ANOMALY DETECTION IN ELECTRICITY CYBER INFRASTRUCTURES

Xuan Jin, John Bigham, Julian Rodaway, David Gamez, Chris Phillips

Electronic Engineering Department, Queen Mary, University of London[1]

Keywords: Electricity cyber infrastructure, anomaly detection, invariant induction, artificial ants

Abstract: This paper presents a novel anomaly detection methodology for the protection of electricity critical infrastructures that learns the normal behaviour of the system, builds up a profile and detects anomalous operations which deviate from the profile. This can be used to identify attacks, failures and accidents and it can also be used to improve state estimation, correct topology errors and inform the operators about potential discrepancies between their view of the network and its actual state. This paper will cover two of the anomaly-detecting techniques that we have been developing for electricity networks - invariant induction and simulated ants – and a Bayesian methodology for integrating the output of these detectors. The results presented in this paper demonstrate that this technique could make a significant contribution to the security of electricity critical infrastructures.

1. Introduction

With the increasing interconnectivity between electricity management network, corporate network and the internet, electricity cyber infrastructures are becoming more and more exposed to outside attackers. This tendency has been accelerated by the widespread introduction of commercial off the shelf software and standard TCP/IP networks. Although traditional intrusion detection systems, anti-virus software and firewalls are used to protect the infrastructure, these signature-based solutions have a limited ability to detect and defend against rapidly emerging new attacks, as the spread of Slammer into a nuclear control centre

[1] Electronic Engineering Department, Queen Mary University of London, Mile End Road, London E1 4NS. Contact author: xuan.jin@elec.qmul.ac.uk

made clear [1][2]. To complement and improve the performance of these signature-based methodologies, we are working on an anomaly detecting approach that learns the normal behaviour of the system, builds up a profile and detects anomalous operations which deviate from the profile. This profile can also be used to improve state estimation, correct topology errors and inform the operators about potential discrepancies between their view of the network and its actual state – one of the factors that contributed to the recent U.S. blackout [3].

There are a number of ways in which anomaly-detecting methodologies could enhance the integrity and security of electricity data. To begin with, it could act as a useful complement to existing techniques for verifying the likely correctness of electricity measurements and give operators constant feedback about the integrity and reliability of the data. Anomaly detection detects any abnormal changes, whether they are caused by software bugs, attackers, or strange network conditions, and this would be extremely useful for bringing operators' attention to problem areas before they start to threaten the stability of the system. This would also be invaluable in the face of attackers who attempt to manipulate or corrupt the electricity data. A second application of anomaly detection is the improvement of standard protection devices, such as the IDS and virus checker. These are generally signature based, and even when a limited form of anomaly detection is deployed it is usually blind to the type of data that is being transmitted over the network. Electricity-specific anomaly detection would be much more accurate than generic techniques and it could be correlated with other security components to improve the integrity of the electricity cyber-infrastructure. This cross-correlation could also be used to reduce the large number of alerts that are generated by intrusion detection systems. Electricity-specific anomaly-detecting methodologies could also help with some of the traditional tasks carried out by electricity operators, such as the generation of pseudo-measurements and state estimation.

This paper will cover two of the anomaly-detecting techniques that we have been developing for electricity networks - invariant induction and artificial ants – and a Bayesian methodology for integrating the output of these detectors which dramatically reduces the number of false positives.

2. Invariant Detection and Artificial Ants with Bayesian Reasoning Approach

2.1 Invariant Detection in Electricity Cyber Infrastructures

The automatic identification of invariants was introduced by Ernst [4] who used test data samples to discover invariants at points in a program, e.g. immediately after entry to a procedure and prior to returning a value. In electricity networks, invariants are discovered by looking for relationships between the different data readings. This approach is particularly effective in the data from electricity networks since most of the data is interrelated in a systematic manner. For example, in the networks that we have experimented on, the relationship between the power flow readings at either end of a line are, to a high degree of accuracy, of the form P1 = kP2 + C, where k and C are constants. Some relationships, such as the one just mentioned are based on physical relationships and the unknown constant depends on the length of the line, but others could be empirical relationships that are found in the training data such as the ranges of specific power readings. As more data comes in, some of these relationships are discarded because they no longer hold and eventually one is left with a set of relationships which hold for all of the training data. The strength of using invariants is that they hold for all states of the network, which is a great advantage for complex systems which have many normal states and where learning all the different models of normality is not a viable option. One of the weaknesses of this approach is that invariants do not cover all relationships in the data, particularly those empirical relationships that are not anticipated from simple consequences of structure and may be particular to a current and perhaps long standing state (e.g. switch setting.)

We have implemented three invariant checkers. They are a Range Checker, a Linear Invariant Checker and a Bus-zero-sum Checker.

1) The Range Checker

The range checker learns a range for each reading from the normal electricity data set in the training phrase. In the detection phase, a reading P is checked to see if it is in the range or how far it deviates from the upper boundary or lower boundary: $P \in [P_{\min} - e, P_{\max} + e]$. The range threshold e measures the tolerance allowed in the range for it still to be considered normal. This approach can be powerful with many types of corruption. However it is not adaptive for network topology changes. When switch condition changes, the Range Checker may generate a lot of false positives and false negatives because the learnt normal profile is no long valid.

2) The Linear Invariant Checker

According to Ohm's Law, a linear relationship exists between two power readings on each end of a transmission and transformation line. This is not strictly true of real electricity lines, but it is approximately true in many cases. In the training phase regression coefficients are calculated against the normal electricity data set. When the Linear Invariant Checker examines the power reading on a line, it takes one reading and the previously calculated regression coefficient and calculates the reading at the other end of the line: $P_1 = a * P_2$. If the deviation from the calculated P_1 to the expected reading $\hat{P}$ is within the bounds defined by the 99% confidence interval, it will be considered as a normal reading. Notice that if any of the switches is open then both power readings should be zero. The linear Invariant is an approximate invariant in an electricity network, and so it continues to hold even when the network topology changes it do not need to be trained again. This makes the Linear Invariant Checker a good complement to the Range Checker. However because linear invariants only exist on transmission and transformation lines, their effect is limited.

3) The Bus-zero-sum Checker

According to Kirchoff's Law, current that flows into a bus must equal to the current that flows out. If the voltage is constant, the power that flows in should match the power that flows out. When the Bus-zero-sum Checker exams each bus, it adds the inflow power and outflow power together as vectors. If the result lies outside the threshold range [-e, e], then it indicates an anomaly, *i.e.:* $P_1 + P_2 + P_3 \in [-e, e]$

The Bus-zero-sum is a true invariant in the electricity network and it does not even need to be trained. The disadvantage of the Bus-zero-sum checker is that it is not very specific about which reading is corrupted because significant errors in any of the readings on one bus will cause the sum of readings on the bus to deviate from zero.

The performance of these three invariant checkers will be discussed in section 4 of this paper.

2.2 Artificial Ant Approach

In Artificial ant clustering [5], a model of normality is built up by clustering according to the real and reactive power values of a group of substations. The artificial ants move around with the assistance of pheromone trails on a two dimensional grid in which the high dimensional vectors of electricity data readings are represented as items that are picked up and dropped. A distance function is used to compare a particular item with the eight surrounding cells of the cell where the ant is located and if there is a high density of similar items, then the probability of that item being dropped or left in that location is high. Once the clusters have been learnt, they can be used to identify anomalous data by comparing the new data item with each of the cells on the grid. The new data item does not fit into the model if its dropping probability remains low over all the cells. One of the advantages of the ant-based approach is that the ants continually reshuffle all the data to optimise the clustering over time and so it exhibits a natural form of homeostasis. However, this can create problems with false positives when the degree of change is high and this is one of the reasons why we elected to integrate this anomaly detecting method with other techniques.

2.3 Bayesian Reasoning

To take advantage of the different strengths of the linear invariant and artificial ant detectors we have developed a Bayesian framework, which takes the output from the invariant and artificial ant anomaly detectors and reasons about the likelihood that they are producing a true or false positive. In this way, the limitations of the individual anomaly detecting methodologies are overcome and the accuracy of the method is substantially increased. This Bayesian approach is especially useful in electricity networks since it can combine the output of the anomaly detectors with information about the range of the readings, missing readings and known properties of the network. This more comprehensive framework greatly increases the accuracy of the results and the quality of the information that is passed to the human operator when an anomaly is detected. An extract from one of the Bayesian networks used in our experiments is shown in Figure 1. In this network the probability that reading *i* is correct is calculated by combining the invariant detector's evaluation of reading *i* with the artificial ant's evaluation of reading *i* along with other information about the network, such as the switch state and whether the sum on the bus is zero.

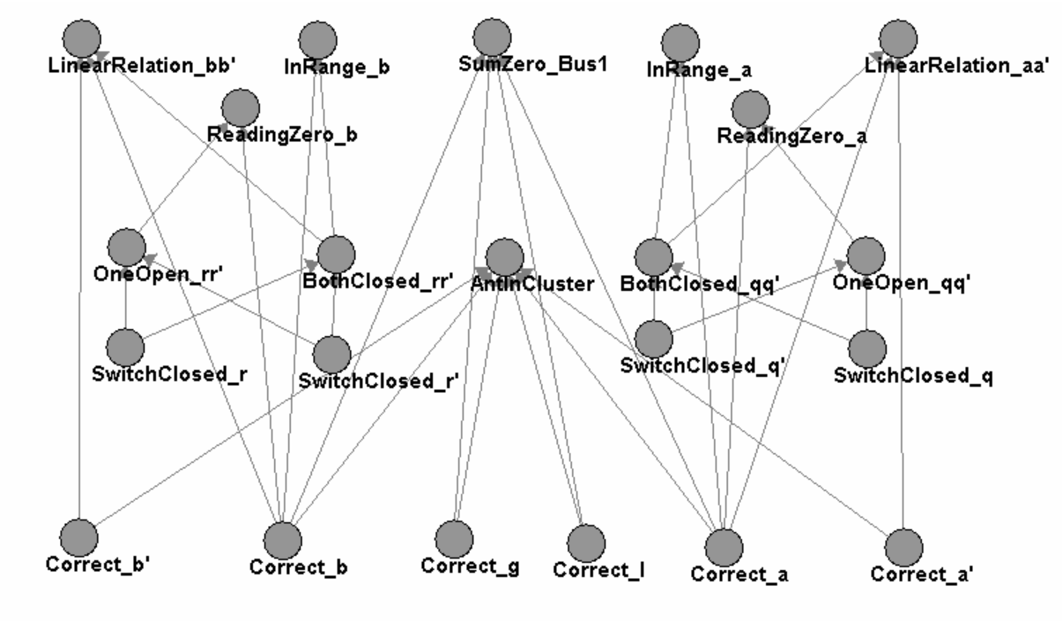

Figure 1 Bayesian network integrating information for a bus. The nodes in the bottom half are the assertions for which the beliefs are to be determined. Nodes near the top are observations about the truth of invariants and from the anomaly detectors with other information about the electricity network.

3. Previous Work

Many of the methods for identifying bad measurements and topology errors use the past state of the system to make their judgments and in this sense they could be described as a form of anomaly detection. For example, forecasting-aided state estimation (FASE) makes predictions about the future state of the system based on its past behavior and uses the difference between the latest acquired measurements and the forecast data to detect inconsistent data [6, 7]. Some systems based on this technique use statistical models and there have been a number of systems based on artificial neural networks (for example, [8, 9]), which are often combined with other techniques, such as gap statistical algorithms [10]. However, the motivation behind all of these approaches has generally been to identify and eliminate bad measurements, rather than develop a general measure of the normality level of the electricity data, which could also be used to secure the cyber-infrastructure. There has also been little application of the types of anomaly-detecting techniques which have been developed within other areas, such as artificial immune systems, artificial ants, n-grams and invariant induction.

Within SCADA systems, Balducelli et al. [11] have developed a technique for monitoring the event sequences and timing relationships of the transmission and polling of telesignals and telemeasures. This uses case-based reasoning to build up a model of the different event courses and compare them with the sequences currently in progress. This system operates on the sequences and timing information of the telesignals and telemeasures that are sent through the cyber-infrastructure and not on the electricity measurements that are transmitted and so has no notion of correctness of the data.

We have also carried out some preliminary work on anomaly detection in electricity data using some of the methods that have been developed within other areas [12]. Using data from a simple six bus network we tested the ability of the n-gram technique [13] and invariant induction [4] to build up a model of the electricity data and detect errors in it. In this simple set up the n-gram method identified 98 % of the corrupt data sets with a 1% false positive rate and the invariant induction identified approximately 85% of the corrupt data with almost no false positives. Since the true and false positives of these two methods were not coincident, we carried out some unpublished experiments, which combined the output from the two anomaly-detecting techniques using a Bayesian network. This substantially improved the overall performance.

4. Experiments and Discussion

To evaluate our methodology, we generated a year's worth of normal and corrupted electricity data on a test bed which emulates the cyber infrastructure of an electricity network in some detail. This was based on the test network developed by the IST-Safeguard project [14], but with a number of extensions and modifications. The QMUL test bed consists of a simulation of the physical electricity network, software that corrupts and filters electricity data, an emulation of a SCADA system, an emulation of the state estimation process at the control centre and a simulated operator, which monitors the state estimation and applies manoeuvres to the simulated physical network. For the state estimation and load flow calculations we used e-Agora (Advanced Grid Observation Reliable Algorithms) and the IEEE RTS 96 24-Bus network model [15]. Five different types of corruption were applied to the electricity data. These reflect the fact that electricity measurements can be altered by random noise, attacks, software bugs, meter failures, electromagnetic interference and transmission errors.

1) **Constant bias with normally distributed deviations**. This added a constant bias to the telemeasure values and a bias error sampled from independent normal distributions with zero mean. The corrupted value *X'* of the telesignal *X* is computed according to equation (1), where $E_{N(0,1)}$ is a standard random value following a *N(0,1)* distribution and sigma (*S*) is standard deviation of the bias error (*B*).

$$X' = (X + B) * E_{N(1,S^2)} = (X + B) * ((E_{N(0,1)} * S) + 1) \quad (1)$$

2) **Loss of decimal point (Mantissa).** This simulates the situation in which an analogue value loses information about the position of the decimal point. The corrupted value of the telemeasure is calculated by 'taking out' the decimal point.

3) **Sign switch.** This type of corruption is used to simulate the situation in which an analogue telemeasure changes its sign.

4) **Fixed at fixed value.** The value of a telemeasure is 'frozen' for one week, instead of changing along with the rest of the electricity data. This could result from a software bug or meter failure.

5) **Fixed at random value.** The value of a telemeasure is replaced by a random value, which is held constant for a week. This could result from a software bug, meter failure

or a deliberate attempt to corrupt the system.

This data generated by QMUL electricity test bed was used to train and test the invariant and the artificial ant anomaly detectors; both individually and with their output integrated using the Bayesian network. This data has also been made available online for the benefit of other researchers [16].

4.2 Results of the Range Checker

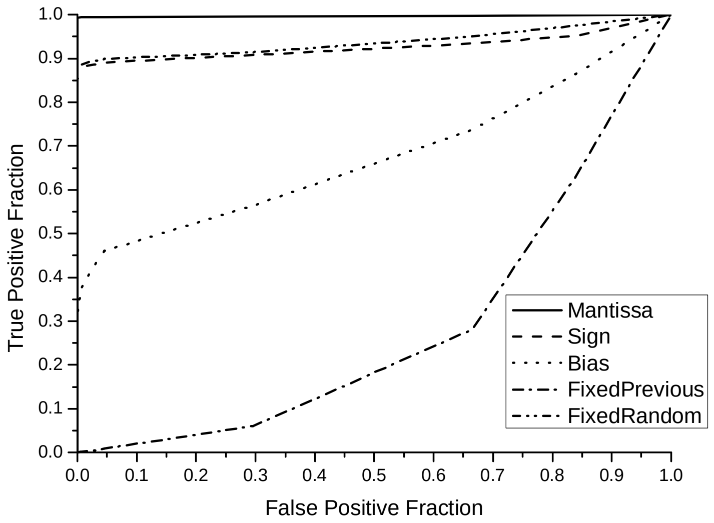

Figure 2 ROC curve of Range Checker on five types of attack corruptions

From the ROC curve, it can be seen that the Range Checker is good for Mantissa, Sign Switch and Fixed to Random Value types of corruption, which shows that it is capable of detecting obvious abrupt changes. The Range Checker does not perform as well on Sign Switch corruption as for the Mantissa corruption because when small, near zero values are corrupted by a sign switch, the changes are too subtle to be detected by the Range Checker.

The Range Checker does not perform well on Bias corruption because the bias and sigma values (5 and 0.15 respectively) do not lift the corrupt readings far enough above the background noise to be detected by a simple range check. The Range Checker performs poorly on Fixed to Previous Value corruption because fixed previous values are still within the learnt ranges, even though they are inconsistent with the rest of the electricity data.

4.3 Results of the Linear Invariant Checker

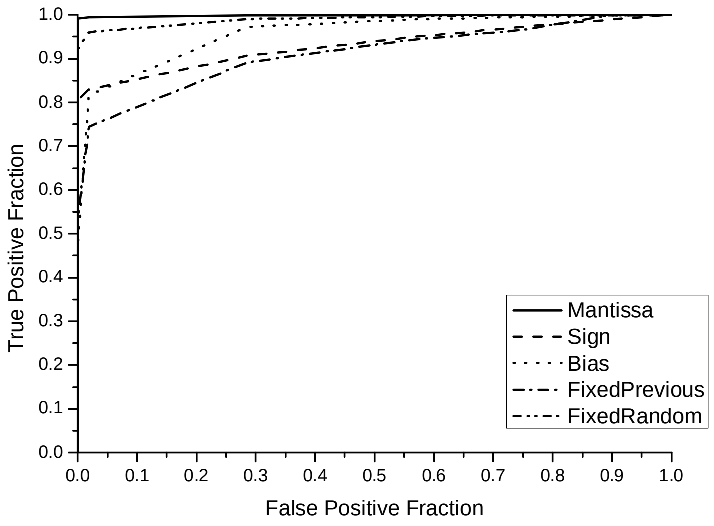

Figure 3 ROC curve of the Linear Invariant Checker on five types of attack corruptions

The ROC curve of the Linear Invariant Checker shows reasonable performance on all the types of corruption. The limitation of this technique is that linear relationships only hold on the real power readings at either end of the lines and so it cannot be used to detect errors in the other readings. It is also unable to identify which of the two power readings is causing the linear relationship to fail.

4.4 Results of the Bus-zero-sum Checker

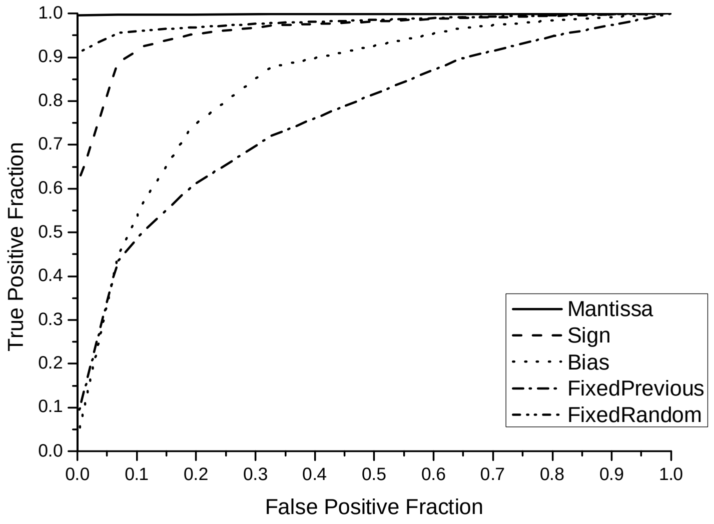

Figure 4 ROC curve of the Bus-zero-sum Checker on five types of attack corruptions

Like the Range Checker, the Bus-zero-sum Checker did not work well against sigma and bias corruption because the changes introduced by these quite close to the background noise and so they can only be picked when the threshold is extremely strict, *i.e.* the checker is very sensitive, which caused a high false positive rate. The Bus-zero-sum Checker had a poor performance on the Fixed to Previous value corruption, probably because the "stuck" power readings are still close to their actual value.

4.5 Results of Ant Clustering Anomaly Detection

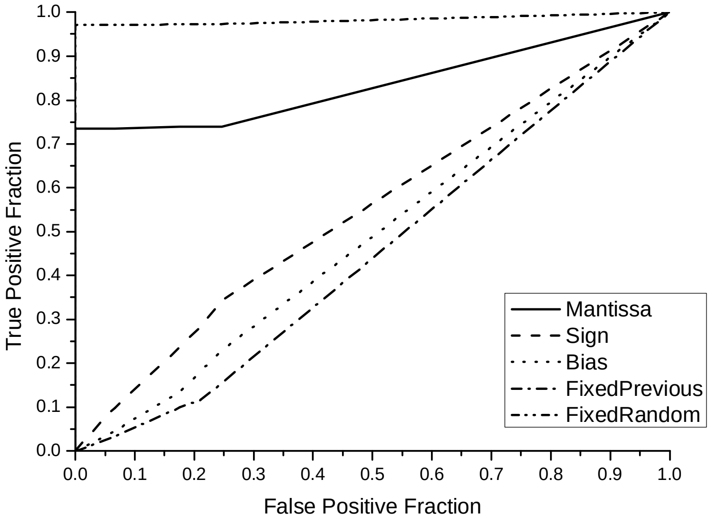

Figure 5 ROC curve of Ant Clustering Anomaly Detection

The results for Fixed Previous corruption showed true positive detection to be consistently above 97% for the ant clustering. This is an indication of the early promise of ant clustering for electricity specific anomaly detection. Although mantissa corruption detection is lower than all the other forms of anomaly detection tested, since ant clustering is intended to complement and not replace other forms of anomaly detection this could be accepted but is likely to be improved. Sign switch corruption detection is also lower than expected given the obviousness of the corruption. The detection rates for Mantissa and Sign may be a result of the normalising that the ants apply to the data, which does not tolerate linear changes. We are working on a way for ants to better recognise big discrepancies between data points while retaining normalising tendencies.

The advantage of Ant Clustering is that it is able to process diverse types of data in the same way. However, the distance function possibly needs to be more electricity data specific to improve the detection of more subtle corruption such as Bias corruption. At the same time it is important to retain the incremental features of the ants as they are able to adapt the normal model in accordance with changes in the state of the electricity network through the seasons.

4.6 Results of Bayesian Network Reasoning

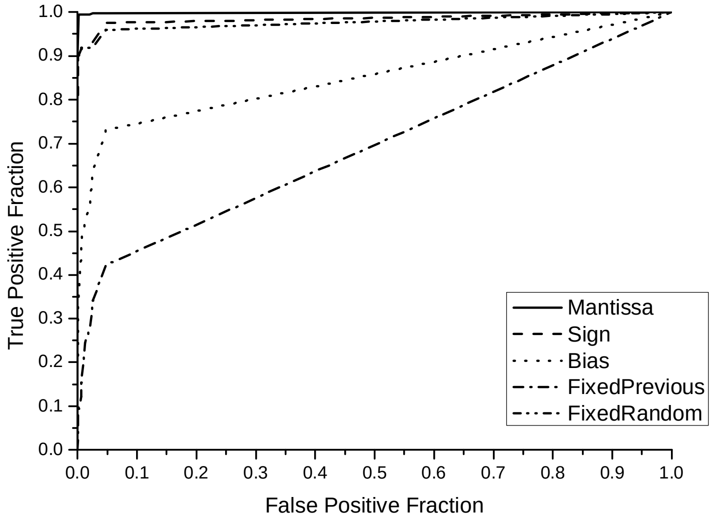

Figure 6 ROC curve of Bayesian networking reasoning with the Range Checker, the Linear Invariant Checker and the Bus- zero-sum Checker

Although these results look significantly worse in some respects than those for the zero sum and linear invariant, these results are for the identification of corruptions in individual power readings, whereas the linear invariant induction results are only for the detection of errors in pairs of readings, the ant results are for a collection of readings and the zero sum results are for the detection of errors in all of the readings on a bus. This means that it is the comparison with the range checker that is the most important, which is the only other methodology to work on all of the readings. The most obvious gains from using the Bayesian network are that false positive rate has decreased and the true positive rate for detecting Fixed to Previous Value corruption has greatly improved. The Range Checker used to be blind to this type of corruption but with the help from the other two checkers and Bayesian network, the true positive rate has substantially increased. There has also been a significant improvement in the true positive rate for the Bias and Sigma corruption.

5. Conclusion

This paper has outlined some of the contemporary threats to the electricity cyber infrastructure and investigated two novel approaches to anomaly detection on electricity data. One uses underlying physical knowledge to express and capture invariants in the data and the other uses artificial ant clustering to capture relationships between elements that are not covered by the invariants and which are more subject to change as the system configuration changes. A Bayesian network has been used to integrate the output from these three anomaly detectors and results in less false positives and significantly improved detection rates for some of the corruption types.

Acknowledgements

We would like to thank the other members of the Safeguard project for their collaboration on this work. These include Wes Carter (QMUL), Simin Nadjm-Tehrani (LIU), Kalle Burbeck (LIU), Tobias Chyssler (LIU), Stefan Burschka (Swisscom), Michael Semling (Swisscom), Thomas Dagonnier (Swisscom), Giordano Vicoli (ENEA), Sandro Bologna (ENEA), Luisa Lavalle (ENEA), David Moyano (AIA), Carlos López Ullod (AIA) and Oleg Morajko (AIA). We would also like to thank the European IST Programme (IST Project Number: IST-2001-32685) and the EPSRC for their support.

References

[1] Poulsen, K. (2003). Slammer worm crashed Ohio nuke plant network. *SecurityFocus News: http://www.securityfocus.com/news/6767*

[2] North American Electricity Reliability Council (2003). SQL Slammer Worm: Lessons Learned for Consideration by the Electricity Sector Slammer. *available at: http://www.esisac.com/publicdocs/SQL_Slammer_2003.pdf*

[3] U.S.-Canada Power System Outage Task Force (2004). Final Report on the August 14, 2003 Blackout in the United States and Canada: Causes and Recommendations.

[4] Ernst, M.D. (2000). Dynamically Discovering Likely Program Invariants. *PhD thesis*, Dept. of Computer Science and Eng., Univ. of Washington, Seattle, Washington.

[5] Ramos, V., Merelo, J.J. (2002). Self-Organized Stigmergic Document Maps: Environment as a Mechanism for Context Learning. In *Proceeding of 1st Spanish Conference on Evolutionary and Bio-Inspired Algorithms*, Merida, Spain.

[6] Do Coutto Filho, M.B., Duncan Glover, J., Leite da Silva, A.M. (1993) "Forecasting-aided state estimation", In *Proc. IIth PSCC*, Vol. 2, pp. 689495. Avignon

[7] Leite da Silva, A.M., Do Coutto Filho, M.B., de Queiroz, J.F. (1983). State forecasting in electrical power systems. *IEE Proc. C*, Vol. 130, No. 5, pp. 237-244

[8] Alves da Silva, A.P., Leite da Silva, A.M., de Souza, J.C.S., do Coutto Filho, M.B.(1993) State forecasting based on artificial neural networks, *Proc. IIth PSCC*, pp. 461467

[9] Souza, J.C.S., Leite Da Silva, A.M., Alves Da Silva, A.P. (1996). Data Debugging for Real-time power System Montering Based on Pattern Analysis. *1996 IEEE PES winter meeting*, p. 96, Baltimore

[10] Huang, S., Lin, J. (2002). Enhancement of Power System Data Debugging Using GSA-Based Data-Mining Technique. *IEEE Transactions on Power Systems*, Vol. 17, No. 4

[11] Balducelli, C., Lavalle, L., Vicoli, G. (2004). Novelty detection and management to safeguard information intensive Critical Infrastructures. In *Proceeding of 11th Annual Conference of The International Emergency Management Society*, Melbourne, Australia,

[12] Bigham, J., Gamez, D., Lu, N. (2003). Safeguarding SCADA Systems with Anomaly Detection. *Lecture Notes in Computer Science*, vol. 2776, pp.171-182. Springer Verlag,

[13] Damashek, M. (1995). Gauging Similarity with n-Grams: Language-Independent Categorization of Text. *Science*, Vol. 267, pp. 843-848.

[14] IST Safeguard project website. http://www.ist-safeguard.org/

[15] Grigg, C., Wong, P., Albrecht, P., Allan, R., Bhavaraju, M., Billinton, R., Chen, Q., Fong, C., Haddad, S., Kuruganty, S., Li, W., Mukerji, R., Patton, D., Rau, N., Reppen, D., Schneider, A., Shahidehpour, M., Singh, C. (1999). The IEEE Reliability Test System-1996: A report prepared by the Reliability Test System Task Force of the Application of Probability Methods Subcommittee. *IEEE Transactions on Power Systems* , Vol. 14, Issue 3 , pp.1010 – 1020.

[16] Electricity data download website: *http://www.elec.qmul.ac.uk/electricitydata*

Author Biographies

Xuan Jin graduated from Beijing University of posts and Telecommunications of China in 2001. Then he received his MSc with distinction in e-Commerce Engineering at Queen Mary, University of London in 2002. He is now a PhD research student in Electronic Engineering Department, Queen Mary University of London. His research interest includes intrusion detection system and information correlation.

John Bigham studied Mathematics and Statistics at Edinburgh University. After working in industry as a statistician for several years he obtained an M.Sc. in cybernetics and a Ph.D. in artificial intelligence and now is a Reader in the Department of Electronic Engineering, Queen Mary University of London. His current research interests are in application level security and resource management of wireless networks.

Julian Rodaway is a PhD research student at Queen Mary, University of London in the Electronic Engineering Department. His research interests include anomaly detection and agent technology.

David Gamez is a Researcher at Queen Mary, University of London and he is also studying for a PhD in computer science at the University of Essex. His research interests include anomaly detection, agent technology, robotics, machine consciousness and knowledge engineering.

Chris Phillips received a B.Eng. degree in Telecommunications Engineering from Queen Mary College in 1987 followed by a Ph.D. on concurrent discrete event-driven simulation. After working at Bell Northern Research, Siemens, and Nortel Networks, Dr. Phillips returned to Queen Mary as a Reader to explore Internet protocols and application developments, including the infrastructure to support rapidly emerging e-commerce systems.