Proposals for Combating Cyber Terrorism through Preventive Active Security

by Radu Sion, Mikhail Atallah, Sunil Prabhakar

Center for Education and Research in
Information Assurance and Security,
Purdue University, West Lafayette, IN 47907

Proposals for Combating Cyber Terrorism through Preventive Active Security (CERIAS TR 2002-32) *

Radu Sion, Mikhail Atallah, Sunil Prabhakar
Center for Education and Research in Information Assurance,
Computer Sciences, Purdue University
West Lafayette, IN, 47907, USA
Phone: (765) 494-6008
Fax: (765) 463-7310
http://www.cs.purdue.edu/homes/sion
[sion, mja, sunil]@cs.purdue.edu

Abstract

Unfortunate recent events clarified the absolute requirement for a unified, concerted, scientifically proven strategy for combating forms of actual or potential cyber-terrorism. In this paper we present related solutions based on some of our ongoing and proposed future research in the broader areas of data and system security. More specifically we focus on preventive techniques for content and system security. In the framework of content security we discuss document tamper-proofing, watermarking and generic information hiding detection, essential tools required in the combat against attacks in the current distributed, heterogeneous, networked world. System security issues address new intrusion detection mechanisms using biometrics as well as new concepts such as "data access patterns", in the framework of structured content and "network breath" in the case of secure computer networks. Finally we present our research in the area of secure multi-party cooperation, an essential component in any inter-party contingency interaction scenario where trust issues might prevent complete cooperation. In the end we introduce some of the main conclusions and propose immediate-future research and focus points.

Keywords: Cyber Terrorism, Active Security, Information Hiding Detection, Intrusion Detection, Biometrics, Secure Cooperation

1 Introduction

In this paper we present solutions to some essential issues and open problems in the framework of combating forms of actual or potential cyber-terrorism. Our ideas are based on some of our ongoing and proposed future research in the broader areas of data and system security. More specifically we focus on *preventive* techniques for content and system security.

Content Security comprises a series of areas that evolve mainly around the concept of digital content. This is relevant to the case as probably most of the communication between parties involved in cyber-attacks and cyber-protection occur through exchange of digital objects, data.

*Portions of this work were supported by Grants EIA-9903545, IIS-9985019 and IIS-9972883 from the National Science Foundation, Contract N00014-02-1-0364 from the Office of Naval Research, and by sponsors of the Center for Education and Research in Information Assurance and Security.

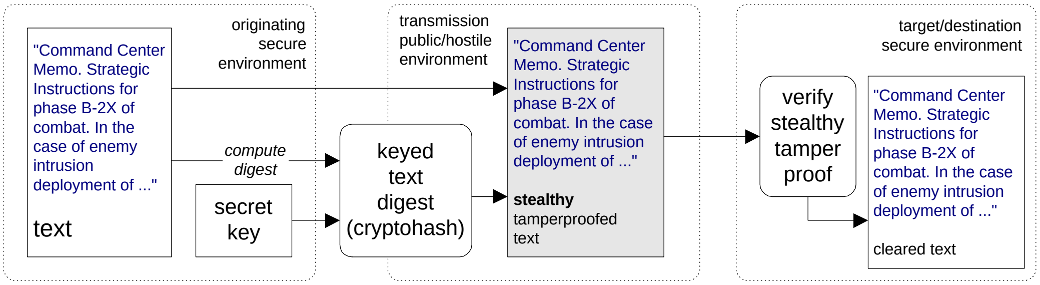
originating
secure
environment
"Command Center Memo. Strategic Instructions for phase B-2X of combat. In the case of enemy intrusion deployment of ..."
text
compute digest
secret key
keyed text digest (cryptohash)
transmission public/hostile environment
"Command Center Memo. Strategic Instructions for phase B-2X of combat. In the case of enemy intrusion deployment of ..."
stealthy tamperproofed text
target/destination secure environment
verify stealthy tamper proof
"Command Center Memo. Strategic Instructions for phase B-2X of combat. In the case of enemy intrusion deployment of ..."
cleared text

Figure 1: XTX provides a solution for stealthy tamper-proofing secure documents that are to be transmitted across hostile environment boundaries.

In this framework we discuss document tamper-proofing, watermarking and generic information hiding detection, essential tools required in the combat against attacks in the current distributed, heterogeneous, networked world.

System security deals with security aspects related to digital content processing and access to computing resources and secure data. More specifically this is extremely relevant to the case of cyber-terrorism and associated attacks in any scenario in which unauthorized access to data and/or resources can lead to potential active/concerted attacks.

Here we address new intrusion detection mechanisms using biometrics as well as new concepts such as ‘data access patterns”, in the framework of structured content and ‘network breath” in the case of secure computer networks.

Finally we present our research in the area of secure multi-party cooperation. Broadly speaking, a secure multi-party computation problem deals with computing any function on any input, in a distributed network where each party holds one of the inputs, ensuring independence of the inputs, correctness of the computation, and that no more information is revealed to a participant in the computation than can be computed from that participant's input and output. This problem is thus an essential component in any inter-party contingency interaction scenario where trust issues might prevent complete cooperation.

The paper is structured as follows. Section 2 discusses content security addressing issues such as document tamper-proofing, watermarking and generic information hiding detection. In Section 3 detection mechanisms for system security are presented and new associated concepts are introduced. Section 4 addresses research in the area of secure multi-party computing. Finally, Section 5 presents some main conclusions and positions our ideas as proposed future research.

2 Content Security

Content Security comprises a series of areas that evolve mainly around the concept of digital content. This is relevant to the case as probably most of the communication between parties involved in cyber-attacks and cyber-protection occur through exchange of digital objects, data. In this framework we are discussing several security issues and techniques, including stealthy tamper-proofing of documents, watermarking and information hiding detection.

2.1 Digital Tamper-proofing

Often times a mechanism of stealthy authenticity assessment for digitally transmitted documents is required, for example in scenarios involving a non-secure third party or transmission medium as depicted in Figure 1. This can become vital in battlefield or crisis situations where several decision, intermediation and target factors cooperate toward a common crisis solution goal.

A simple to use, yet powerful algorithm is to be desired, allowing hiding (and corresponding verification) of document authenticity information preferably self-contained within the data itself.

We are currently working on a software solution for document tamper-proofing, XTX. Without entering into too much detail, the main idea behind XTX is the use of keyed special one-way cryptographic hashes (digests) of parts of the original content, digests "embedded" into the transmitted version, functioning effectively as "witnesses" to its integrity,

Additionally, for each document, the actual digest method as well as other algorithm parameters are selected by a secret result from a sender-receiver (initially) synchronized random key generator.

Any form of tampering with the original document, be it at the semantic/syntax level or even at the text indentation level triggers the XTX detection mechanism. Once triggered, XTX attempts to actually also discover what changes have been made to the document from its original version.

The system has reached the proof-of-concept implementation stage and currently allows stealthy (virtually undetectable) tamper-proofing of text documents. Extensions to other types of media are envisioned.

2.2 Watermarking. Information Hiding Detection

Detection of information hiding deals with the problem of determining whether data (e.g., media, text, etc) contains a secret message hidden in it or not. The importance of this area stems from the belief that evil-doers will increasingly use information-hiding technologies to communicate secretly without even appearing to be hiding their communication. The reason they may prefer this to using encrypted message exchanges is that the latter immediately reveal that this particular pair of communicating parties are hiding something, which would expose them to either technological or judicial and legal counter-measures [1].

The goal of detecting whether an exchange contains a hidden message is, of course, only a first step toward the next stage (of determining what the secret message is), but it is immensely challenging as a separate research topic in its own right.

For example, using an online auction service [2] (e.g. ebay) to publish hundreds of watermarked entries (e.g. multi-typed documents advertising the sale of fake non-existent items), a (cyber) terrorist element could spread huge amounts of information to followers in a very hard to detect manner [3].

One approach to this problem is statistically based: The injection into a "normal" media of a secret message will typically disturb it enough to make such an approach viable - the main challenge is of course the traditional one of feature selection (which characteristics of the media to use for the detection mechanism). Another envisioned approach would make use of Information Theory concepts, particularly the notion of Entropy.

In the framework of the Internet and the Web, having a formal model of multi-typed documents (e.g. HTML/XML) would constitute a first start toward a design for an automated network crawler that could perform a search for public watermarked multi-typed documents as well as other tasks such as copyright violation detection.

We propose initial steps in defining a comprehensive theoretic domain model of information hiding and watermarking with immediate applicability in the generic framework of information hiding detection for multi-type/media documents. We also propose the analysis of possible workings of an information

[1] E.g. the use of supercomputers and sophisticated algorithms for attacking their encryption, search warrants and forensic analyses of their computer hardware

[2] Ebay.com features a huge amount of new auction items every month, often reaching hundreds of millions.

[3] Example brought up by a participant at a CERIAS Security Seminar.

hiding detection process and the relation to the estimated domain model. Envisioned attacks should be outlined from a mark-destructive/recovery perspective.

Multi-type/media documents are characterized usually by value lying both in the structure ("Graph") *and* in the non-structured content ("Nodes"). Examples include XML documents, complex Web Content, Software, Natural Language, relational DBMS data, VRML and similar environmental representations, structured financial and B2B interaction data, work-flow and planning descriptions etc.

In most traditional/media watermarking techniques, the noise-band is usually the main recipient for any watermarking techniques. The amount of noise in well defined structured data tends asymptotically to zero increasing the complexity of our challenge.

In this case the available noise-band width is very low, making it necessary to discover new data properties that allow for a resilient, un-noticeable entropy increase, maintaining all other properties of interest in place, within usable boundaries.

One idea considers the fact that "nodes" in semi-structures/documents are value-carrying, and a watermarking algorithm could make use of their encoding capacity, by deploying traditional watermarking techniques. Apparently, bandwidth is available from capacities associated to properties of both the structural and the content domains. Further insight shows that many limitations apply and the task is not trivial. In many applications *stealthiness* is paramount even compared to encoding capacity.

Discovering appropriate bandwidth channels (watermarking capacity) is one of the challenges. Their use in a most efficient and resilient way is another. The above-described nature of documents determines the existence of multiple bandwidth sources. In general we can distinguish between different types of channels as shown below. Further research is required to differentiate among them and focus on the types of channels that offer the least distortion and maximum bandwidth with respect to a desired usability domain.

Semantic channels derive from the higher level semantics of the documents. *Content semantics* deals with the actual document content whereas *structural semantics* capture meaning in the structural (e.g. relation semantics) definition of the document (e.g. in XML Schemas).

Aggregation channels are linked to information content in the actual aggregation of various content within the document. We plan on developing a generic information hiding model of aggregation and define main bandwidth sources. We immediately envision several sub-channel carriers: *topology*, *relation attributes*, *node correlation.*

Raw content channels are defined by the actual types of content that are aggregated into a document. We distinguish between *media* and *non-media* content. In the media realm we identify *image*, *audio* and *video* types. Non-media content, one of our main area of interest, includes *various types of text* and *functionality/behavior.*

Note: We mentioned functionality here to basically model approaches where a certain document distributed on an interactive medium (e.g. e-books) has an extra available information channel, namely its behavior, allowing Easter-Egg approaches where the watermark itself is an intrinsic property of the behavioral aspects of the document.

Text content can be of many types (e.g. natural language, functional descriptions, structured meta-information, software code) and further research needs to identify bandwidth channels for each of them as well as a generalized theoretic framework that can be applied to multiple types of text.

Having identified a set of generic channels as above is helpful but still does not offer the power of a higher level theoretic framework. A higher level theoretic model for multi-type/media documents needs to be developed, allowing identification of available channels given existing content types and aggregation constraints.

We envision one generic detection technique that uses certain *metrics of normality* in identifying violated domain constraints. Thus our domain model needs to be augmented with elements allowing the specification of domain constraints and "normality".

Note: On the other hand, in the area of actual information hiding for covert communication, further research needs to actually develop a method which "covers" the hidden data by identifying normality metrics in each of the considered hiding channels and then uses those metrics to model the hidden mark (e.g. make the channel look like "white" noise to an attacker).

More details on our research in the areas of information hiding and watermarking can be found at **http://www.cerias.purdue.edu** and **http://www.cs.purdue.edu/homes/sion/wm**.

3 Systems Security. Intrusion Detection.

System security deals with security aspects related to digital content processing and access to computing resources and secure data. More specifically this is extremely relevant to the case of cyber-terrorism and associated attacks in any scenario in which unauthorized access to data and/or resources can lead to potential active/concerted attacks.

In this section we discuss ongoing research as well as proposed ideas augmenting traditional authentication techniques with new intrusion detection mechanisms based on concepts such as biometrics and learning techniques for system normality (see Figure 2 (a)).

3.1 Keystroke Timing

Biometrics is based on the implicit assumption of the existence of characteristic patterns, features that uniquely identify a certain individual/user, differentiating him/her from other individuals/users. One of the first biometrics ever used are human fingerprints, particularly in forensics and criminology.

In the broad area of computing, various other human features have been claimed to qualify as unique identifiers for system users. We experimented with one of these, namely keystroke timing patterns.

xKey is a software solution that uses keystroke timing patterns in modeling a personal profile for the purpose of aiding authentication. It relies on several novel ideas, including the use of neural networks, dictionary training and normal profile matching.

xKey is in the stage of a proof-of-concept implementation software package that enables the storage, retrieval and matching of user keystroke timing information biometric profiles (see Figure 2 (b)).

The system models keystroke timing behavior by constructing so-called "profiles", basically recipients of timing information for user-entered tokens (words). Profiles are composed of timing information of several types including inter-letter time intervals and keystroke durations.

Once profiling data is gathered in an initial data aquisition step, a training step follows in which main characteristics are extracted from the raw profiles resulting in "refined" versions of the initial profile data.

This refinement process is built around two alternative solutions. One alternative uses a simpler model of confidence interval matching to allow the construction of a certain profile distance metric. The other solution does not rely on any predefined model but rather uses neural network training techniques in training a neural network to the data of a certain profile, the negative examples being derived from synthesized data and/or different observed/previously-seen users' profiles.

At authentication time, after optionally performing an initial username/password step, the user could be asked to type in a number of words in a particular order, determined by the algorithm (given the training data seen so far from all inputs to the system) to be of highest relevance in differentiating this particular user from other users and/or an intruder.

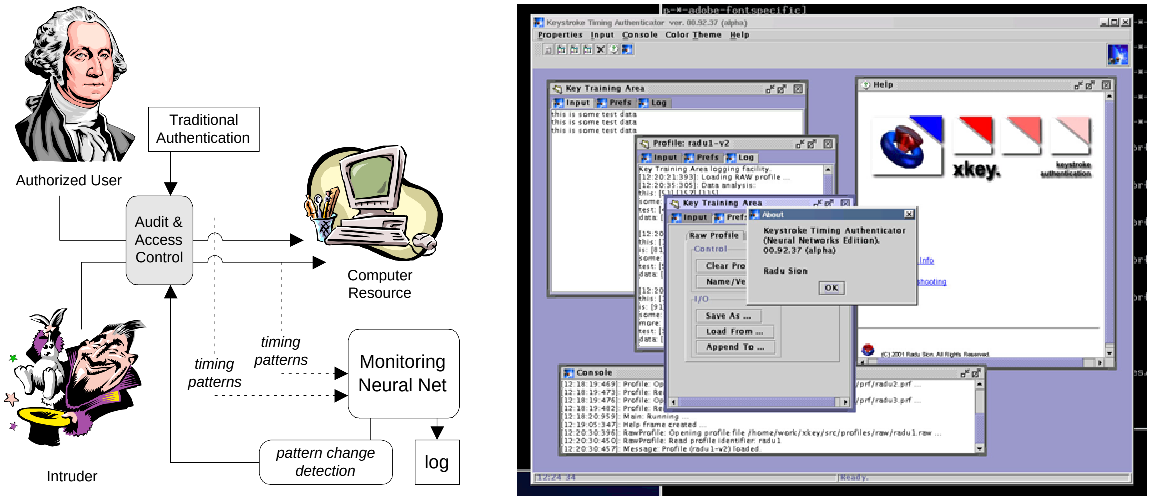

```
this is some test data
this is some test data
this is some test data
```

```
Key Training Area logging facility.
[12:20:21:393]: Loading RAW profile ...
[12:20:35:305]: Data analysis:
```

About

Keystroke Timing Authenticator (Neural Networks Edition). 00.92.37 (alpha)

Radu Sion

(C) 2001 Radu Sion. All Rights Reserved.

```
[12:18:19:469]: Profile: Op                /prf/radu2.prf ...
[12:18:19:473]: Profile: Re
[12:18:19:476]: Profile: Op                /prf/radu3.prf ...
[12:18:19:482]: Profile: Re
[12:18:20:959]: Main: Running ...
[12:19:05:347]: Help frame created ...
[12:20:30:396]: RawProfile: Opening profile file /home/work/xkey/src/profiles/raw/radu1.raw ...
[12:20:30:450]: RawProfile: Read profile identifier: radu1
[12:20:30:457]: Message: Profile (radu1-v2) loaded.
```

Figure 2: (a) Intrusion Detection by Training for Normality using Neural Networks. (b) A sample screenshot of xKey in operation. It allows for user keystroke timing profile training, refining and matching.

Given the nature of the implemented algorithm[4], a large number of individuals need to be assessed/tested before statements about the accuracy of the method can be provided.

Despite the lack of resources currently allocated to this topic, we performed testing on a group of 15-20 people with excellent recognition accuracy results (in some cases close to 99%) and a surprising false positive rate.

Note: We envision also other biometric "domains" that could be researched and integrated in to xKey, in the computer access framework, including *mouse motion behavior* and *desktop item/icon access patterns*.

More research needs to be performed on the underlying algorithm, on actually assessing method goodness per-se as well as in actually implementing it inside an existing operating system as a run-time module allowing for continuous authentication. A preview version of **xKey** can be found online at **http://www.cs.purdue.edu/homes/sion/projects/xkey**.

3.2 Hyper-Data: System Access Behavioral Patterns

Remote access to distributed information proves to be one of "the" killer applications for computer networks. More and more content in current inter and intra nets is available as hyper-linked data, a form easing its distribution and semantic organization.

In the framework of the Internet's Web-Portals and Pay-Sites, mechanisms for login based on username and password enable the dynamic customization as well as partial protection of the content. In other applications (e.g. commercial intra-nets) various similar schemes of authentication are deployed.

Stolen passwords are an easy avenue to identity theft, in both public and commercial data networks. Once a perpetrator enters the system, assuming a user's identity, the task of actually detecting this intrusion becomes non-trivial and is in most cases ignored completely.

We would argue that after the initial authentication step, intrusion detection mechanisms are required to continue a virtual run-time user authentication, and detect identity theft and password misuse.

[4] As no synthetic data production method is available for training.

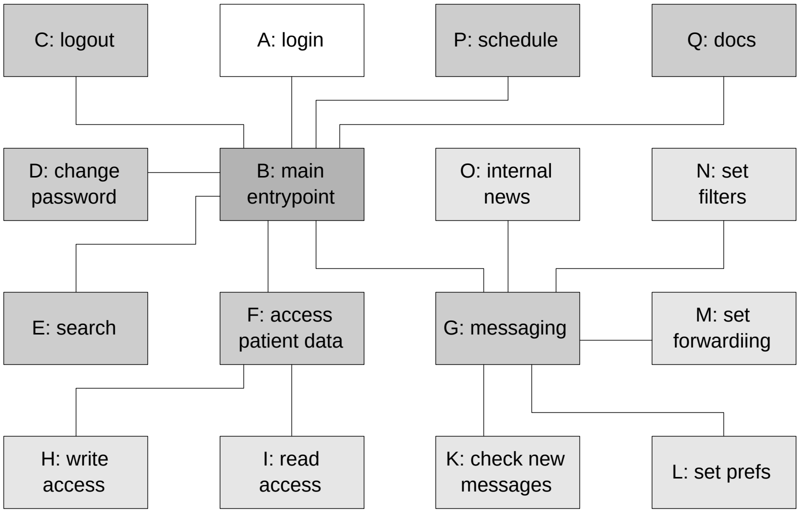

Figure 3: System Access Patterns for semi-structured content can be trained for and used as intrusion detection tools. Figure depicts a possible hospital intranet portal. A likely authorized access pattern would present a trace of A,B,G,O,P,F,I. A suspicious pattern could show up as A,B,D,F,H,G,M,L,N.

We propose research on designing a pervasive intrusion detection method based on analysing trusted access patterns to hyper-linked data, aiming at detecting intruders and raising a red flag at the content provider end.

Such a method is to be used in conjunction with current username- password protection schemes. It provides an additional level of security, allowing detection of identity theft or misuse. It could be deployed either as a mandatory shut-off authorization step or as a simple preventive warning mechanism.

We can relate our efforts to research in the framework of biometrics as a method for authentication through pre-training and run-time metric matching. In a remotely similar fashion our algorithms are designed to analyse and train for trusted access patterns, the "browse-metrics" of the data consumer.

We envision a broad applicability of our research. An implementation as a web-server module could be used in additionally enforcing current weak pay-site authentication mechanisms. Deployment inside a company or government agency intra-net could detect identity theft and immediately revoke credentials for access to sensitive internal information. An additional level of privacy control can be guaranteed by deployment in the framework of a hospital medical records access software.

There are currently no resources specifically allocated for this project. Nevertheless we started an initial problem assessment and made good progress by coming up with a preliminary algorithm draft for intrusion detection using access patterns for this type of content.

Without entering into much detail, the main idea behind our algorithm is the use of graph-theoretic concepts and primitives in training for normality and then matching authorized users and detecting intruders (see Figure 3).

Further research should focus on an actual proof-of-concept implementation in the framework of an existing content providing software server (e.g. Web Server) and on assessing its effectiveness through training and real-life testing. Different types of hyper-linked data and associated theoretical and deployment issues should be analysed.

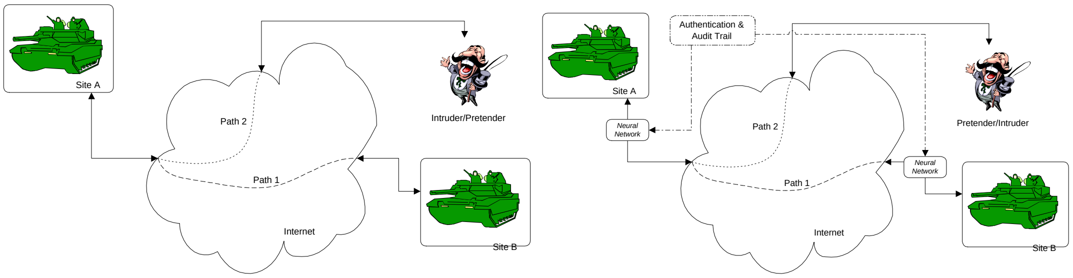

Figure 4: Network Breath. (a) An intruder enters the system by having gained password access ("pretender"). (b) An intrusion detection mechanism is able to detect difference in network packets timing patterns ("network breath") and signal a potential breach by a pretender.

3.3 Network Breath Intrusion Detection

Network Breath is a novel concept based on the claim that *medium to long-term* inter-site network timing properties (e.g. delay, jitter) feature a certain level of stability and are not influenced by short term bursts of high bandwidth traffic. These properties are seen to be network-link specific, hence the "network breath" name.

In other words network breath is a concept similar to a "system biometric" for inter-site network links. Thus, given two established secure sites (see Figure 4 (a)) the network breath of the network link between them is claimed to be a medium to long-term timing invariant.

The main idea then is to use (real or fake/generated) network breath to identify potential "pretenders" (see Figure 4 (b)), system intruders connecting from the outside to one of the secure sites and pretending to be the other site. This, in addition to anti-spoof mechanisms, can virtually eliminate this type of malicious attacks.

We again envision the use of a training (e.g. using neural networks or similar technology) device ("network breath meter") with memory, non-accessible through any external protocol (black box) "planted" on the incoming/outgoing connections of each site. The "network breath meter" undergoes periodic training phases in which it observes normal traffic patterns between the two sites, followed by detection phases in which it simply detects anomalous timing patterns on the given logical connection. If such patterns are observed consistently, a certain warning factor increases and ultimately triggers a log entry as well as a remote warning-mechanism.

The security of this scheme relies on the fact that the network breath meter is entirely non-accessible and logs every warning securely, internally, and remotely through a direct connection to a logging device. In extreme cases the box could be programmed to simply cut the traffic automatically permitting flow reactivation only upon a restart through a direct secure connection to it (e.g. direct serial cable).

Network breath meters could increase the security of any secure inter-site links, in both dedicated military and commercial (e.g. intra-nets) applications. Our limited software experiments performed in a local network framework yielded very encouraging results, as is intuitively expected. Research should focus on associated base concepts and on building a proof-of-concept implementation.

4 Secure Multi-Party Cooperation

Consider the following scenario. Two military organizations plan to cooperatively work on a combat situation for their mutual benefit. Each organization would like its own requirements being satisfied (these requirements could be modeled as linear equations or linear inequalities). However, their

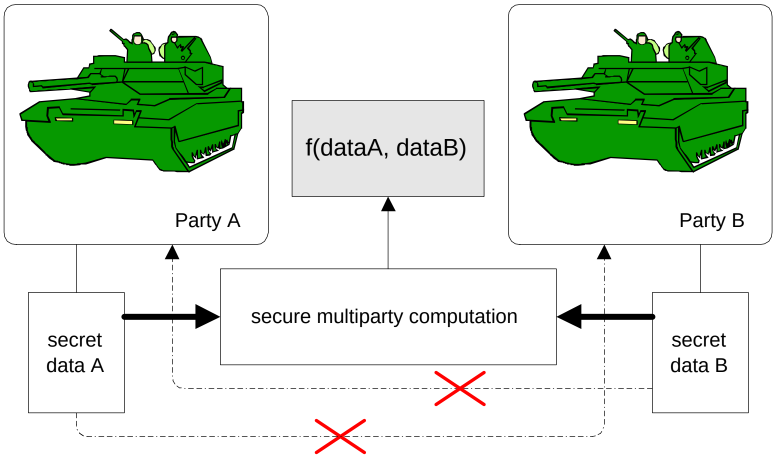

Figure 5: Secure Multi-Party Cooperation. No internal secure data has to be shared with other parties.

requirements are formulated using secure internal data which is not to be disclosed to any external party (see Figure 5).

An apparent deadlock situation is presented. Nobody likes to disclose their requirements to the other party, or even to a “trusted” third party. How could they cooperate so as to satisfy the requirements while preserving data privacy ?

This problem is conveniently named “Secure Multi-Party Computation Problem” (SMC) in literature. Research in the SMC area has been focusing on only a limited set of problems, while privacy concerning cooperative computations call for SMC studies in a variety of computation domains.

Generally speaking, a secure multi-party computation problem deals with computing any probabilistic function on any input, in a distributed network where each participant holds one of the inputs, ensuring independence of the inputs, correctness of the computation, and that no more information is revealed to a participant in the computation than can be computed from that participant’s input and output.

Currently, to solve the above problems, a common adopted strategy is the use of a trusted third party, assumption often unacceptable. Therefore protocols that can support joint computations while protecting the participants’ privacy are required.

Before we can study the problems, we need to identify and define the specific SMC problems for those computation domains. This is a non-trivial task, and we have developed a framework to facilitate this task. Based on our framework, we have identified and defined a number of new SMC problems for a spectrum of computation domains. Those problems include privacy-preserving database queries, privacy-preserving data mining, privacy-preserving intrusion detection, privacy-preserving statistical analysis, privacy-preserving geometric computations, and privacy-preserving scientific computations.

Instead of searching for new SMC problems in a random fashion, we have proposed a transformation framework that allows us to systematically transform normal computations (not necessarily security related) to secure multi-party computations. Further research on these resultant problems shows some of them are new problems, while some have been proposed earlier, although not discussed in the secure multi-party computation context.

In the past, secure multi-party computation research has mostly been focusing on theoretical aspects. Very few applied problems have been studied. In the literature, there are a limited number of examples of secure multi-party computation problems, such as the Private Information Retrieval Problem (PIR), joint signatures, joint decryption, networked elections, electronic bidding, privacy-preserving statistical databases and privacy-preserving data mining.

Encryption is a common technique to protect privacy; however it works only when the encrypted data is solely used for communication or storage. Although there are several works on conducting computation on encrypted data, the type of allowed computation is very limited and it is also difficult to implement and use. For example, the comparison operation is a quite simple computation for unencrypted data, but it is difficult to compare two encrypted data items without decrypting them first. Therefore, encryption cannot be directly used to solve secure multi-party computation problems; more sophisticated protocols are expected.

Anonymous communication protocols were designed to achieve somewhat related goals. However, anonymity techniques only help to hide the identity of the information originator, rather than the information being sent. In situations where sender identity needs to be protected, anonymous communication protocols are appropriate. However, there are situations where the content data, rather than the originator identity has to be protected in order to preserve privacy.

In many cases, in order to ensure the integrity and thus the trust-ability of the inputs from the other party, the participants need to prove their identity, rather than hide it, which renders anonymity useless. Therefore anonymity does not solve our problems, and cannot replace secure multi-party computation. Rather, by combining anonymity techniques with secure multi-party computation techniques, one can achieve better overall privacy more efficiently.

We propose the extension of our previous research in the area of secure multi-party computations in order to provide actual solutions for many of the issues associated with the scenarios presented earlier. In the extreme, the above scenario can be extrapolated to represent an actual limited required cooperation of a certain national army with a non-friendly or even hostile enemy organization (e.g. pre-armistice scenario). This becomes of extreme relevance in a cyber-attack framework where the attacker's identity is known and negotiations or any other type of hostile-party cooperation are to take place. Uses of secure multi-party cooperation can be envisioned also in post-attack emergency situations/contingencies.

More details on previous research in the area of secure multiparty computing can be found at **http://www.cerias.purdue.edu** or **http://www.cis.syr.edu/˜wedu**.

5 Conclusions.

Cyber Terrorism proves to be a real threat for our freedoms and privileges. The free flow of information facilitated through recent developments in computer networking is under daily explicit and implicit siege from negative de-stabilizing elements with potentially unrestricted access to its resources.

We believe that a series of extensive preventive measures need to be implemented immediately in order to protect positive cyber-experiences at one extreme and ultimately strategic national resources at the other extreme. A delicate balance needs to be observed and constitutional rights and duties of all participating parties need to be considered.

Nevertheless we feel that there exists a certain associated imminence as to the necessity of measure deployment and thus advocate rapid support and a decisive move toward a concerted anti-terrorism strategy. In this paper we presented some current and proposed research issues we find are of paramount relevance to the case in point.

In the broader area of Content Security we addressed issues such as information hiding detection and digital tamper-proofing, two vital tools in today's hostile and polluted networks. Intrusion Detection in the framework of System Security is a required yet often neglected issue in many of today's computing systems. We propose mechanisms based on human biometrics for run-time authentication as well as new concepts such as "network breath", in the area of inter-site network authentication. Securing multi-party cooperation is of paramount importance in scenarios of contingency and required interaction with hostile parties.

Most of the proposed research issues are in early stages of analysis, design and development. Given the recent shift in national priorities we feel more attention and possibly support is required to push them forward.