Dynamic Trust Management of Semi-Automated Complex Systems

John Bigham, Xuan Jin, David Gamez, Ivan Djordjevic, Chris Phillips
Department of Electronic Engineering,
Queen Mary University of London,
253 Mile End Road,
London E1 4NS, UK

ABSTRACT

Supervisory Control and Data Acquisition (SCADA) systems are widely used in industry to manage systems ranging from electricity networks to chemical plants. However they are vulnerable to security breaches from insiders (the people operating the system). In addition, the increased complexity of the managed system may result in unreliable and false feedback to the operator. This paper describes a framework for automated trust management and privilege re-allocation within the monitoring system, based on the feedback about the performance of the monitored system. The framework uses an agent-based architecture for system monitoring, together with Bayesian networks and workflows for reasoning about the trust levels of different actors in the system.

Keywords: Workflow, Correlation, Trust, Intrusion Detection

1. INTRODUCTION

Within most industries the usual management policy is to trust operators by default and attempt to rectify problems when they arise. However this is not ideal since insider attacks are often a problem. A recent survey [7] reports that 48% of large businesses blame their worst security incident on insider activity – and the increased mobility of the workforce in deregulated industries has only increased the probability of insider attacks [14]. The same default trust policy applies to the management system itself. Once it has been installed and tested, operators generally trust what it says about their system and believe that it has carried out the requested actions upon this system. Again, this trust can often be misplaced. For example, one of the contributing factors to the recent U.S. blackout was the fact that the operators were acting on the basis of information that was in some cases several hours out of date [18].

What is needed to address to these difficulties is a more dynamic model of trust that can evaluate both the trustworthiness of the operators and the system that they are operating. This paper describes the research that we are pursuing on a workflow trust management architecture that uses agents to gather information about the system and actuate changes on it in response to beliefs about the trustworthiness of the operators and their actions.

Workflows have been used in business for a number of years to model the flow of information within an organisation and the operations carried out on that information. In our research we are applying workflows to monitor both normal and abnormal flow of activities within an organisation and to build up a model of the trustworthiness of the system and the operators using it, independently of whether workflows are already in place. Each action within a workflow has certain consequences and one of the aims of our work is to anticipate what will happen as the result of each action on the system. This information is gathered using an agent system that looks for anomalies as well as pre-defined events. From this we can evaluate the trustworthiness of actors, the trustworthiness of actions and the trustworthiness of the system as a whole.

The methodologies that we are developing are applicable to any complex system involving a number of actors and actions. However, in the first instance our work is being applied to the management of trust within electricity networks operated by Supervisory Control and Data Acquisition (SCADA) systems and this will be the main focus of this paper. The methodologies described in this paper are being developed as part of the SAFEGUARD project [17] and will be used to protect large complex critical infrastructures such as electricity and telecommunication networks.

2. SCADA SYSTEMS

In an electricity distribution network, information about the physical network is gathered using a large number of sensors attached to breakers, voltage meters and so on. This information is passed on to remote terminal units (RTUs), which in turn pass it on to concentrator devices placed in a wide area network. A number of control centres also sit on this wide area network, gathering and processing the data and sending out control signals across the network back to the RTUs. These control centres typically contain a local area network with machines running databases and specialised software, such as the SCADA software, energy management systems and state estimators, and these local networks often have connections to the corporate network as well as software and hardware vendors.

In an electricity network the standard operations include bringing generators in and out of service, adjusting the voltage level and phase angle using reactors and capacitors, shedding load and general maintenance tasks such as repairing or replacing breakers and transformers. Each of these tasks involves one or more human operators and a large number of functioning components. A diagram of a typical electricity management network is shown in Fig. 1.

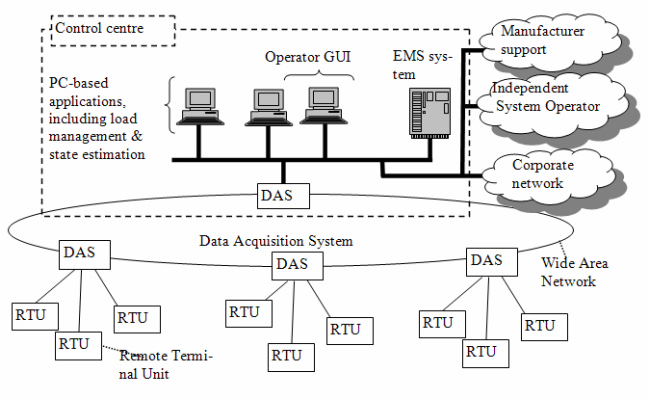

Fig. 1. Electricity Management Network

3. WORKFLOW

3.1. Workflow Overview

Workflows are defined by the Workflow Management Coalition as follows: " The automation of a business process, in whole or part, during which documents, information or tasks are passed from one participant to another for action, according to a set of procedural rules" [9]. The advantage of workflows is that they are very good at tracking a complex series of temporal events and it is also easy to automatically carry out actions at different points in them. It is for this reason that we have chosen to use them to evaluate the trustworthiness of the operators and management system.

3.2. Workflow Management System

Workflow management systems are used to define manage and execute workflows using software whose order of execution is driven by a computer representation of the workflow logic [19]. In our research, the workflow management tool mainly plays the role of executing and monitoring predefined workflows. There are many workflow management tools available such as Cosa [4] and OpenWFE [15]. We use the Bossa Workflow System [1] because it has the following advantages:

1) Bossa uses an extended Petri net model which provides an intuitive way of modelling workflows and a way to verify workflow correctness [28]. Extended Petri nets even allow users to model time and include a hierarchy of workflow models.
2) Bossa is designed to be embedded and it is easy to define and dynamically load workflows in Bossa.
3) Bossa is written in Java, which can be platform independent.
4) Bossa is lightweight and fast.

Bossa workflow system is integrated into the correlation agent, which controls the start, execution and monitoring of each workflow instance.

4. AGENT SYSTEM

The agent system dynamically monitors the operators and system components within an electricity management network in order to evaluate their performance. Operators who create a large number of anomalies, perform unusual actions, or create instability in the system will have their actions treated more circumspectly. In addition problems within the system, such as anomalous or out-of-date data or unreliable software, will be identified. This information can be either passed to the operator or automatically acted upon, in order to prevent or limit inappropriate behaviour. The most important agents will now be covered in more detail.

4.1. Hybrid Detector Agents

Hybrid detector agents (HDA) are effectively sensors that are used to gather information about operators and their activity and the behaviour of the system. Typically, their role does not exceed passive monitoring, although some may perform certain actions on the managed system, but only if explicitly permitted by the action agent. HDA agents combine known information with a dynamic model of the system's normal behaviour. A large number of different types of dedicated agents are placed in the system to monitor many aspects of system activity:

1) *Event course anomaly detector*. This looks at the sequences of events within the control centre and the SCADA system. The sending and receiving of data and control messages involves a whole sequence of tasks with their own timing constraints. This agent builds up a model of these sequences using case-based reasoning and identifies anomalies.
2) *Keystroke anomaly detector*. This examines the keystroke patterns of the different operators. This has been shown to be a reliable biometric identification mechanism provided the person is consistently using the same keyboard. [12]. Significant anomalies in an operator's keyboard patterns could indicate that someone else is using their terminal or password.
3) *Data hybrid detector*. This agent builds up a model of the normal data patterns being passed by the SCADA system about the electricity network. Large deviations from this model could indicate that a trusted operator is behaving abnormally or that an inaccurate model of the electricity network is being presented to the human operator.
4) *Network anomaly detector*. This monitors the connections between the machines and the traffic levels to identify unusual patterns in the network.

4.2. Wrapper Agents

Anomaly information is extremely useful, but it needs to be combined with other information in order to reach a definitive conclusion. This is gathered from wrapper agents attached to parts of the SCADA system that evaluate what control actions each operator is attempting to do. There are also wrappers attached to the intrusion detection system that provide information about possible attacks on the system including attempts to escalate privileges by trusted insiders.

4.3. Correlation and Action Agents

As mentioned before, the correlation agent contains the embedded workflow system and is responsible for integrating information from the different agents, reasoning about the state of the network using Bayesian networks and firing transitions

within the workflow based on this reasoning. Some of these transitions are used to control the action agent. In this way the correlation and action agents work together to provide a quick response that rectifies problems as they arise. The available responses include adjusting the privilege level of, for example, the human operators, or changing the system topology (e.g. detaching of part of the network, or replacing a major switch).

The correlation passes tasks on to the action agent by issuing attribute certificates [8] that specify the set of privileges and the target object that an action agent needs to interact with. Due to a nature of the system, where instant and appropriate response is essential, these certificates have a short validity period or are single-use only. This alleviates the need to maintain Certificate Revocation Lists (CRL) and improves the scalability of the agent communication system. When in possession of a valid certificate, action agents are permitted to perform tasks.

4.4. Man Machine Interface Agent

The Man Machine Interface (MMI) agent is used to manage the agents and define the scope of their legitimate activity. It represents the root authority in the local Privilege Management Infrastructure (PMI) [3], responsible for setting the privilege level of the correlation and action agents, as well as those of the operators. As such, it can also modify the level of the privileges that is granted to each of these entities, based on global observations of the system.

4.5. Functioning of the System

The generic architecture of the agent system and the system being monitored is given in Fig. 2. Different hybrid detector agents are positioned in the system based on the type of the activity they are monitoring. Information from these is passed on to the correlation agent, which makes an assessment of the trustworthiness of actors, actions and the system components. Based on this, the privileges of the operators or topology of the system is modified over time by the action agent with the authorisation of the administrator operating through the man machine interface agent.

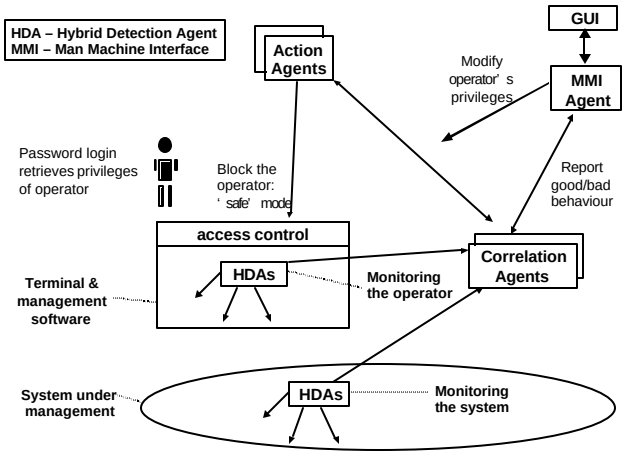

Fig. 2. System monitoring through agent interaction

.5. DYNAMIC TRUST MANAGEMENT CASE STYDY

The benefits of our approach will now be illustrated through a transformer maintenance scenario taken from an electricity network.

This scenario covers a situation in which an untrustworthy operator makes a series of adjustments to the system that progressively put it into an unstable state. In an electricity network transformers are often oil cooled. If this cooling system ceases to function properly then the transformer will overheat when the load is high. For this reason, a repair engineer has the ability to reduce the power rating of such a transformer pending replacement or repair of the cooling system. This is done by adding a constraint to the network data model that limits the load through this transformer to a value that is considerably lower than the value when the transformer could be cooled. This network data model is used by the electricity management system to decide what control actions to take, and so if this constraint is added when it is not in fact true, then it could be used as a mechanism for a malicious attack. If the power rating of several transformers is altered in this way, perhaps by a person who does not normally do this kind of adjustment, then this could force the operators to limit the flow of electricity between sections of the network and incorrectly restrict the control actions that they can make.

This scenario presents a double trust problem. Operators incorrectly trust the maliciously adjusted transformer power ratings and the system incorrectly trusts the malicious operator who makes adjustments to the system. Our architecture deals with these trust violations by tracking both the anomaly level of the operators and the anomaly level of the actions carried out by the operators on the system. This information is fed into a workflow that is set up to track the progress of each transformer maintenance operation and which checks for trust violations at each stage of transformer maintenance. Fig. 3 shows the workflow for trust management in transformer maintenance. Simple Bayesian belief networks are constructed to control work item routing at OR-split transitions of the workflow. In this example, the Bayesian network B3 is used to evaluate the probability of the operator being "honest". There are three observables in the Bayesian network. Each observable represents a factor that reflects an operator's honesty. If the probability of the operator being honest is below a certain threshold, for example 0.6, we rate the operator as not trustworthy. This Bayesian network controls the workflow routing at three transitions:

1) At transition "Change Settings in Database". To begin with this workflow will expect only certain members of the repair work group to perform the maintenance actions and updates of the network data model. If the keystroke patterns of the person who is updating the transformer database do not match those of the operator who usually carries out this operation an impostor will be suspected and the correlation agent can issue a certificate to the action agent instructing it to restrict the privileges of the operator.

2) At transition " Check Update Logs". Checks on the update logs to the network data model over the recent weeks

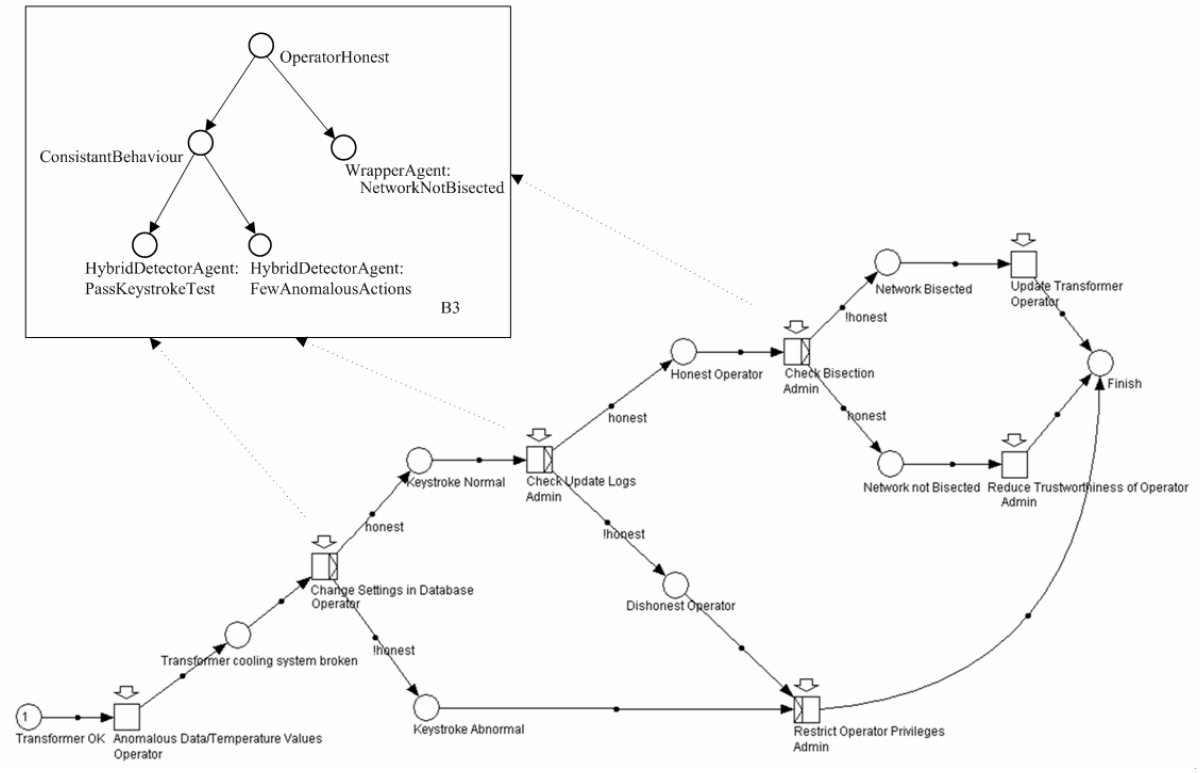

Fig. 3. Workflow for trust management in transformer maintenance

will be used by the workflow to establish how many similar actions any particular repair engineer performed. A operator that has caused many abnormal adjustments to the system in the recent past should be treated with suspicion

3) At transition " Check Bisection". The SCADA wrapper agent will provide the correlation agent with information about the degree to which the network has become bisected by the transformer power rating adjustments. The higher the risk of bisection, the lower the belief in the trustworthiness of the engineer needs to be before taking action to inspect the transformers or accelerate repairs. If the actions of the repair engineer(s) cause increasing bisection of the system, there will be reason to reduce the privileges of the engineer.

When the Bayesian network B3 is working out the operator's honesty, the threshold may change. This is because at certain transitions, not all the observables will be observed. For example, at transition "Change Settings in Database", only the observable "PassKeystrokeTest" will be observed, while at the transition " Check Bisection" all three observables will be observed. So the threshold of the operator to be honest will vary for each transition.

6. SIMULATION PLATFORM

To test our system we are currently building a simulation platform as shown in Fig. 4 This uses a mixture of real and simulated electricity management software to manage a simulated electricity network. A number of trust breaching scenarios are being defined and executed using an attack tool developed by one of the partners. Our agent system uses the methodologies outlined in this paper to identify the breaches of trust within the normal operations of this network and respond appropriately to them.

Load changing scripts alter the electricity network and normal use scripts control the electricity network as it would be controlled by an operator under normal circumstances. The tests are controlled by an attack tool developed by one of our project partners, which alters the properties of the operators or the system. It is then the task of the correlation agent to adjust its trust model in response to these changes in behaviour.

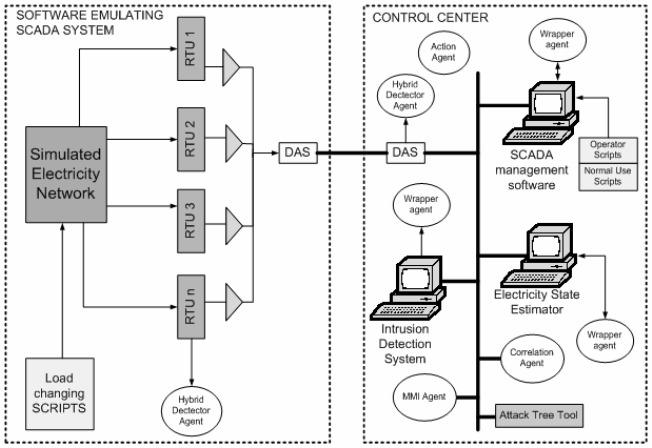

Fig. 4. Simulation platform.

7. RELATED WORK

A major influence in the development of our system has been the work of Konstantin Knorr [11], who used Petri nets to model the flow of confidential information within business processes. A second related influence is the recent research that has been carried out by Ning et. al [13] among others on correlation in intrusion detection. The major problem in intrusion detection at the moment is that a very large number of alarms are generated by each attack, which have to be processed manually by an operator in order to identify an underlying cause. By modelling the stages of an attack it is possible to link a large number of alarms with the underlying activities of an attacker. Agents have also been used in a number of intrusion detection systems (see [2] and [9] for example) to detect and respond to break-ins by external attackers. However the focus in this work is on the mistakes and attacks of insiders and the trustworthiness of the management system. The problem of external attacks fits within our general approach, but it is not our main focus. In addition, none of the current agent-based intrusion detection systems have been applied to SCADA systems and there has been no use of the Petri-net modelled workflows and a Bayesian network correlation mechanism as deployed in our system.

Although not the main focus of this paper, it is important to note that in the proposed system, trust is supported through controlled access based on the authorisation policies, as well as certification-based trust. KeyNote [23] and Simple Public Key Infrastructure (SPKI) [24] are examples of trust management systems aiming to control security of distributed environments by delegation of the permissions through the use of credentials.

8. CONCLUSIONS AND FUTURE WORK

This paper describes a framework for the automated trust management based on agent architecture for system monitoring and Bayesian Networks and workflows for reasoning about the trust levels of different actors in the system. Although motivated by the application within the electricity networks operated by SCADA, the framework presented can be easily extended to cover any complex system involving a number of actors and actions.

One such extension of this proposed framework would be within the management system for the virtual collaborative environments, developed in the EE department of QMUL [5], [6]. In this system, dynamic formation and self-management of the large groups of peer clients is supported via few security manager nodes. Security perimeters of the collaborative groups are supported by centralised definition and distributed enforcement of security policies, achieving protection of the group from outsiders as well as group members from each other within the collaboration. However, one of the realistic concerns is the clients' implicit trust in their security manager, and their limited ability to protect themselves and the group from a malicious manager. One item of further work is to apply the proposed agent system to monitor the managers' actions and corresponding reactions within the collaboration environment, and to evaluate them by the means of workflows and Bayesian networks, in order to reduce vulnerability 'hot-spots' within the system.

A second area that will be addressed by future work is the security robustness of our architecture. Electricity management networks are mission critical systems and mechanisms need to be put in place to reduce the chances of an agent being compromised together with the means to quickly detect and repair security breaches in the agent system when they occur. We plan to address these problems by applying the techniques outlined in this paper to the interactions between the agents. While our system is running, correlation agents will be in contact with a number of action agents (and vice versa). Other interactions also take place between the action and wrapper agents at the time of authorisation enforcement. Feedback from these interactions will be periodically collected by the MMI agent and combined in order to assess the confidence in the agents' performance. For this evaluation process, a number of existing trust models (e.g. [25], [26], [27]) are currently under the consideration.

ACKNOWLEDGEMENTS

We would like to acknowledge the other members of the Safeguard project. These include Julian Rodaway (QMUL), Wes Carter (QMUL), Stefan Burschka (Swisscom), Simin Nadjm-Tehrani (LIU), Kalle Burbeck (LIU), Giordano Vicoli (ENEA), Sandro Bologna (ENEA), Claudio Balducelli (ENEA) and Carlos Ló pez Ullod (AIA). We would also like to thank the European IST Programme for supporting this project.

REFERENCES

[1] Bossa website. http://www.bigbross.com/bossa/

[2] J. S. Balasubramaniyan, J. O. Garcia-Fernandez, D. Isacoff, E. Spafford, and D. Zamboni, "An architecture for intrusion detection using autonomous agents", COAST Technical Report 98/ 05, 11 June 1998.

[3] D. Chadwick et al., "The PERMIS X.509 role based privilege management infrastructure," 7th ACM Symposium on Access Control Models and Technologies, June 2002.

[4] Cosa website http://www.transflow.com/english/

[5] I. Djordjevic, C. Phillips, "Architecture for secure work of dynamic distributed groups," Proc of IEEE Consumer Communication and Networking Conference, Las Vegas, Nevada, USA, January 2004

[6] I. Djordjevic, C. Phillips, "Certificate-based distributed firewalls for secure e-commerce transactions," Journal of the Institution of British Telecommunications Engineers, vol. 2, part 3, 2001, pp. 14-19.

[7] DTI (Department of Trade and Industry, UK), |Information Security Breaches Survey 2002," available at https://www.security-survey.gov.uk/ isbs2002_detailedreport.p df

[8] S. Farrel, R. Housley, "An Internet attribute certificate profile for authorization," RFC 3281, Network Working Group, IETF, April 2002.

[9] L. Fischer (ed), "The Workflow Handbook 2003," Published in association with the Workflow Management Coalition (WfMC), 2003.

[10] V. Gorodetski, O. Karsaev, A. Khabalov, I. Kotenko, L. Popyack, V. Skormin, "Agent-based model of computer network security system: a case study," Proceedings of the International Workshop "Mathematical Methods, Models and Architectures for Computer Network Security, Lecture Notes in Computer Science, vol. 2052, Springer Verlag 2001, pp.39-50.

[11] K. Knorr, "Multilevel security and information flow in Petri net workflows', Proceedings of the 9th International Conference on Telecommunication Systems, Modeling and Analysis, Special Session on Security Aspects of Telecommunication Systems, Dallas, TX, March 2001 pp. 9-20.
[12] F. Monrose, M. K. Rieter, S. Wetzel, "Password hardening based on keystroke dynamics," Proceedings of Sixth ACM Conference on Computer and Communication Security, Singapore. 2-4 November 1999.
[13] P. Ning, Y. Cui, and D. S. Reeves, "Analyzing intensive intrusion alerts via correlation," Proceedings of RAID 2002, Lecture Notes in Computer Science vol. 2516, 2002 pp. 74-94.
[14] P. Oman, E. Schweitzer, and J. Roberts, "Safeguarding IEDs, substations, and SCADA systems against electronic intrusions," available at: http://tesla.selinc.com/techpprs.htm
[15] OpenWFE website http://www.openwfe.org/index.shtml
[16] Petri nets world website http://www.daimi.au.dk/PetriNets/
[17] Safeguard website http://www.ist-safeguard.org
[18] U.S.-Canada Power System Outage Task Force, Interim Report, "Causes of the August 14th blackout in the United States and Canada," available at: ftp://www.nerc.com/pub/sys/all_updl/docs/blackout/814BlackoutReport.pdf
[19] Workflow Management Consortium website http://www.wfmc.org
[20] WFMC, Workflow Management Coalition Terminology and Glossary. Technical report, Workflow Management Coalition, Brussels, 1996.
[21] R. Robbins, "Distributed intrusion detection systems: an introduction and review," SANS InfoSec Reading Room, February 2003.
[22] R. Janakiraman, M. Waldvogel, Q. Zhang, "Indra: a Peer-to-Peer approach to network intrusion detection and prevention," Proc of IEEE International Workshops on Enabling Technologies, June 2003.
[23] M. Blaze, J. Feigenbaum, J. Ioannidis, A. Keromytis, "The KeyNote trust-management system," version 2. IETF, RFC 2704, September 1999.
[24] C. Ellison, B. Frantz, B. Lampson, R. Rivest, B. Thomas, T. Ylonen, "SPKI certificate theory," IETF, RFC 2693, September 1999.
[25] T. Moreton, A. Twigg, "Enforcing collaboration in Peer-to-Peer routing services," Proc of 1st International Conference on Trust Management, Heraklion, Crete, Greece, LNCS 2692, May 2003, pp. 255-270
[26] A. Josang, "A logic for uncertain probabilities," International Journal of Uncertainty, Fuzziness and Knowledge-Based Systems, June 2001, pp. 279–311.
[27] A. Abdul-Rahman, S. Hailes "Supporting trust in virtual communities," Proc of 33 Hawaii International Conference on System Sciences (HICSS), Maui, Hawaii, 4-7 January 2000
[28] W.M.P. van der Aalst, "The application of Petri nets to workflow management," J. of Circuits, Systems, and Computers, 8(1): pp. 21-66 1998