediate risk. These known problems can mask the inception of new, potentially more critical issues when using standard diagnostic techniques such as dissolved gas analysis (DGA) [2]. Therefore, this paper proposes the use of machine learning techniques that allow the ‘normal’ behavior, be it without any faults or with known ones, to be learnt for a given transformer.

Each transformer will age at a different rate due to many factors, including different load profiles experienced and different weather conditions. Since each transformer behaves individually, machine learning techniques can be applied to learn what normal behavior is for a particular item of plant [3]. If the item’s behavior diverges from normal, it can be identified that an anomalous event has occurred within this plant item.

This approach has a second advantage, relating to fault identification, when data has no standard model or approach for interpretation. This type of data can arise, for instance, when new sensors are retrofitted to the transformer. In these

cases, by first learning normal data signals, anomalies in indecipherable data can be detected and may then potentially be identified with recognized causes.

Anomaly detection tends to identify an event that is statistically unlikely; however, it cannot explain why it is unlikely. Identification of the type of fault that is occurring within the transformer can be achieved by automated diagnostic techniques. Therefore, the creation of a flexible, on-line system that incorporates a suite of anomaly detection and diagnostic techniques is required to give detailed health analyses about plant behavior.

This paper describes the on-going intensive monitoring of two 275/132kV 180MVA transformers in the UK. Multi-agent systems technology is used to integrate various diagnostic and anomaly detection modules, with data from approximately 50 sensors captured every five minutes. This paper will show how the collaboration of various sensory data can lead to an enhanced estimation of the transformers health. Section II describes the sensor installation and data infrastructure. Section III introduces the multi-agent system architecture used to mesh together the various detection and diagnostic techniques. Details of the Conditional Anomaly Detection technique, and how it is applied to these transformers is given in Section IV. The integration of previous work on partial discharge (PD) analysis is described in Section V and the paper concludes in Section VI.

II. SITE INSTALLATION

Two 275/132kV 180MVA transformers at a substation in the UK are nearing the end of their design life. The owner utility wants to keep this plant in service for as long as possible without too great a risk of failure in service. To achieve this, the transformers are under intensive monitoring with a wide set of on-line sensors, allowing engineers constant access to the latest measurements.

The aims of the project were threefold. First, the utility wanted to be able to detect anomalous transformer behavior on-line, potentially having early warning of failure. Secondly, the transformers are sister units installed in successive years, yet one is showing more signs of aging than the other. Any correlations or deviations in behavior between the two units should provide information about their relative states of health. Thirdly, this project employs a mixture of conventional and

unusual sensors, both internal and external to the transformers, with the aim of assessing the advantages and disadvantages of cheaper, externally mounted sensors.

To this end, a large number of sensors monitor the more aged unit, with a subset of the same sensors on the other. These include temperature and vibration sensors externally mounted at various positions on the transformer tank and tap changer, external vibration sensors on the cooling bank pumps and fans, and load and auxiliary system currents. The internal sensors include top and bottom oil moisture and temperature. Additionally, a Kelman TRANSFIX unit monitors levels of eight dissolved gases, while a GE Hydran unit measures hydrogen.

The substation environmental conditions are also being monitored. Ambient temperature, wind speed and direction, air pressure, and rain intensity and duration are recorded by an on site weather station.

The TRANSFIX unit was installed before the other sensors, and the data is archived with all the utility's dissolved gas data in a data warehouse. Since this archive is commercially sensitive, on-line access for the researchers was not possible. Files of historical data give an off-line record of the hourly readings. However, the remaining sensors were installed specifically for this project, and the data is archived to a second data warehouse. On-line access to this data is possible using a web site for engineers, or a web service for machine-readable access. These sensors take measurements every five minutes.

The more conventional parameters of dissolved gas levels and top oil temperature have the disadvantage of needing internal sensors. This has safety and scheduling implications, as an outage may be required to get internal access. On the other hand, external sensors may be fitted without an outage, but give a more approximate estimation of top oil and hot spot temperature.

Further, in the case of vibration and weather parameters, there is little knowledge about what such parameters may tell about transformer health. This indicates the strong need for a flexible condition monitoring system, capable of handling various types of parameter, and extensible enough to incorporate new types of data interpretation as more is understood about the data. The design of such a system is described in the following section.

III. MULTI-AGENT SYSTEM FOR CONDITION MONITORING

Based on the mixture of sensors deployed on site and the aims of the monitoring project, the condition monitoring system had to meet the following requirements:

- 1) It must be capable of taking input data from various sources, namely the on-line data warehouse and files of TRANSFIX data;
- 2) It must monitor the transformers for anomalous behavior, and indicate when an anomaly occurs;
- 3) It should attempt to diagnose possible faults, in order to provide more information about anomalies;

- 4) It should include multiple techniques for analyzing data, and be extensible as more techniques are developed;
- 5) Where possible, the same techniques should be applied to both transformers, to allow comparison of performance.

Fundamentally, these requirements mean the system must be capable of flexibly adding and duplicating modules to gather and interpret data, and must be extensible for future uses. To meet these needs, multi-agent system technology was chosen as the platform for system design.

Within the power engineering community, the definition of an agent proposed by Wooldridge has gained most support[4]. This defines an intelligent agent as software that displays flexible autonomy, through a mixture of pro-activeness, reactivity, and social ability. A multi-agent system is a collection of agents operating autonomously but co-operatively, which results in some system goal being achieved.

For a condition monitoring application, that system goal is to derive as much information about plant health as possible. Agent technology allows multiple data collection and interpretation techniques to be created as separate agents, which use their social ability to find others within the system with which to usefully co-operate and share data.

Specifically for the monitoring project reported here, agents were created for a variety of tasks, including interfacing to the data warehouse, applying data analysis techniques to recorded measurements, and alerting engineers to transformer condition. The complete system is shown in Figure 1, with the following sections describing the design and functionality of each group of agents.

A. Architecture Overview

Each agent within the architecture conforms to one of five different types. Data provider agents take data from an external source, such as the data warehouse or historical files, and convert it into a format understandable by other agents in the system. This means that other agents in the system need not be concerned with where data is stored, they can simply query the data provider for the appropriate data type or time period.

Data interpretation and analysis is performed by service provider agents. The data director agents link appropriate data providers with service providers, allowing great flexibility in the order of data processing tasks. This flexibility may be needed when deploying the same monitoring software in different environments, such as adding a filtering stage to data analysis in a noisy EM environment. The filter can be deployed as a separate agent, with the data director deciding when it is appropriate to filter noise, and when data can be sent directly for interpretation.

A fourth type of agent is an archive, for long term storage of the output of service provider agents. Finally, an external interface is required to convert from agent-format data into external formats, such as an engineer's user interface or other systems. This is essentially the opposite of the data provider's task.

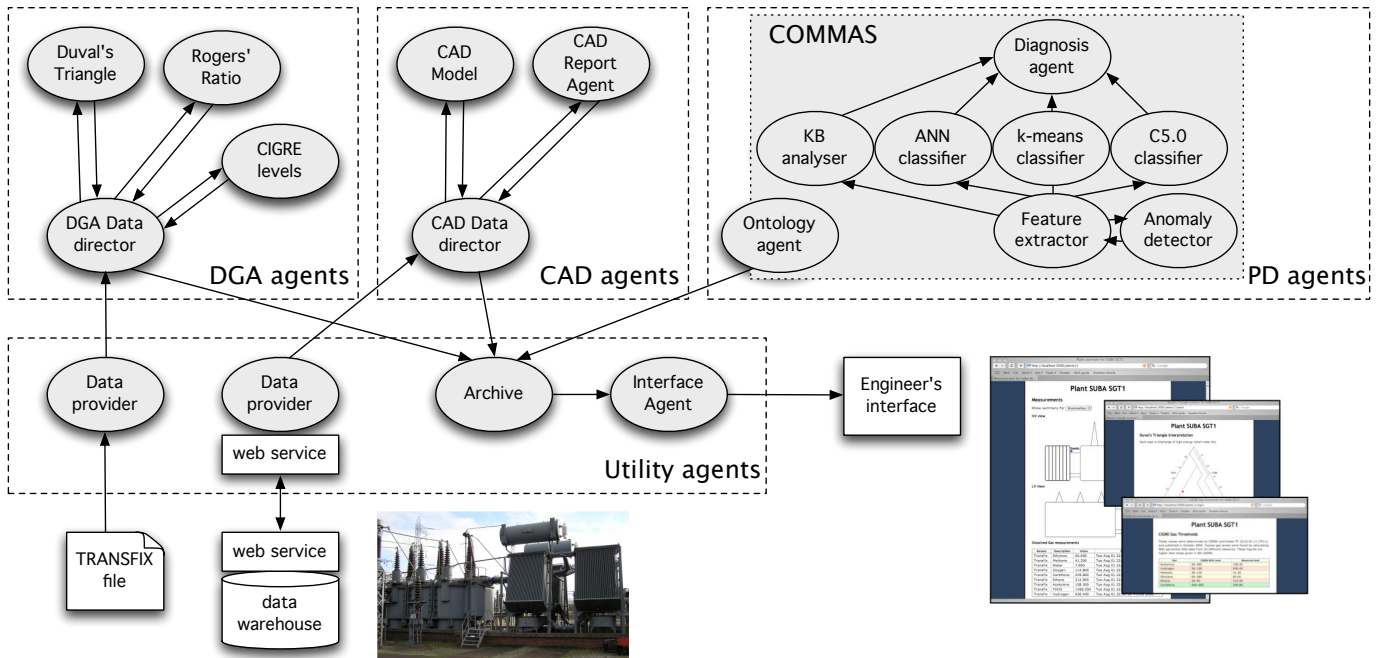


Fig. 1. The complete on-line condition monitoring architecture, with groups of agents split by high-level function.

In order to keep this architecture as open and future-proof as possible, agent messaging conforms to the standards created by the Foundation for Intelligent Physical Agents (FIPA)[5]. Specifically, agents send messages using FIPA Agent Communication Language (FIPA ACL), where message content conforms to FIPA Semantic Language (FIPA SL). However, standards do not cover the ontology—the dictionary of terms used within messages—since this is application dependent. Previous work suggested the Common Information Model (CIM) could form the basis of a standard ontology for the power engineering domain[4], as it is an object-oriented model for data exchange between Energy Management Systems; therefore this architecture uses an ontology based on CIM.

B. Utility Agents

Figure 1 shows four utility agents, which perform tasks related to the operation of the system rather than data interpretation. Two data provider agents are needed: one to parse historical files of dissolved gas data from the TRANSFIX system, and one to interface with the on-line data warehouse via a web service. In both cases, measurements from the transformers under study are converted into the system ontology, and the agents answer queries about the measurements they hold.

One interface agent is required to populate a web-based system interface for engineers. This web application displays the latest information about transformer health, based on the most recent site measurements and any diagnostics or anomaly detection agents running. Additionally, one archive agent stores the information generated by other system agents.

C. Dissolved Gas Analysis

Dissolved Gas Analysis (DGA) is a widely-used technique for diagnosing transformer faults[2]. There are a handful of popular interpretation methods, based on ratios of gases, absolute levels, or the gas with the highest concentration. For this project, three were implemented as separate agents: Duval's Triangle, Rogers' Ratio, and CIGRE thresholds[6]. One data director agent is required to collect gas measurements and forward to these techniques. The resulting collection of agents is shown in Figure 1.

The output of this group of agents is information about the levels of dissolved gases in the transformer. Specifically, the CIGRE threshold agent identifies which gases have an absolute level within standard limits, and which are higher. Duval's Triangle and Rogers' Ratio give diagnoses of the likely cause of these deviations, such as PD or thermal fault.

D. Conditional Anomaly Detection

To complement Dissolved Gas Analysis, which diagnoses specific fault types, a technique for multivariate anomaly detection was implemented. Anomaly detection looks for deviations from normal or expected plant behavior, where 'normal' can include known low-level faults. It can therefore detect changes due to incipient faults at an earlier stage than diagnostic techniques such as DGA, which can be confused by multiple faults.

The specific technique chosen for this application is Conditional Anomaly Detection, where anomalous plant behavior is assessed in the context of the environmental conditions experienced by the plant, such as ambient temperature and load current. Detail of how this is achieved follows in Section IV.

The group of agents deployed for anomaly detection is shown in Figure 1, and described further in Section IV.

E. Partial Discharge

In addition to the agents and techniques developed specifically for monitoring the application transformers, the authors have previous experience with a multi-agent system for PD analysis[7]. Developed as the stand-alone system COMMAS, this set of diagnostic and anomaly detection agents were integrated with the full monitoring system to show its extensibility to future developments and updates. Shown in Figure 1, the operation of the COMMAS agents, as well as the means of agent integration, is described in detail in Section V. In short, this group of agents offers classification and explanation of partial discharge-causing defects, with detection of anomalous PD behavior.

F. Installation

On-line operation of this system was required to deliver timely health information about the transformers under study. To this end, the multi-agent system was installed on an ABB substation computer and deployed in a lab at the University of Strathclyde.

Every five minutes, the on-line data provider agent connects to the data warehouse’s web service and collects new measurements. As new files of dissolved gas data become available, they are fed to the historical data provider agent. The data director agents pick up new data as appropriate, and pass it to service provider agents, collecting service output for archival.

Engineers can view the system through the engineer’s web interface, to see up-to-date results of the diagnostic and anomaly detection agents.

IV. CONDITIONAL ANOMALY DETECTION

Conditional Anomaly Detection (CAD)[8] is a type of *contextual anomaly detection*[9]; that is, the context of measurements is considered before identifying something as anomalous. This is of particular importance for transformer monitoring, as network events such as reconfiguration may cause transient unusual operating conditions that could be reflected in monitoring data. A technique that considers the environmental and operating context of the plant can reduce the number of false anomalies by recognizing that unusual conditions may cause unusual behavior without being indicative of a problem or fault.

To detect anomalous transformer measurements, the CAD technique learns two models: the statistical behavior of the monitored plant, and the statistical behavior of the environment. The probabilistic links between the models are also learned, giving a combined model of likely data that co-occurs in the environment and the plant. True anomalies can then be defined as statistically unlikely events in the plant parameters that occur when the environment is normal. Within the CAD literature, the parameters from the plant under study are termed the *indicator* parameters, while those from the surrounding conditions are the *environment* parameters.




Fig. 2. The CAD model comprises Gaussian mixture models U and V for indicator and environment parameters, and the correlation model $p(U|V)$.

The technique used in [8] for learning the indicator and environment models is Gaussian mixture modeling. This is a method of representing a dataset as a mixture of Gaussian (normal) components, giving a probabilistic model of correlations between different measurement parameters. The parameters of the model are $\Theta = \langle \theta_0, \theta_1, \dots, \theta_k \rangle$ for k components. Each θ comprises the Gaussian parameters $\langle \mu, \Sigma, \pi \rangle$ for mean, covariance matrix, and mix proportion.

A. Learning the CAD Model

The process of learning a CAD model has three stages. First, a Gaussian mixture model, U , of the indicator parameters is trained. Next, a Gaussian mixture model, V , of the environment is trained. Lastly, the correlation model, $p(U|V)$, is trained. These three parts together form the CAD model, shown in Figure 2.

All three stages of training can be performed using the Expectation Maximization algorithm: an iterative technique that converges towards a locally optimal set of values for unknown model parameters, which maximize the likelihood of the overall model[10]. In the case of Gaussian mixture models, the unknown model parameters are the $\langle \mu_i, \Sigma_i, \pi_i \rangle$ for all components i . For the correlation model, the unknown parameters are the probabilities of a component in the indicator model co-occurring with a component in the environment model, $p(U_i|V_j)$ for all components i, j .

With this model, new site measurements can be assessed for anomalousness. Given particular values for the indicator and environmental parameters, the CAD model returns a value, f_{CAD} , which indicates the likelihood of these indicator conditions weighted by the likelihood of the environment. If the value of f_{CAD} is below a certain threshold, the measurements represent an anomaly; if it is above the threshold, either the indicator measurements are within normal ranges, or the environmental conditions are also anomalous, likely due to transient unusual operating conditions.

The threshold for anomalousness must be chosen through experimentation with an individual CAD model, as it is dependent on the particular application. The anomaly detector can be made more or less sensitive by varying the threshold, allowing control over the model output even after training.

B. Transformer Monitoring

CAD was applied to the transformers under study. An engineer’s visual inspection of data from September and October

TABLE I
ANOMALIES FOUND IN THE SEVEN MONTH TEST SET

Month	Anomalies ($f_{CAD} < 1 \times 10^{-20}$)
November 2008	0
December 2008	14
January 2009	0
February 2009	1
March 2009	5
April 2009	0
May 2009	0

2008 revealed normal transformer behavior during that time, so these two months of data were used for training the CAD model.

The environmental model consists of five parameters: ambient temperature, solar radiation, wind speed, wind direction, and Y-phase load current. The indicator model focusses on transformer oil parameters: top oil temperature, bottom oil moisture, and levels of hydrogen from the GE Hydran unit. In both cases, a set of Gaussian mixture models ranging in size from five components to 17 were trained, with the best fit model being used for CAD. The correlation model was trained using the selected Gaussian mixture models: a 14-component indicator model, and a 16-component environment model.

C. Results

The complete CAD model, comprising the indicator, environment, and correlation models, was wrapped as a service provider agent for deployment within the monitoring architecture. Data from the seven months following the training period was used for model testing (November 2008 to May 2009). Based on this testing, the threshold for anomalous values was set to $f_{CAD} = 1 \times 10^{-20}$.

The results of this testing are shown in Table I. This period contains great differences in the ambient weather conditions, from cold winter in December and January through to (unseasonably) warm summer conditions in May. However, there is no pattern to the occurrence of anomalies that suggests performance is affected by the season.

Indeed, the 20 anomalies detected during this time are all due to temporary sensor and logging failures for one or two measurements in a row, not due to transformer behavior. This suggests that individual occurrences of anomalous measurements are not of great interest to engineers. Rather, it is continuous anomalous measurements or an increasing frequency of anomalies that are important.

As a result, a second stage of processing was added as the CAD Report Agent shown in Figure 1. This agent aggregates statistics about the frequency of anomalies, looking hourly over the last seven days, and daily over the last 30 days. This information is presented as histograms in the engineer's interface, so that trends are immediately apparent. If particular anomalies warrant further investigation, the engineer can view the CAD model output and operation for the associated measurements. This gives detail of which parameters are causing the anomaly.

By deploying this set of agents, the system gains multivariate anomaly detection to supplement the diagnostic techniques. A great benefit of this particular technique is that it reduces the potential for 'false alarm' anomalies caused during unusual operating conditions. With the addition of the CAD report stage, engineers can see at-a-glance the history of anomalous behavior.

V. INTEGRATING PARTIAL DISCHARGE ANALYSIS

For a condition monitoring system to be useful in the long-term, it must be both flexible, allowing easy modification to integrate new sensory data, and extensible, providing a way of augmenting new diagnostic techniques. This transformer monitoring project exemplifies why these requirements are so important. Here, a number of sensors with no known diagnostic techniques, such as accelerometers detecting vibration levels, have been retrofitted. In this case, the system must be flexible to integrate the new data sources and, as the data becomes more familiar over time, must be extensible to allow new diagnosis tools to be added. This flexibility and room for expansion, along with the ability to collate interpretations from various diagnostic techniques, is key to this project. In order to test the extensibility of the full condition monitoring architecture, the system was integrated with a multi-agent system for PD anomaly detection and diagnosis, COndition Monitoring Multi-agent System (COMMAS).

A. Prior Work on PD Analysis

Within the University of Strathclyde, various PD analysis techniques have been the focus of research. These include data-driven techniques [11], knowledge-based interpretation [12], and anomaly detection [3]. It was discovered that whilst, individually, the data-driven diagnostic techniques could each diagnose faults in the transformer to some degree, certain diagnostic approaches would lead to a more accurate diagnosis of particular types of fault. This led to the creation of the COndition Monitoring Multi-agent System (COMMAS) [7], which combined and interpreted outputs from various data mining and intelligent techniques.

To add explanation and therefore confidence to the overall classification result that COMMAS produced, a knowledge-based approach was introduced [12]. Although originally created as a stand-alone system, it was felt that the integration of this system into the COMMAS architecture would provide the engineer with varying degrees of explanation concerning the diagnosis. The flexible and extensible architecture offered by the multi-agent approach meant it could be easily extended and the knowledge-based approach integrated by wrapping the technique within an agent (shown in Figure 1 as the KB analyzer agent).

The knowledge-based method utilizes expert knowledge regarding phase-resolved patterns [13] and PD phenomena to identify defects causing PD. A five step method is utilized mimicking the approach an expert would take to diagnose a defect when examining a phase-resolved pattern. This process provides the engineer with an explanation regarding the

features of the pattern used to distinguish the classification, along with the identified PD behaviors, defect characteristics and potential location of the defect.

To further improve the original COMMAS, an anomaly detection agent was created as the first stage of PD analysis [3]. This agent learns normal PD behavior for individual transformers, after which the detection of anomalous events can warrant classification by the diagnostic agents and explanation through the knowledge-based approach.

Together this set of agents implement a suite of PD analysis techniques, informing the engineer of specific plant anomalies, providing a classification and diagnosis of faults, and supplying a detailed explanation of how the diagnosis was achieved (shown as the COMMAS block within Figure 1). This suite of agents could be usefully integrated into the overall architecture, supplementing dissolved gas analysis and multivariate anomaly detection with complete PD interpretation.

B. Integration

Since the detection and diagnostic techniques in COMMAS were developed as agents, the whole suite could be easily deployed with the agents of the full system. The only difficulty arose with agent message ontology.

Since the COMMAS agents predate the new architecture, they also predate the CIM-based ontology (discussed in Section III-A). Agent ontology is often fundamental to the way agents represent data internally, in addition to their external communication, and so changing the agent ontology requires a significant amount of development time. Rather than re-implementing the ‘legacy’ PD agents, or forgoing the benefits of a standards-based ontology, the solution was to deploy an ontology translation agent capable of understanding both the older and newer ontologies.

This agent is shown in Figure 1 straddling the COMMAS boundary, as it translates messages from all of the COMMAS agents into the CIM-based ontology for archival, while translating messages from the wider community of agents into appropriate COMMAS-ontology format. In this way, it appears as a service and data provider in wider system, allowing the COMMAS system to fully integrate without redevelopment.

The addition of COMMAS to the overall system for transformer monitoring demonstrates that the architecture is indeed flexible and extensible, due to the use of multi-agent system technology. The addition of PD analysis for fault detection and diagnosis provides further monitoring capabilities for the complete system.

VI. CONCLUSION

This paper details the design and operation of an on-line agent-based system for the monitoring of two in-service transmission transformers. This application had particular requirements of handling data from disparate sources, including multiple techniques for data interpretation, and replicating data analysis across both units. Agent technology was chosen to support the delivery of a flexible and modular system to meet these needs, leading to the creation of a condition monitoring

architecture of five different types of agent with a CIM-based ontology.

In addition to flexibility requirements, the application called for anomaly detection and diagnostic capabilities. To support this, groups of agents handling dissolved gas diagnosis, Conditional Anomaly Detection, and partial discharge anomaly detection and diagnosis were integrated into the architecture. The system is installed and deployed on a substation PC running in the laboratory, collecting and interpreting measurements from the site every five minutes.

ACKNOWLEDGEMENT

This work is funded through the EPSRC Supergen V, UK Energy Infrastructure (AMPerES) grant in collaboration with UK electricity network operators working under Ofgem’s Innovation Funding Incentive scheme—full details on <http://www.supergen-amperes.org>.

REFERENCES

- [1] R. James and Q. Su, *Condition Assessment of High Voltage Insulation in Power System Equipment*. The Institution of Engineering and Technology, 2008.
- [2] T. K. Saha, “Review of Modern Diagnostic Techniques for Assessing Insulation Condition in Aged Transformers,” *IEEE Trans. Dielectr. Electr. Insul.*, vol. 10, no. 5, pp. 903–917, Oct. 2003.
- [3] A. Brown, V. Catterson, M. Fox, D. Long, and S. McArthur, “Learning models of plant behavior for anomaly detection and condition monitoring,” in *Intelligent Systems Applications to Power Systems, 2007. ISAP 2007. International Conference on*, Nov. 2007, pp. 1–6.
- [4] S. D. J. McArthur, E. M. Davidson, V. M. Catterson, A. L. Dimeas, N. D. Hatziaargyriou, F. Ponci, and T. Funabashi, “Multi-Agent Systems for Power Engineering Applications—Part 1: Concepts, Approaches and Technical Challenges,” *IEEE Trans. Power Systems*, vol. 22, no. 4, pp. 1743–1752, Nov. 2007.
- [5] Foundation for Intelligent Physical Agents, “FIPA Standards Repository,” <http://www.fipa.org/>, 2007.
- [6] CIGRE, “Recent Developments in DGA Interpretation—CIGRE D1.01/A2.11 296,” Tech. Rep., 2006.
- [7] S. D. J. McArthur, S. M. Strachan, and G. Jahn, “The Design of a Multi-Agent Transformer Condition Monitoring System,” *IEEE Trans. Power Syst.*, vol. 19, no. 4, pp. 1845–1852, Nov. 2004.
- [8] X. Song, M. Wu, C. Jermaine, and S. Ranka, “Conditional Anomaly Detection,” *IEEE Trans. Knowl. Data Eng.*, vol. 19, no. 5, pp. 631–645, May 2007.
- [9] V. Chandola, A. Banerjee, and V. Kumar, “Anomaly Detection: A Survey,” to appear in *ACM Computing Surveys*, Sep. 2009.
- [10] Z. Ghahramani and M. I. Jordan, “Learning from incomplete data,” MIT Center for Biological and Computational Learning, Technical Report 108, available from <http://hdl.handle.net/1721.1/7202>.
- [11] S. Strachan, G. Jahn, S. McArthur, and J. McDonald, “Intelligent Diagnosis of Defects Responsible for Partial Discharge activity Detected in Power Transformers,” in *Intelligent Systems Applications to Power Systems, 2003. ISAP 2003. International Conference on*, 2003, p. ISAP03/109.
- [12] S. Strachan, S. Rudd, S. McArthur, M. Judd, S. Meijer, and E. Gulska, “Knowledge-based diagnosis of partial discharges in power transformers,” *Dielectrics and Electrical Insulation, IEEE Transactions on*, vol. 15, no. 1, pp. 259–268, February 2008.
- [13] J. Pearson, O. Farish, B. Hampton, M. Judd, D. Templeton, B. Pryor, and I. Welch, “Partial discharge diagnostics for gas insulated substations,” *Dielectrics and Electrical Insulation, IEEE Transactions on*, vol. 2, no. 5, pp. 893–905, Oct 1995.