

# Data Mining for Security and Crime Detection

Gerhard PAAß<sup>1</sup>, Wolf REINHARDT, Stefan RÜPING, and Stefan WROBEL

*Fraunhofer Institute for Intelligent Analysis and Information Systems, Sankt Augustin, Germany*

**Abstract.** As the internet becomes more pervasive in all areas of human activity, attackers can use the anonymity of the cyberspace to commit crimes and compromise the IT infrastructure. As currently there is no generally implemented authentication technology we have to monitor the contents and relations of messages and internet traffic to detect infringements. In this paper, we present recent research on internet threats aiming at fraud or hampering critical information infrastructure. One approach concentrates on the rapid detection of phishing email, designed to make it next impossible for attackers to obtain financial resources or commit identity theft in this way. Then we address how another type of internet fraud, the violation of the rights of trademark owners by faked merchandise, can be semi-automatically solved with text mining methods. Thirdly, we report on two projects that are designed to prevent fraud in business processes in public administrations, namely in the healthcare sector and in customs administrations. Finally, we focus on the issue of critical infrastructures, and describe our approach towards protecting them using a specific middleware architecture.

**Keywords:** internet security, email fraud, phishing, trademark infringement, counterfeit merchandise, internet auctions, critical infrastructure

## Introduction

Classically, the most severe and dangerous security threats to countries, organizations and individuals are considered to be physical acts of violence, such as the ones we have sadly had to observe in recent years. Consequently, besides classical police work, efforts in the online realm have concentrated strongly on finding individuals and organizations who are offering material that incites towards such acts of violence, or who are using electronic means to arrange for committing them. At the same time, however, it has become clear that the nature of the threats, and the structure of organizations and individuals that deal with the threat has started to change.

Firstly, the distinction between criminal activity carried out for purely financial gain and security threats motivated by political or ideological reasons is beginning to dissolve. More and more, we are seeing that the internet is being used to provide financing for politically or ideologically motivated offences, and that techniques that up to now have been associated only with financial crime, such as *phishing*, are being used in this context. Drying up these sources of funds through measures designed to prevent *phishing* and other financial fraud on the internet thus becomes an important contribution not only to fighting crime in general, but also to defend against larger scales security threats. Secondly, the nature of the threats is changing dramatically. While previously, explosives or other chemical, biological or nuclear weapons were needed to seriously threaten the general public or infrastructure of a country, recent examples such as the attacks against Estonia in the spring of 2007 [Landler, Markoff 07; Rhoads 07] show that nowadays, critical infrastructures can be destabilized by entirely non-physical methods, simply by attacking country's infrastructure with suitable means.

In this paper, we therefore present recent research that is not primarily directed at identifying individuals conspiring to threaten countries or organizations, but which is focused the general internet threats that provide the financial basis for such activities. In particular, in the following sections, we first present a research project centrally focussed on the rapid detection of *phishing* email, designed to make it next to impossible for attackers to obtain financial resources or commit identity theft in this way. In the next section, we then show how another type of internet fraud, the violation of the rights of trademark owners by offering faked merchandise, can be semi-automatically detected with text mining methods. Thirdly, we report on two projects that are designed to prevent fraud in business processes in public administration, namely in the healthcare sector and in customs administrations. Finally, we focus on the issue of critical infrastructures, and describe our approach towards protecting them using a specific middleware architecture.

---

<sup>1</sup>Corresponding Author: Gerhard Paaß, Fraunhofer Institute for Intelligent Analysis and Information Systems, IAIS, Schloss Birlinghoven, 53757 Sankt Augustin, Germany, <http://www.iais.fraunhofer.de>; Email: [gerhard.paass@iais.fraunhofer.de](mailto:gerhard.paass@iais.fraunhofer.de)

## 1. Fighting Phishing: The Anti-Phish Project



Dear valued customer of TrustedBank,

We have received notice that you have recently attempted to withdraw the following amount from your checking account while in another country: \$135.25.

If this information is not correct, someone unknown may have access to your account. As a safety measure, please visit our website via the link below to verify your personal information:

<http://www.trustedbank.com/general/custverifyinfo.asp>

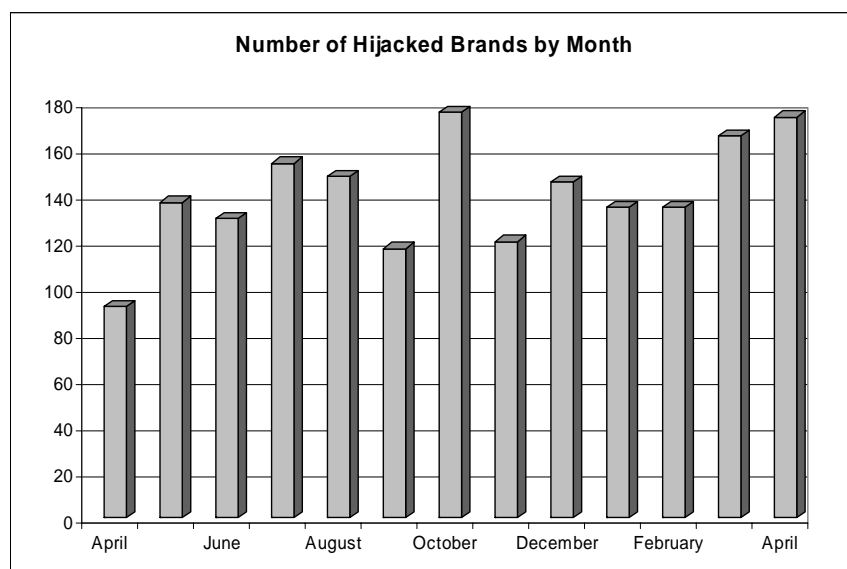
Once you have done this, our fraud department will work to resolve this discrepancy. We are happy you have chosen us to do business with.

Thank you,  
TrustedBank

Member FDIC © 2005 TrustedBank, Inc.

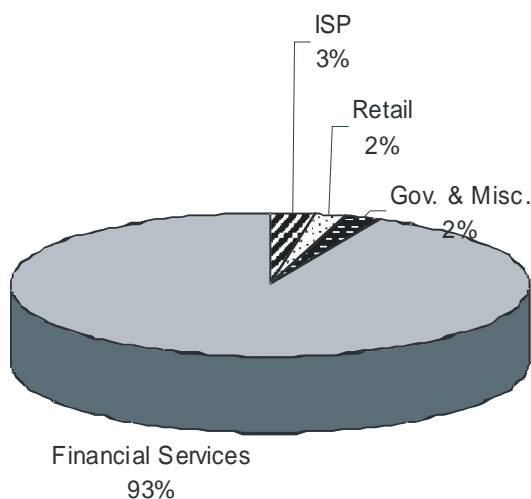
**Figure 1**, Example of a phishing email (© wikipedia <http://en.wikipedia.org/wiki/Image:PhishingTrustedBank.png>)

One of the most harmful forms of email spam is phishing. Criminals are trying to convince unsuspecting online customers of various institutions to surrender passwords, account numbers, social security numbers or other personal information. To this end they use spoofed messages which masquerade as coming from reputable online businesses such as financial institutions (see Figure 1 for an example). It includes a web link or a form to collect passwords and other sensitive information. Subsequently this information is used to withdraw money, enter a restricted computer system, etc.



**Figure 1**, Brands used in phishing emails April '06 – April'07 according to AntiPhishing Working Group (© Fraunhofer IAIS)

The AntiPhishing Working Group reports [APWG 07] that phishing has increased enormously over the last months and is a serious threat to global security and economy (see Figure 2). Most often banks and online payment services are targeted. Others phishers obtain sensitive data from U.S. taxpayers by fraudulent IRS-emails. Recently more non-financial brands were attacked including social networking (e.g. myspace.com), VoIP, and numerous large web-based email providers. (see Figure 3, source of graphics and information: APWG 07).



**Figure 3,** Phishing emails attacks by business sector of attack brand (© Fraunhofer IAIS)

There is an upward trend in the total number of different phishing emails sent over the whole internet. They grew from 11121 in April 06 to 55643 in April 07. There is a massive increase of phishing sites over the past year. In addition there is an increasing sophistication of phishing e-mails, e.g. by link manipulation, URL spelling using unicode letters, website address manipulation. Finally there is an evolution of phishing methods from shotgun-style email to phishing using embedded images and targeted attacks to specific organizations (“spear phishing”).



**Figure 4,** Anti-Phish project logo (© AntiPhish Consortium)

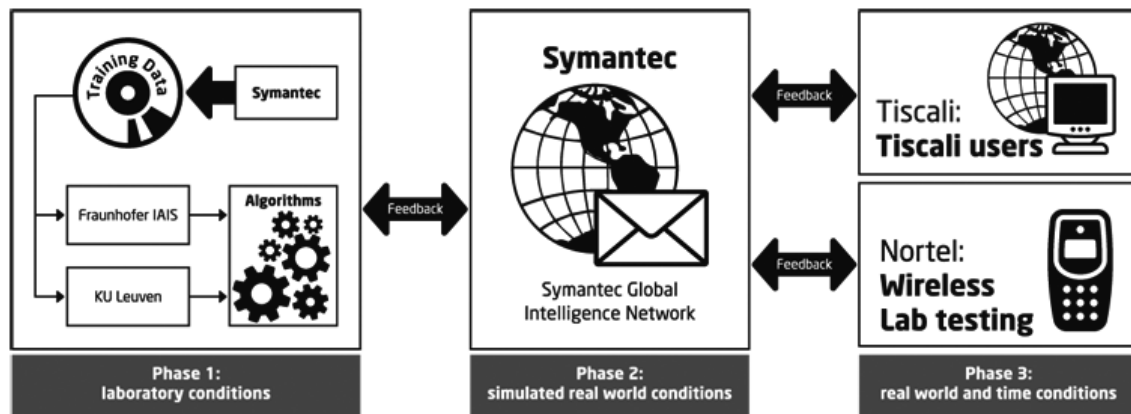
AntiPhish (see figure 4, <http://www.antiphishresearch.org/press.html>) is a specific targeted research project funded under Framework Program 6 by the European Union. It started in January 2006 and will run until December 2008. AntiPhish is an acronym for “Anticipatory Learning for Reliable Phishing Prevention”. The consortium consists of six partners. AntiPhish aims at developing spam and phishing filters with high accuracy for use both on traditional emails and mobile messaging services.

The scientific focus of the project is on trainable and adaptive filters that are not only able to identify variations of previous phishing messages, but are capable anticipating new forms of phishing attacks. Such technology does not exist yet, but could greatly improve all existing methods used in spam and phishing filters. So the project does not pursue new legal regulations or the change of email protocols, but is concentrating on technical recognition approaches.

Besides Fraunhofer IAIS the consortium comprises a reputable university (K.U. Leuven), a world leading company (Symantec), a major internet service provider (Tiscali) and a leading provider of mobile phone infrastructure (Nortel). The research will be driven by the hands-on expertise of our industrial participants, who have years of success fighting spam on a global scale. The AntiPhish project not only aims at developing the filter methodology in a test laboratory setting, but has the explicit goal to implement this technology in real

world settings at our partners sites, to be used to filter all email traffic online in real time, as well as content filtering at the edge of wireless networks.

Figure 5 shows the structure of the project [AntiPhish 06]. From the stream of emails appropriate features are extracted to train and estimate classifiers in a laboratory setting. Subsequently these filters are simulated in real-world conditions in the Symantec labs. Finally, they will be deployed into large-scale real world environments at an internet provider and a mobile phone provider.



The whole project will be driven by the continuous exchange and interplay between the research and industrial partners. The outcome will be a real-world product.

Figure 5, Anti-Phish project structure (© AntiPhish Consortium)

If successful all of this technology will be implemented in Symantec's Global Intelligence Network where people are working in three shifts, 24 hours a day, around the globe to help fight spam and phishing. A success against these threats will always require a combination of human labor and intelligent algorithms. The main approach of this project is to obtain training data from the email stream, to extract features, to estimate and update classifiers and, at the end, to deploy them at the internet service provider. A central challenge to this task is concept drift and the project will develop adaptive approaches to detect this concept drift and react on it.

## 2. Detecting Illegal Products: Text Mining to Identify Fake Auction Offers

According to a current study by the University of Mainz [Huber et al. 06] the number of faked products seized by German customs between 1998 and 2004 rose by about 1000 percent. The real number of fake products is estimated significantly higher. Whereas previously mainly watch, clothing and perfume industries had to face this problem, now the pharmaceutical, automobile and aeronautical industries are affected as well. For a perfume web auction in a specific period 84.4 percent of products in the sample survey could be identified as fakes, and only 7 percent as the genuine article. Besides the financial losses for the original manufacturer the poor quality of fake products can adversely affect the original brand.



**Figure 6,** Counterfeited Luxury Watch (© wikipedia <http://en.wikipedia.org/wiki/Image:Frolex.jpg>)

Often the fake products are openly offered as replicas, remakes, etc. and the customer knows that he will get a fake product, but nevertheless with the prestige of the genuine product. For a well-known manufacturer of luxury watches Fraunhofer IAIS developed a filter to detect such fakes in Ebay watch offer auctions. First a training set of genuine and fake watches offers were compiled. Then a classifier was trained to detect the different classes using the text of the offer and format features as inputs. On a training set these classifiers showed very good performance and needed very little time for actual filtering.

Therefore it may be used in the workflow of an internet auction platform. Whenever a new offer is entered the classifier checks the text and immediately bans a suspicious auction. The additional computational effort for the auction provider is minimal. Using the results of the project the highest German court, the Bundesgerichtshof, ruled that a brand product manufacturer can require an internet auction provider to set up a filter to identify fake offers and exclude them from the auctions [Bundesgerichtshof 07].

### **3. Preventing Fraud: The iWebCare and RACWeB Projects**

The assessment of fraud risks associated with business processes is a very interesting field of application for data mining. With respect to security and crime detection, those risks can, for example, be the risk of customs offences, or the risk of fraudulent activities in public administration. This has a significant economical impact; for example according to the Counter Fraud Service of the UK's National Health System, fraud accounts for almost 3% of public healthcare expenditures in the UK.

The projects iWebCare and RACWeB (Risk Assessment in Customs in the Western Balkans), in which Fraunhofer IAIS takes part, are aimed at implementing anti-fraud measures in public administration based on data mining techniques in the domains of healthcare and customs, respectively, in order to improve the efficiency and transparency of risk assessment and fraud detection and prevention.

The problem of fraud detection is characterized by the adversarial setting of a fraudster trying to hide from fraud detection in a large amount of data. This setting has important implications that distinguish fraud detection from standard data mining tasks:

**Skewed class distribution:** there are many different estimates of how many cases of fraud exist in different domains. What they all have in common is that fraud is always the exception to the rule, that is, the percentage of fraud cases is usually quite low. This poses a problem for many supervised learning algorithms, which usually work best when there are equally many positive and negative examples. Skewed class distributions can be addressed in data mining by forms of sampling [Scholz, 2005].

**Sparsely labeled data:** fraud is a complex organizational problem, and it is usually not trivial to determine whether a given case actually constitutes fraud or not, as many other legal, organizational, and practical issues have to be considered. As a result, one cannot expect that fraud experts can produce large labeled data sets that

would be necessary to train supervised methods. Instead, one has to take into account the fact that expert feedback is a scarce resource and one must make optimal use of the limited time that experts have by using methods like active learning and semi-supervised learning [Chapelle et al., 2006].

**False negatives:** a false negative is a case of fraud which the system has labeled as non-fraud. While a case the system labels as fraud will be investigated by a fraud operator, such that false positives will be corrected later in the fraud discovery process, false negatives will not be inspected any further and disappear in the large amount of cases which are seemingly correct. Hence, completely novel, unknown methods of committing fraud will not be detected by a supervised learning scheme and one must rely on unsupervised methods to detect new trends and developments in the data that may turn out to be fraud on further inspection. The general problem with this is that while it is possible to confirm a detected case of fraud with high confidence, it is much more difficult to decide whether a given case is not fraudulent or whether the fraud is simply hidden too well.

**Trivial rules:** a standard approach for fraud detection would be to collect an initial set of known fraud cases from fraud inspectors and then to use a supervised learner to generalize these fraud cases into rules to detect future cases of fraud. The problem with this approach is that fraud inspectors already use rules for finding fraud. These rules may either be given explicitly as a set of procedures to apply to new cases (an example of this would be to check that no two invoices have the same serial number), or implicitly by the way fraud inspectors work (an example being that when fraud inspectors come upon a novel case of fraud, they try to find similar cases, e.g., by checking all the invoices submitted by the same company). Consequently, the set of known fraud cases given to the system is not randomly selected, as would be assumed by standard statistical techniques, but are biased in the sense that only fraud cases that are detectable by the known fraud rules are included. When this happens, the best thing the data mining algorithm can do is to re-construct the old rules that have been used to construct the data set. Returning these rules to the fraud inspectors will be disappointing to them, because they will only get back what they already know.

**Concept drift:** as a direct result of the adversarial nature of fraud and anti-fraud measures, whenever there is a good way to detect certain cases of fraud, fraudsters will sooner or later adapt to the detection strategy and will try to find novel ways to cheat the system. This means that fraud is evolving over time and hence one cannot expect fraud rules to remain constant over time. This scenario is called concept drift in the machine learning literature and must be treated by careful validation of the learning results [Widmer, Kubat, 1996].

**Interpretability:** data mining is only one step in the process of fighting fraud. Human fraud inspectors and automatic data mining have to interact closely to have the maximum effect. In order to achieve this interactivity, it is important that the rules returned by the system are interpretable to the user, such that he can decide whether the new pattern that the algorithm has discovered actually describes fraud or not [Rüping, 2006].

From this discussion it follows that one has to distinguish between two different goals of data mining fraud detection.

**1) Detection of new fraud cases of a known fraud pattern.** When a new case of fraud is detected, the goal of fraud detection is not only to stop and prosecute this particular instance of fraud, e.g., by dismissing an employee who was involved in procurement fraud, but also to prevent similar cases of fraud from occurring, e.g., by finding indicators that facilitate the identification of such cases earlier and with higher confidence. Supervised learning can be used to find these indicators by generalizing the single cases into high-quality rules, and prevent the same type of fraud from happening again.

**2) Detection of new fraud patterns.** It is safe to assume that new types of fraud are being developed all the time. Hence, fraud detection cannot only rely on tracing known types of fraud, but must also incorporate methods to find new fraud patterns. In order to do this, one cannot rely on known fraud labels, but must identify unusual patterns based on other properties of the data. Once a statistical significant deviation is found, the pattern can then be reported back to a fraud officer for investigation. An example might be that a certain type of doctor spends much more money per patient than the rest. While it can be statistically confirmed that such a deviation in the budget is indeed significant and not random, one can usually not decide from the given data whether there is a valid reason for the higher spending (e.g., the doctor treating a special group of high-risk patients that require more expensive treatment) or whether this is a sign of fraud. This makes unsupervised fraud detection more challenging, because it needs to combine high statistical significance of the found patterns with interpretability, such that the experts can understand and judge the validity of the patterns.

#### **4. Fighting Threats Against Critical Infrastructures: The IRRIS Project**

The EU Integrated Project "Integrated Risk Reduction of Information-based Infrastructure Systems" (IRRIIS) is carried out under the motto: Enhance substantially the dependability of Large Complex Critical Infrastructures (LCCIs) by introducing appropriate Middleware Improved Technology (MIT) components. IRRIS increases dependability, survivability and resilience of EU critical information infrastructures based on Information and Communication Technology (ICT) and has the objectives to:

- determine a sound set of public and private sector requirements based upon scenario and related data analysis;
- design, develop, integrate and test MIT components suitable for preventing and limiting cascading effects and supporting automated recovery and service continuity in critical situations;
- develop, integrate, and validate novel and advanced modeling and simulation tools integrated into a synthetic environment (SimCIP) for experiments and exercises;
- validate the functions of the MIT components using the SimCIP environment and the results of the scenario and data analysis;
- disseminate novel and innovative concepts, results, and products to other ICT-based critical sectors.

IRRIIS addresses the challenges of Complex Information Infrastructure Protection (CIIP) by a "diagnosis - therapy strategy" and "therapy implementation and validation approach" starting with the electrical power infrastructure and its supporting telecommunication infrastructure. After thoroughly analysing these infrastructures and their interdependencies, the synthetic simulation environment (SimCIP) is build. MIT components are developed, tested and validated inside SimCIP to demonstrate their capabilities before dissemination to potential stakeholders. The approach subsequently includes additional critical infrastructures.

The interdisciplinary research is performed by a European consortium of fifteen partners, ranging from academia to key stakeholders from the fields of energy supply and telecommunication. The project is partly financed by the European Union's Sixth Framework Programme for a term of three years.

#### *4.1. Large Complex Critical Infrastructure Analysis and Requirements*

Up until now there is a lack of deep understanding of Large Complex Critical Infrastructures' (LCCIs) dependability and interdependency particularly with regard to the use of Information and Communication Technology (ICT). Although some models and tools dealing with these issues exist, LCCI complexity cannot yet be tackled properly. Basic research is necessary to understand the phenomena of interdependency, dynamic behaviour and cascading effects in order to support the development of solutions for protecting and managing existing LCCIs in case of incidents. IRRIIS will perform in-depth research regarding the topological structure of LCCIs and the interdependencies between different LCCIs. Appropriate analytical approaches will be applied such as simulation models or analytical models suitable to investigate interdependency, network dynamics and cascading effects.

Starting from a thorough analysis of LCCIs, incorporating the stakeholder's views regarding ICT tools and models, a sound set of public and private sector requirements can be determined. These requirements will be the base for the development of the SimCIP simulation environment and the Middleware Improved Technology (MIT) components.

In order to enhance the understanding of LCCIs and to gain a sound foundation for the development of the SimCIP simulation environment and the MIT components IRRIIS will:

- Survey LCCI stakeholder's requirements on technology and tools needed for understanding and mitigating cascading effects
- Survey and analyse existing tools and models
- Analyse current research gaps to identify relevant research and development efforts
- Provide detailed scenario and risk analysis
- Perform in-depth topological analysis of LCCIs
- Analyse the interdependency between different LCCIs
- Analyse the upcoming Next Generation Networks (NGN), i.e., networks based on IP connectivity or wireless connections with mainly software based services


This work will not only help ensure that the SimCIP environment and the MIT components meet the stakeholder's needs but also contribute to the ongoing world-wide research efforts concerning LCCIs.

#### *4.2. Middleware Improved Technology*

Starting with the knowledge gained from the LCCI analysis and the survey of stakeholder's requirements and existing tools, Middleware Improved Technology (MIT) components are developed. These MIT components facilitate the communication between different LCCIs and allow the identification and evaluation of incidents and malicious attacks and responding accordingly.

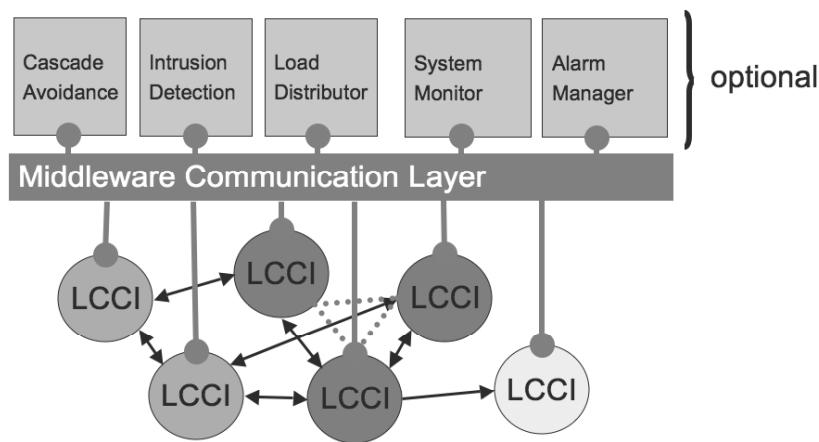
Currently, one big problem for the dependability, security and resilience of LCCIs is the high interdependence between different LCCIs within the same sector and also between different sectors. The consequence is that problems within one LCCI can lead to severe problems in dependent LCCIs. The resulting cascading effects are not limited to one kind of infrastructure and do not stop at national borders. To make things

worse there is often a lack of appropriate communication structures between the dependent LCCIs (see Figure 7). This results in a lack of awareness of problems occurring in other infrastructures and therefore mitigating actions cannot be performed in time.



**Figure 7:** Interdependent LCCIs of the same and different sectors. The arrows indicate interdependencies and the lines communication links using different standards. (© IRRIS Consortium)

To facilitate the communication between different infrastructures, IRRIS will develop appropriate middleware communication components. All communication between different LCCIs should run via this middleware. The advantage is that each LCCI only needs one communication link to the middleware and does not have to interface with several other LCCIs (see Figure 8).

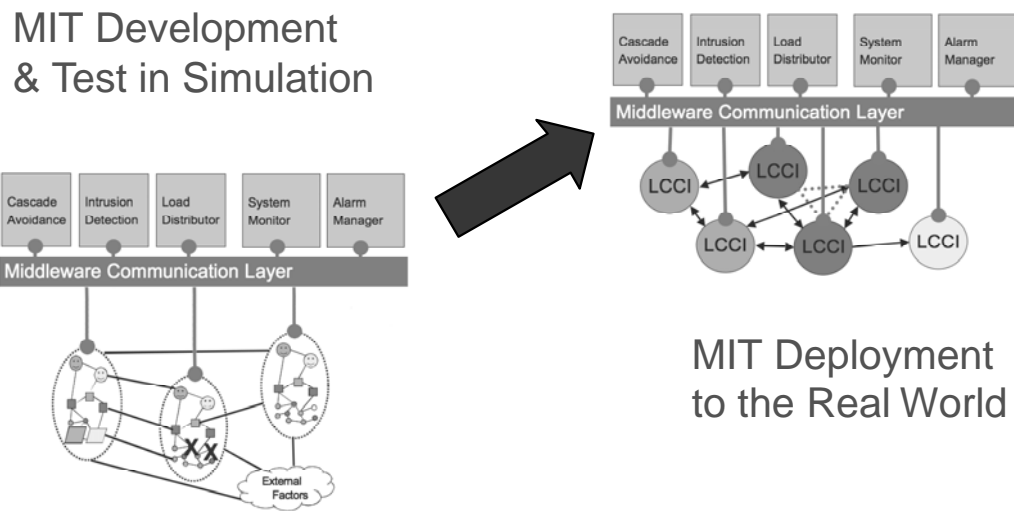


**Figure 8:** Interdependent LCCIs with Middleware Communication Layer and Middleware Improved Technology components. (© IRRIS Consortium)

The middleware will also be used by the optional MIT add-on components which have some kind of build-in “intelligence”. These add-on components will monitor data flowing within and between the infrastructures and raise an alarm in case of intrusions or emergencies and then take measures to avoid cascading effects. They will be able to detect anomalies, filter alarms according to their relevance and support recovery actions and will thus contribute to the security and dependability of LCCIs. MIT components will interface with existing systems and will not require major replacement of existing hardware or software. The flexibility of the middleware allows the easy integration of new LCCIs or the exchange of new kinds of information.

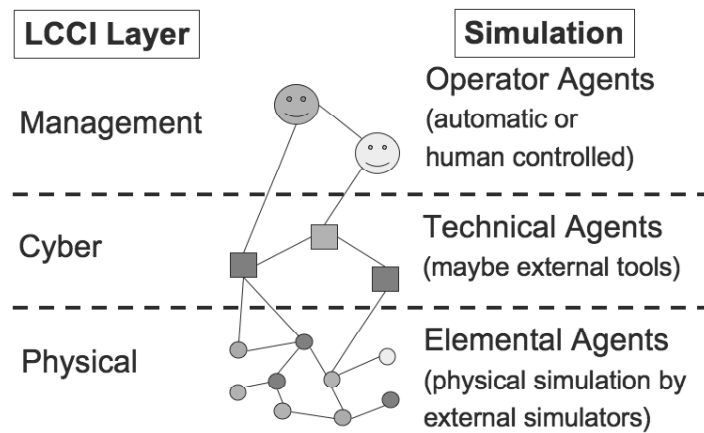
#### 4.3. SimCIP Simulator for Critical Infrastructure Protection Application

The purpose of the SimCIP simulation environment is twofold: On the one hand, side simulation can be used to improve the understanding of interdependent LCCIs. On the other hand, the MIT components will be tested and validated in experiments using SimCIP. Furthermore, their applicability and usefulness will be demonstrated to stakeholders within the SimCIP environment before deployment to the “real world” systems (see Figure 9).



**Figure 9:** The role of the SimCIP simulation environment in the development of the MIT components.  
 (© IRRIS Consortium)

Building the SimCIP environment is a big challenge because the simulation will not only have to include physical simulations but also has to simulate the cyber layer and the management layer of a LCCI as well (see Figure 10). For this purpose SimCIP will use the principle of agent-based simulation. Each object will be modeled as an agent with clear interfaces to its environment and other agents.



**Figure 10:** The different layers of a single LCCI and the according SimCIP simulation objects.  
 (© IRRIS Consortium)

The SimCIP environment will include and interface with existing tools to keep the simulation meaningful with respect to existing technologies and to allow the use of the results gained in current systems. This also means that the SimCIP environment does not have to start from scratch but can rely on already existing and proven technology. To decide which tools and models should be included in SimCIP, an in-depth survey of existing tools and models will be performed.

However, the main strength of SimCIP will be the simulation of interdependencies between different LCCIs. To that end, one will model some objects of the individual LCCIs on more abstract levels. This will ensure the high scalability and flexibility of the SimCIP environment. SimCIP should be as generic as possible to allow its application to various kinds of LCCIs and its adaptation to the specific needs of individual stakeholders.

Knowledge Elicitation and Research will lead to a “diagnosis” of the current and the future status of interdependent LCCIs. The “therapy” will be implemented through the MIT components which can be tested and validated in the SimCIP environment. The main contributions of IRRIS are an enhanced understanding of

LCCIs, the SimCIP simulation environment and the MIT components. To disseminate the results broadly to stakeholders, technology and service providers and the research community, these interest groups will be included within the IRRIS project right from the start. IRRIS also relies on international cooperation and is open to joint efforts of all kinds to achieve its goals. To foster international cooperation IRRIS will establish an international conference and will define scenarios and benchmarks to allow the comparison of different approaches.

## 5. Conclusion

In this paper, we have given an overview of several projects currently carried out at our institute which are aimed at fighting internet related crime. As these projects show, to fully address the issue of internet related security threats, especially when looking towards politically and ideologically motivated attacks, requires consideration not only of physical violence, but also of the criminal activities that provide the funds for such activities. As a primary example, phishing is used on the internet in this way, and we have described how phishing activities can be suppressed using advanced text mining solutions. In related projects, we show how data mining and text mining techniques can be used to prevent fraud with fake auction goods, and in the medical administration system. Finally, we have described how specific software architectures and middlewares can be used to protect critical infrastructures consisting of technical, economic and social systems that are interrelated.

All the above projects form part of what we call “Preventive Security” which is one of the major areas of activity for our institute. In addition to the projects described here, we provide services to military and non-military organizations in the area of organization planning, simulation and process intelligence, as well as research and applied work on physical hardware and robots for security and inspection purposes.

## Acknowledgements

This work was funded in part by the EU FP6 projects AntiPhish (contract 027600), IRRIS (contract 027568), iWebCare (contract 28055), and RACWeB (contract 045101).

## References

- AntiPhish (2006): Website of the project. <http://www.antiphishresearch.org/home.html>.
- APWG (2007): Phishing Activity Trends. Report for the Month of April, 2007. [http://www.antiphishing.org/reports/apwg\\_report\\_april\\_2007.pdf](http://www.antiphishing.org/reports/apwg_report_april_2007.pdf).
- Bundesgerichtshof (2007): Bundesgerichtshof bestätigt Rechtsprechung zur Haftung von eBay bei Markenverletzungen. Press release 45/2007, April 19th 2007.
- Chapelle, O., and Schölkopf, B., and Zien, A. (2006): Semi-Supervised Learning, MIT Press, Cambridge, MA.
- Huber, F., Matthes, I., Vollhardt, K., Ulbrich, D. (2006): Marken- und Produktpiraterie aufdecken und bekämpfen – am Beispiel von Internetauktionen eines Markenparfums Arbeitspapiere Management P6 Center of Market-Oriented Product and Production Management Mainz ISBN: 3-938879-13-0.
- Landler, M., Markoff, J. (2007): Digital Fears Emerge After Data Siege in Estonia. Online edition of New York Times, May 24<sup>th</sup>, 2007. <http://www.nytimes.com/2007/05/29/technology/29estonia.html?pagewanted=1&ei=5070&en=5858f0fda0af7087&ex=1184644800>.
- Rhoads, C. (2007): Cyber Attack Vexes Estonia, Poses Debate. The Wall Street Journal Online. May 18, 2007. Page A6.
- Rüping, Stefan (2006): [Learning Interpretable Models](http://hdl.handle.net/2003/23008), Ph.D. Thesis, Universität Dortmund, URL: <http://hdl.handle.net/2003/23008>.
- Scholz, Martin (2005): [Sampling-Based Sequential Subgroup Mining](#). In Grossman, R. L. and Bayardo, R. and Bennett, K. and Vaidya, J. (editors), Proceedings of the 11th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (KDD '05), Seiten 265--274, Chicago, Illinois, USA, ACM Press.
- Widmer, G. and Kubat, M. (1996): Learning in the presence of concept drift and hidden contexts, Machine Learning 23, pp. 69-101.

**Copyright of images:**

- Figure 1: Public domain image from <http://en.wikipedia.org/wiki/Image:PhishingTrustedBank.png>
- Figure 2 & 3: Own Graphs
- Figure 4: Logo of our own project AntiPhish, copyright AntiPhish consortium
- Figure 5: Graph of our own project AntiPhish, copyright AntiPhish consortium
- Figure 6: Public domain image from <http://en.wikipedia.org/wiki/Image:Frolex.jpg>
- Figures 7-10: Graphs of our own project IRRIS, copyright IRRIS consortium