HELLENIC REPUBLIC
MINISTRY FOR DEVELOPMENT
GENERAL SECRETARIAT FOR RESEARCH & TECHNOLOGY
International S & T Cooperation Directorate
Bilateral Relations Division
14 -18 Messogheion av., 115 10 Athens, Greece
Tel (+3) 010 77 52 222, Fax (+3) 010 77 14 153

APPLICATION FORM

SCIENTIFIC and TECHNOLOGICAL COOPERATION

between

RTD ORGANISATIONS in GREECE

and RTD ORGANISATIONS in U.S.A, CANADA, AUSTRALIA, NEW ZEALAND, JAPAN, SOUTH KOREA, TAIWAN, MALAISIA and SINGAPORE

1. PROJECT TITLE

EAR: Early Warning System for Automatic detection of Internet-based Cyberattacks.

Research fields covered by the present call (*please check where applicable*):

Information and Telecommunication Technologies	***X***
Environmental Technologies	☐
Renewable Energy Sources Technologies	☐
Biotechnology and Genomics	☐
Space Technologies	☐
Nanotechnologies and Nanosciences	☐

This form must be submitted to the G.S.R.T. in four (4) copies

5. PROPOSAL SUMMARY

Project Leader in GREECE: Prof. Evangelos Markatos, Institute of Computer Science – Foundation for Research and Technology – Hellas, PO BOX 1385, Heraklion, Crete, Greece, GR71110

tel.: +30 2810 3891655 fax: +30 2810 391601 e-mail:markatos@ics.forth.gr

Project Leader abroad: Constantinos Dovrolis, College of Computing, Georgia Institute of Technology, Atlanta, Georgia, 30332-0280

tel.: +1 404 385 4205 fax: +1 404 385 0332 e-mail: dovrolis@cc.gatech.edu

CO-OPERATING ORGANISATIONS in GREECE (name, address, tel and fax no, e-mail)

1 Institute of Computer Science – Foundation for Research and Technology – Hellas, PO BOX 1385, Heraklion, Crete, Greece, GR71110, tel.: +30 2810 3891655 fax: +30 2810 391601 e-mail:markatos@ics.forth.gr

2.

3.

CO-OPERATING ORGANISATIONS ABROAD (name, address, tel and fax no, e-mail)

1. College of Computing, Georgia Institute of Technology, Atlanta, Georgia, 30332-0280, tel.: +1 202 385 4205 fax: +1 404 385 0332 e-mail: dovrolis@cc.gatech.edu

2.

3.

Research Field(s) (select from the first page)
Information and Telecommunication Technologies

Requested Funding (up to **60.000** €):

5. PROJECT LEADER IN GREECE

3.1 Full Name: Evangelos Markatos

3.2 Profession – Duties: Associate Professor of Computer Science

3.3 Affiliated Institution: Institute of Computer Science (ICS) – Foundation for Research and Technology Hellas (FORTH)

5.3 Mailing Address, tel, fax and e-mail of the project leader:
Institute of Computer Science – Foundation for Research and Technology – Hellas, PO BOX 1385, Heraklion, Crete, Greece, GR71110, tel.: +30 2810 3891655 fax: +30 2810 391601 e-mail:markatos@ics.forth.gr

3.5 % of working time allocated to the proposed project: 20%

Attach resume and list of publications (as an annex)

4. PROJECT LEADER ABROAD

4.1 Full Name: Constantinos Dovrolis

4.2 Profession – Duties: Assistant Professor

4.3 Affiliated Institution: College of Computing, Georgia Institute of Technology,

4.4 Mailing Address, Tel, Fax and e-mail of the project leader:
College of Computing, Georgia Institute of Technology, Atlanta, Georgia, 30332-0280, tel.: +1 202 385 4205 fax: +1 404 385 0332 e-mail: dovrolis@cc.gatech.edu

Attach cv and list of publications (as an annex)

5. PARTICIPATION OF GREEK ENTERPRISE *(if applicable)*

5.1 Trade name and legal status: FORTHnet S.A.

5.2 Mail address: PO Box 2219, Science And Technology Park of Crete, 71003, Vasilika Vouton, Heraklion, Kriti, GREECE

5.3 tel and fax no: Telephone: 0030 2810 391200
Fax No: 0030 2810 391207

5.3 Full name of:

- President of Board of Directors: Kostas Klironomos
- Managing Director or General Director: Pantelis Tzortzakis

5.5 Financial status of the enterprise

5.6 Statement, signed by the Legal Representative of the entreprise:

«The Legal Representative of the entreprise declares that he took cognizance of the submission of the project and as long as it will be approved, the entreprise will cover the proposed financial contribution»

Signature

5.7 Short description of the entreprise (year of establishment, history, structure, products/services):

5.8 Presentation of the research needs and scopes of the enterprise conducting to this project (existence of a research lab, existing international cooperation):

6. PROJECT DESCRIPTION (for all partners)

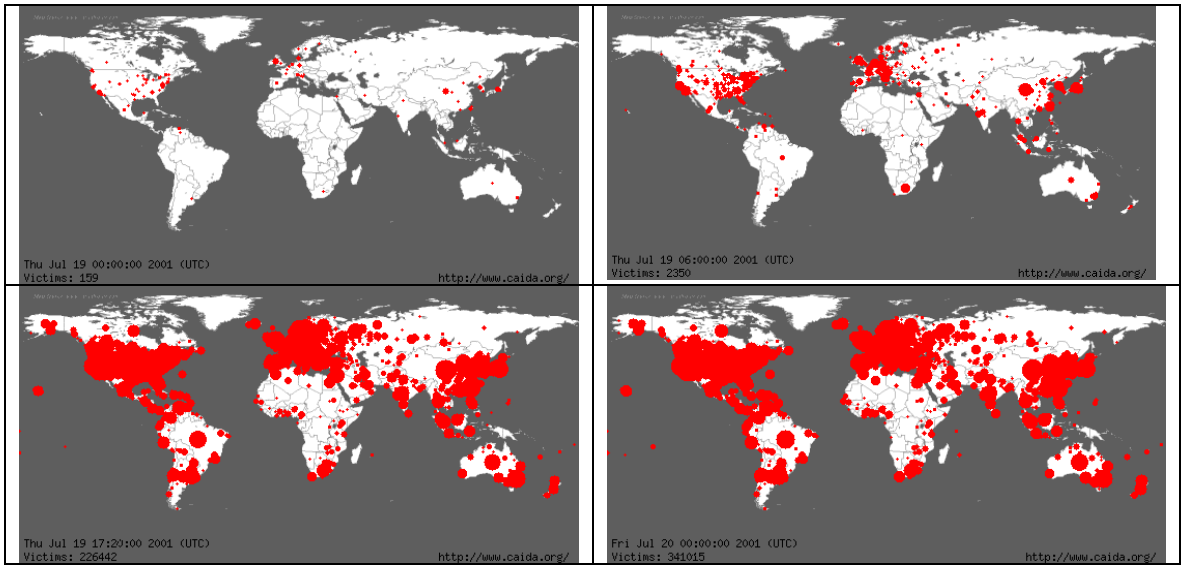

Figure 1 : Spread of the CORE-RED worm/exploit during July 19th 2001. We see that at 00:00 the worm had infected only 159 computers, which were increased to 2,350 by 06:00. By dinner time, the worm had spread to more than 226,442 computers, and by midnight it had infected 341,015 computers covering practically every corner of the earth (the images are courtesy of CAIDA[1]).

5.3 Abstract

Over the last few years, the Internet has been repeatedly used as a medium to launch attacks against computer and communication subsystems. Such attacks, which are usually called cyber-attacks may **disable a large number of computers**, which may in turn **paralyze critical infrastructures** including telecommunications, provision of electric power, transportation, water supplies, athletic infrastructure, and commerce. Such cyber-attacks propagate rapidly and may have profound impact. For example, in 2001 a computer worm/exploit named CODE-RED was released on the Internet and infected more than 340,000 computer systems in less than 24 hours. Indeed, Figure 1 shows that CODE-READ spread very rapidly even during business hours, and within a day it infected computers in practically every corner of the earth. This CODE-RED incident was not an isolated case. Actually, the frequency of such events is currently on the rise. For example, Figure 2 shows that the number of computer-related vulnerabilities reported to CERT (the Carnegie Mellon University Computer Emergency Response Team[2]) is currently increasing exponentially, doubling every year or so for the last 2-3 years. To reduce the number, spread, and impact of Internet-related attacks, we propose to do research towards the creation of **early warning systems** that can detect cyber-attacks quickly and can respond to them efficiently.

[1] http://www.caida.org
[2] http://www.cert.org

This proposal is a step towards the direction of designing, implementing, and deploying early-warning systems that are able to detect computer attacks at their infancy. Capitalizing on the cooperation between the first ISP in Greece, the largest research center in Greece, and one of the most prominent U.S. Universities with significant experience on network monitoring, this proposal aims to develop and deploy systems that detect and respond to cyberattacks as early as possible. To do so, we plan to use novel techniques that process all network packets in real-time searching for (old and new) signatures of viruses, worms, and in general any forms of attacks. We believe that by co-relating cyberattack-related information from various points of an ISP's network we will be able to increase the probability of accurately pinpointing Internet-based attacks as soon as they start to spread, reducing their impact on our critical infrastructures that are connected to the Internet.

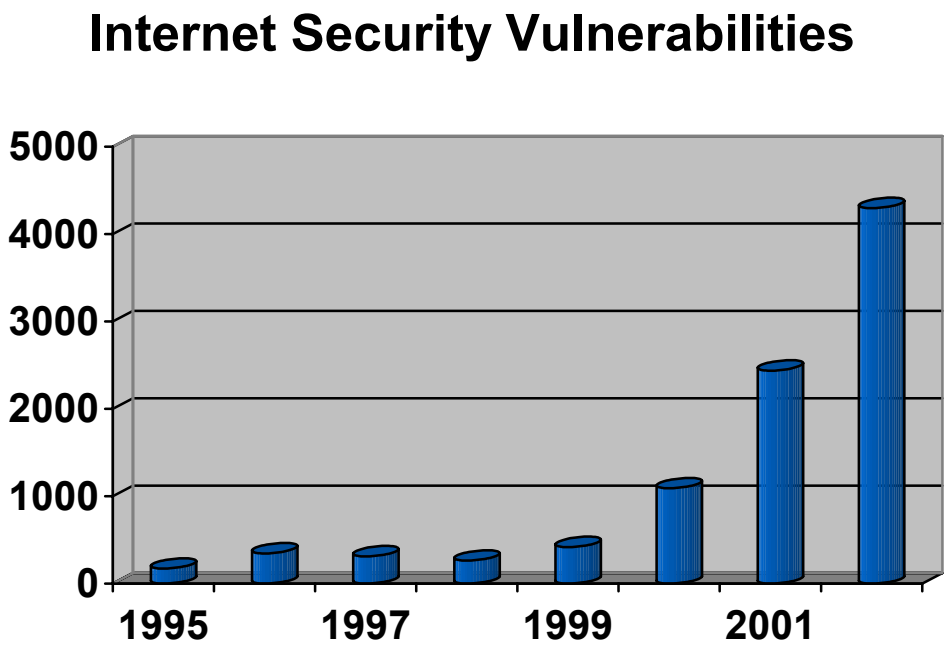

Figure 2 : Number of Internet security-related vulnerabilities reported to CERT (data courtesy of www.cert.org: actual data for 1995-2001, projected data for 2002).

5.4 Objective(s) of the proposal and expected results

The **objectives** of this proposal are:

- To **develop** an **early-warning system** that can detect cyberattacks at their infancy, so that appropriate countermeasures can be taken.
- To **deploy** this **early-warning system** at several strategic points in the Greek cyberspace so as to calibrate its accuracy, test its performance, and explore whether it is beneficial to deploy a larger-scale version of the system.
- To provide the mechanisms that will create a **safer** and **more secure Greek cyberspace** through careful and persistent network monitoring
- To develop and evaluate the techniques that are necessary to identify new (previously unknown) worms and viruses, and automatically alert system and network administrators of the existence of these new threats

We expect this project to create a large number of interesting scientific, technological, and commercial **results**.

On the scientific side, we expect that one outcome of the project will be the design and calibration of a **network-based distributed methodology** for **identifying new cyber-attacks**. This methodology will be tested in the lab, and demonstrated in a **real-world prototype**.

Having proven its robustness, the methodology and the prototype may be later used as the basis of creating a **product or service** that will be provided to end users who want to be protected from cyberattacks. All the partners, under the successful coordination of FORTHnet will explore the best ways to turn the developed prototype into a service that will be supplied in the rapidly growing market of Internet security. We believe that the consortium is in a unique position to guarantee the **viability** of the project and its **industrial exploitation**. That is, both FORTH and FORTHnet have created successful spin-off companies that have commercially exploited results produced in part within research projects.

Besides the above direct results of the project we believe that there will also be several indirect results. We expect that this project will form the basis of a **closer cooperation** between ICS-FORTH and Georgia Tech. We believe that **security** in the **cyberspace** is a topic that **concerns both Greece and the United States** and transfer of knowledge in this important topic will be beneficial to both countries. The United States in particular, have recently invested heavily in research and technology that may lead towards reducing or even eliminating the threat of a significant attack in their computer and communication systems that are connected to the Internet. This project will be a first step towards the close cooperation between the Greek and the U.S. partners. This collaboration can be further strengthened through multi-national large-scale research projects, and possibly through commercial cooperation.

5.5 Methodology to be used / justification

In this project we will develop a methodology for detecting new cyber-attacks as early as possible. To do so, we will both use proven technologies, but we will also conduct state-of-the-art research to develop and evaluate new ones. Our research contributions will revolve around the following issues:

- Payload Inspection
- Packet Header inspection
- Connection/Session characteristics
- Honeypots
- Networks of packet sniffers

Payload Inspection

Network Packet payload inspection has been traditionally used to detect worms/viruses hidden within network packets. Most known worms/viruses consist of a known string of bytes, usually referred to as the worm/virus "signature". Intrusion detection systems inspect all packets for such signatures in order to find whether the packet contains (parts of) a known virus [Roesch99]. In this research project we plan to use payload inspection not only to detect known worms/viruses, but also to find new (unknown) ones. Obviously, the most significant challenge in automatically finding new worms/viruses is in identifying which network packets belong to a virus/worm distribution. To so do, we plan to capitalize on the fact that in order to be effective, worms and viruses need to infect several computers, a task that requires a significant amount of traffic and a large number of transfers of the virus/worm over the network. By inspecting the network packets' payload we plan to detect this large number of worm transfers as follows:

> We will inspect the payload of network packets and record all N-bytes long sub-strings they contain. Some of those N-byte sequences will have higher frequency of occurrence than others. Since viruses and worms, in order to be effective, need to generate (and send) several copies of them, packets that contain viruses and worms will appear amongst the most frequent of the N-byte sequences.

Of course, amongst those frequently occurring sequences, along with the new worms, there will be several other sequences that do not belong to new worms, including sequences that belong to known worms/viruses, to popular web pages, to popular software downloads, etc. However, we believe that we can easily filter those non-new-worm sequences out. For example, known worms have known signatures, and it will be rather easy to determine whether a popular N-byte sequence belongs to a known worm, or not. We can even filter easily out popular N-byte sequences that belong to known pages: traditionally popular web sites are generally known (or can be easily found after a short training period) and thus it will be rather easy to determine whether an N-byte sequence belongs to traffic generated by a popular web site. However, besides known popular web sites, there also exist web sites that become popular only for a few days (e.g. after the release of some new software distribution), generate a large number of hits, and attract a large number of clients, which are usually called "flash crowds" [SMB02]. Fortunately, new popular web sites and new popular software downloads can

easily be distinguished from new viruses because of their different traffic source/destination distributions. For example, popular content tends to be served from a **small number of servers** (i.e. a single server or a Content Distribution Network) to a very **large number of clients**. On the other hand, worms and viruses tend to originate from many computers and are destined to several others. Therefore, by comparing the ratio between the senders and receivers of packets where popular signatures are being found, we can distinguish between popular content (served from few servers to lots of clients) and harmful viruses (Send from several clients to several others).

Obviously the proposed approach (i.e. register and examine popular N-byte sequences) needs to be appropriately defined and calibrated against real packet traces. For example, we expect that registering all N-byte sequences for all packets will require significant memory and computational resources. To reduce the memory and computational needs of our method we plan to use sampling. Fortunately, recent results suggest that sampling methods can have a very high accuracy (larger than 99% in some cases) [EV02].

There are cases, however, where payload inspection seems to be insufficient for the detection of new (and old) viruses. For example, polymorphic computer viruses continually change themselves, possibly encrypting their code with a different key each time they are being transferred, making detection by payload examination practically impossible. Even in these cases, however, we can use payload inspection to help us identify new viruses as follows: encrypted viruses appear as an (almost perfectly) random sequence of bits, while ordinary packets appear to have a known (non-random) distribution of bits. Thus, calculating the entropy of the payload of each packet we can identify sessions that seem to be encrypted. These sessions can be inspected further by other tools, which will later determine whether they carry an encrypted virus or legitimate encrypted traffic.

Packet header inspection

If payload inspection is not enough, we can complement it by exploiting inspection of packet headers. By inspecting the headers of the packets we can easily focus on specific subsets of traffic, like email traffic only, reducing the strain on a system that focuses on identifying viruses that spread using email attachments. By inspecting the headers we can also identify the spread patterns of viruses. For example, worms and viruses tend to spread from many computers to many others leaving a very characteristic traffic pattern that identifies them like a signature. On the contrary, most of the other traffic is directed from a small number of servers to a large number of clients (e.g. web traffic, ftp traffic etc.). Inspecting the traffic patterns of the packets where popular N-byte sequences appear will help us identify which of those packets might belong to a worm/virus spread.

Connection/Session Characteristics

Although payload and header inspection will help us identify a significant percentage of cyberattacks, it is possible that very clever attackers use various forms of polymorphism to encrypt their worm and hide their attack. Although it is

possible to identify encrypted network packets, these packets besides polymorphic viruses may also contain legitimate traffic like secure banking transactions, encrypted personal information, etc. We believe that we can use connection and session characteristics to distinguish legitimate traffic from polymorphic viruses. One such characteristic may be the size of the transferred data during a session. All instances of a polymorphic virus will probably have the same size and will transfer the same data. Actually, for each network session we expect to be able to define a "signature" of the session. The signature will depend upon the spatial and temporal distribution of the session's packets. We believe that polymorphic viruses will change this distribution and several different transmissions of the same polymorphic virus will depict similar distributions.

Honey-pots

One way to track the behavior of intruders and their intrusion tools is through innocent-looking computers called honeypots. Honeypots are constructed so that they are easy to break in, and thus hackers break in them and use them as a means to launch attacks against third-party computers. In this way hackers can cover their tracks. However, the honey-pot can be set up in such a way as to lure and carefully study hackers (without them knowing it), and find about their new worms. In this proposal we are going to install such a honeypot in order to study the spread of viruses the moment they are being released.

Networks of packet sniffers

Although all the above methods can be used in any single point in the network, we plan to deploy them in several different points – in several different "sniffers" as they are usually called. By combining information from several sniffers we will be able to more accurately identify new worms. The network of sniffers allows us to detect the distribution patterns of worms and viruses. The network will also enable us to detect slowly spreading viruses that do not generate significant amounts of traffic and thus are difficult to detect in any single point of the network.

However, we believe that using a network of sniffers we will probably be able to find some of the polymorphic viruses as well. Traditional anti-virus systems detect (known) polymorphic viruses by simulating the execution of the victim application so that the virus decrypts itself. Once the virus is decrypted, anti-virus systems transform it into a canonical form, then compute a signature out of this canonical form, and compare it against their database of known signatures. Our approach (of computing all N-tuples and fiding the most popular of them) may also be used in this case, but instead of applying it into the raw network packets, it is being applied in the canonical form. By combining the popular N-tuples from several different sniffers, our approach will be able to make a more informed decision on whether an encrypted programs corresponds to a virus/worm or not.

Advantages

The main advantages of our approach is that it can be easily implemented and can help network administrators deal with new viruses/worms very fast, because our approach not only **detects new viruses and worms**, it also **finds their signature**, that is the N-byte sequence that identifies them. This automatic discovery of the virus/worm signature can be very important for administrators, who

traditionally, even if they identify a new worm they can do very little about it. They usually wait for a patch, or the release of the next version of the software the worm is exploiting. In the meantime (which can be several days) they manually shut down every running copy of the exploited software. Our approach, however, empowers them with a signature (the N-byte sequence) that they can add to their Intrusion Detection System, or even to their firewall. Thus, all packets that contain the offending signature will be banned from their system.

Another advantage of our approach is that it can be deployed in the computer network instead of the end hosts. Computer networks are usually managed by security-conscious system administrators. On the contrary, end-hosts usually belong to individual users that may not have the knowledge (or the time) to install and maintain security-related software. However, even in the case where they do install and maintain such software, they usually can do very little to respond to a worm attack. Thus, we believe our approach is more easily deployable and may have a more significant impact than approach that is being deployed at end-hosts.

5.6 Main phases of the project, part of each side

The main phases of the project will be as follows:

- Phase 1: Requirements Analysis
- Phase 2: Design
- Phase 3: Implementation - Integration
- Phase 4: Deployment – Evaluation
- Phase 5: Commercial Evaluation

Phase 1: Requirements Analysis
Phase leader: FORTHnet
Phase participants: FORTH, GA Tech
Deliverables: D1.1: Requirements Analysis (document)

In this phase, leaded by FORTHnet, the requirements of the system will be defined. These will include both functionality requirements as well as performance constraints. At the same time, a survey of the state-of-the-art systems, mostly done by GA tech with the participation of FORTH, will identify existing solutions and their shortcomings.

Phase 2: Design
Phase leader: FORTH
Phase participants: FORTHnet, GA Tech
Deliverables: D2.1 System Design (document)

In this second phase we will focus on the design of the early-warning system, i.e. define the main components, their interfaces, and their combinations. Using various different methods of evaluation, including trace-driven simulations we will identify the prediction accuracy of the system, both in identifying new attacks, as well as reducing the number of false positives (i.e. identifying attacks that never happened). Based on the results of these simulations we will be able to calibrate and combine the basic intrusion detection algorithms that will be eventually deployed.

Phase 3: Implementation – Integration
Phase leader: FORTH
Phase participants: FORTHnet
Deliverables: D3.1 System Implementation

Based on the results of the previous phase, this third phase will implement and test the initial functionality of the system. FORTH will take the lead in the implementation, while FORTHnet will provide assistance at the integration. Will deal with the integration of the system.

Phase 4: Deployment – Evaluation
Phase leader: FORTH
Phase participants: FORTHnet, GA Tech
Deliverables: D41 System deployment and evaluation (document)

This phase, leaded by FORTH, will deploy the system in strategic places of the Greek Cyberspace and evaluate its effectiveness. FORTH will take the lead in both the deployment of the system and its evaluation. GA Tech will participate by designing the experimental environment and by giving feedback on the experimental results. FORTHnet will participate during the installation phase. This will complement the lab-based evaluation of phase 2, with an evaluation on a real environment using real traffic and (possibly) real attacks.

Phase 5: Commercial Evaluation
Phase leader: FORTHnet
Phase participants:
FORTH
Deliverables: D5.1: Commercial Viability study

This phase will study the commercial viability of the prototype developed. The partners will work towards the creation of a model that describes the operations that are necessary for the cyberattack detection process. The model will contain the entire process including existing activities, information, use of personnel, and use of equipment. The proposed methodology may include business process models based on standards set by ISO and CEN, such as the existing CEN ENV40003/12204 and their newly drafted successors ISO/CEN prEN19439/19440 for Enterprise Modeling. FORTHnet, the phase leader, will deploy its expertise and its modeling tools in order to produce a consistent model at different abstraction levels. In addition, the phase will develop a business plan and the commercial agreements that would ensure the smooth introduction of these results into the market. The work in this phase is analysed as follows:

5. Development of **Business Plans**: The market status will be traced periodically (semi-annually or at most annually)and the results will be evaluated.
6. **Commercial Agreement**: This task is devoted to the development of a Commercial Consortium Agreement. The agreement will tackle all the IPR and exploitation issues. The Commercial Agreement duration, terms and conditions will be defined during the course of this phase and will be activated after the end of the project.

Deliverables:

Phase	Name of deliverable	Deadline
1	D1.1: Requirements Analysis	M6
2	D2.1 System Design	M12
3	D3.1 System Implementation	M18
4	D4.1 System deployment and evaluation	M24
5	Deliverables: D5.1: Commercial viability study	M24

5.7 State of the art at international level

Cyber-attacks as are currently being exemplified by worms and viruses are a rather recent phenomenon. The word “worm” was introduced as late as 1982 [SH82] to describe a new model of distributed computing, a model based on self-replicating programs that moved and computed through the network. The self-replicating nature of worms guaranteed their robustness and their longevity. Unfortunately, soon after their first introduction, worms were used as a robust intrusion mechanism. The first major attack based on a worm was launched in the evening of November the 2nd, 1988 [Spafford89]. That evening the first large-scale worm was release in the Internet, leading to the infection and shutdown of more than six thousand computers. Since then almost all computers connected to the Internet are under increasing attacks of worms, viruses, and other cyberthreats.

To reduce the number and effect of these attacks, several levels of protection are being used. Most systems and organizations today are protected against viruses through anti-virus software [Nachen97]. These applications scan computer memory, computer files, and several sources of traffic (like email and web pages) to identify whether they contain harmful viruses. To do so, the have a database of known viruses. For each virus the database keeps a “signature” (a sequence of bits) that identifies the virus. This “signature” may be a portion of the virus code, the virus name, or something else. Since new viruses are being release every day, this database is usually automatically updated every few weeks or so.

Although anti-virus systems are able to detect and eliminate known viruses they are of little help in detecting worms and other kinds of intrusions. To alert administrators of such attacks, Intrusion Detection Systems are currently being deployend in networks and in end-hots [Base99]. Network-based Intrusion Detection systems “sniff” all the traffic in a segment of the network and compare it against known attack signatures. When a network packet matches this signature, it is logged and (optionally) an alert is being sent to the system’s administrator. Host-based Intrusion Detection Systems try to identify possible intruders by examining system files, operating system audit trails, and various system parameters. All the above Intrusion Detection Systems depend on the existence of a “signature” database that contains a signature for each possible type of intrusion. Obviously their main disadvantage of systems that depend on such a database is that they cannot find new (not included in the database) attacks. Anomaly detection systems try to remedy this situation by measuring various system parameters and reporting any anomalies observed in these measurements. For example, some systems measure network traffic, or CPU load. If the traffic exceeds a threshold, this probably indicates an anomaly. Although anomaly detection systems are able to identify Denial of Service attacks, they also suffer from a large number of false positives: i.e. an increase in the network load does not always signify an intrusion attempt.

To reduce the number of false positives and increase the accuracy of detecting new attacks, several researchers have proposed the use of neural networks. Cannady and Mahaffey proposed the use of neural networks to successfully detect ftp-based intrusions [CM98]. Lippmann and Cunningham used neural networks to improve accuracy using better signatures and discriminate training [LC00].

Given that computer viruses resemble real viruses, Forrest suggested combating them using systems inspired from the Human immune system [FHS96]. Forrest several host-based systems that monitor various sources on information including system calls to detect changes in usual patterns, which are indication of intrusions.

Our approach shares similar goals with much of the previous work. For example, we are interested in exploring the possibility of automatic detection of (new) worms/viruses. However, in contrast to most previous systems, we are interested in automatic detection based on **information available in the network**. We are interested in investigating new worms, viruses, and other cyberattacks, before they infect a computer and before they leave an audit trail on it. Moreover, we are interested in developing an approach can be easily deployed and can potential protect a large number of computers. That is why we base our system on several points in a network and not in the end (victim) hosts. On the contrary, systems that depend on pieces of software that run on the end hosts are much more difficult to deploy since they depend on the cooperation of the end-host owners. In addition, our methodology differs from previous approaches as well. For example, we plan to detect worms/viruses by combining information from several different sniffers on a network. We plan also to use novel techniques that combine payload inspection and packet header inspection, that to our knowledge, have not been used and evaluated for this purpose in the past.

6.6 Common interest of both sides and relevance for cooperation

We believe that security-related research is **particularly important for both Greece and the United States**. Greece needs to monitor and secure its cyberspace, against possible attacks, especially in light of the upcoming 2004 Olympic Games. The United States, on the other hand, have already taken several steps towards securing their cyberspace, and are particularly active in reducing cyberattacks. To underline the importance of such an activity, the "President's Critical Infrastructure Protection Board", in the "National Security Strategy to Secure Cyberspace"[3] strongly encourages that "***the Federal government should work with … international organizations to foster the enstablishment of … international watch and warning networks to detect and prevent cyberattacks as they emerge***". Underlining the important of computer security, on November 27, 2002, President Bush signed $900 million "Cybersecurity Act", legislation dedicating more than $900 million over five years to security and education to protect the US infrastructure against hackers and terrorists.

But, besides the countries, this proposal is important for each participating partner as well. **FORTHnet**, historically the first, and currently one of the largest ISPs in Greece operates one of the largest Greek networks and is particularly interested in reducing the attacks originating from or destined to its customers. In addition by operating such a larger network, FORTHnet will provide the most representative traffic that can be found in the Greek cyberspace. **GA Tech**, on the other hand, has a long history with distributed monitoring of large-scale networks, involving nodes all over the world including Greece and the United States. By participating in this project, GA Tech will provide the necessary expertise to design and carry out distributed detection of cyberattacks. On the other hand, by participating in these experiments, the researchers of GA tech will be able to calibrate their methodologies and test them in more realistic conditions. Finally, **FORTH**, by being one of the largest research centers in Greece, and by being a pioneer of the Internet in Greece is very interested in this project because it will enable it to strengthen its expertise on Internet security. FORTH will capitalize on its expertise on network monitoring for performance and extend it to include network monitoring for safety and security as well.

Last, we should emphasize that all partners of the project have demonstrated the **ability to cooperate** successfully in the past. For example, FORTH and FORTHnet cooperate in several externally funded research projects, some of then in the area of network monitoring. In addition, FORTH and GA Tech cooperate by sharing network resources and data necessary for large-scale experimentation. This project will enable all partners to strengthen and formalize their cooperation under a single project that is important both to them and to their countries.

[3] http://www.whitehouse.gov/pcipb/

PROJECT DESCRIPTION for FORTH

FORTH's participation was described in the previous sections 6.1-6-6

PROJECT DESCRIPTION for Georgia Tech
Georgia Tech's 's participation was described in the previous sections 6.1-6-6

PROJECT DESCRIPTION for FORTHnet
FORTHnet's participation was described in the previous sections 6.1-6-6

Bibliography

[Base99] R. Base: “Intrusion Detection”, Pearson Higher Education, ISBN 1578701856, 1999.

[CM98] J. Cannady and J. Mahaffey “The application of Artificial Neural Networks to Misuse detection”. In proceedings of the First International Workshop on the Recent Advances in Intrusion Detection, 1998.

[EV02] C. Estan and G. Varghese: "New Directions in Traffic Measurement and Accounting”, in Proceedings of the ACM SIGGCOMM Conference, 2002.

[FHS96] S. Forrest, S. Hofmeyr, and A. Somayaji: "Computer immunology". *Communications of the ACM*, 40(10), pp. 88-96, 1997.

[LC00] R. Lippmann and R. Cunningham: “Detecting Computer Attackers : recognizing patterns of malicious, stealthy behavior” - MIT Lincoln Laboratory - Presentation to CERIAS 11/29/2000

[Nachen97] Carey Nachenberg: “Computer virus-antivirus coevolution”. *Communications of the ACM*, 40(1):47-51, Jan. 1997.

[Roesch99] M. Roesch: "Snort: Lightweight Intrusion Detection for Networks", in Proceedings of the 1999 USENIX LISA Systems Administration Conference, November 1999.

[SH82] J. Scoch and J. Hupp: “The “worm” programs – early experiments with a distributed computation”. *Communications of the ACM*, 22(3):172-180, March 1982.

[SMB02] T. Stading, P. Maniatis, and M. Baker: “Peer-to-Peer Caching Schemes to Address Flash Crowds”. In Proceedings of the "Fist International Peer To Peer Systems Workshop (IPTPS 2002)", March 2002.

[Spafford89] E. Spafford: “Crisis and aftermath (the Internet worm)”, Communications of the ACM, 32(6):678-688.

7. DISSEMINATION AND EXPLOITATION OF RESEARCH RESULTS

Academic and Research dissemination

The results of this project will be disseminated throughout the research community via a variety of mechanisms including, publications in **conferences**, **journal** paper publications, and **workshop** presentations. All the involved researchers have an excellent publication record, and plan to improve it through their involvement in this project. Besides, however, the traditional dissemination channels, this project will disseminate information through the **web**, through **mailing** lists, and through participation in **working groups**, like the Internet Engineering Task Force working groups. The partners will also study the possibility of preparing workshop or a BoF session in association with an established conference in informatics, like the TERENA networking conference.

Commercial Exploitation

The prototype that will be developed in this project will be thoroughly evaluated and tested under real conditions for the automated and coordinated detection of cyberattacks. Testing will take place within both core data center activities of FORTHnet and also at selected customer premises. If the experiments indicate that the system is successful, and robust over a large period of time, then the partners will consider whether it will be beneficial to develop a commercial version of the prototype and sell it as a product, and operate it and sell it as a service.

Early warning systems are offering the ability to enhance the SLA offering of the service provider and also increase the credibility of network services. The intensive problems faced by continuous and annoying intrusions, are revealing the need to upgrade the quality of security systems, and to invest towards preventive mechanisms rather than reactive systems and procedures. Henceforth, any commercial provider is interested to enrich their tools, to efficiently protect their networks and customers, to update their competitive advantage, to add value to their provided service, and to reduce their management overhead and troubles. FORTHnet will quantify the interest of its customers to such services, through investigating into its major accounts customer base, and define an exploitation plan for early warning methods and tools.

Fortunately, the **market** is very **favorable** for the development of security-related products and services. For example, IDC, the world's leading provider of technology intelligence industry analysis and market data, predicts that the security market will increase from $66 billion in 2001 to $155 billion in 2006[4]. In particular intrusion detection software is expected to have a compound annual growth rate of 37%, reaching a market size of more than one billion dollars by 2003[5]. Given the world's concern about security, we expect that the Internet security-related market will continue to increase for the years to come, and thus, once our system is tested, the market will be favorable for it.

[4] http://www.newsfactor.com/perl/story/19809.html
[5] http://www.infoworld.com/articles/hn/xml/01/04/18/010418hnsecuritymarket.xml

8. PROJECT TEAM

5.3 Composition of the Greek research team: project leader and researchers, % of working time of each researcher allocated to the project:

Name	Title	Percentage of time dedicated to this project	Organization
Evangelos Markatos	Associate Professor of Computer Science	20%	FORTH
Angelos Bilas	Associate Professor of Computer Science	10%	FORTH
Manolis Katevenis	Professor of Computer Science	10%	FORTH
Kostas Xinidis	Graduate Student	100%	FORTH
Yiannis Haritakis	Graduate Student	100%	FORTH
Kostas Polychronakis	Undergraduate Trainee	50%	FORTH
Spyros Antonatos	Undergraduate Trainee	50%	FORTH
New graduate student	Graduate Student	100%	FORTH
New graduate student	Graduate Student	100%	FORTH
Vasilis Spitadakis	Technical Manager	10%	FORTHnet
Manolis Petsagourakis	Technical Manager	40%	FORTHnet
Maria Manasaki	MTS	50%	FORTHnet
Constantine Dovrolis	Assistant Professor of Computer Science	10%	GA Tech
Graduate Student	Graduate Student	50%	GA tech
Graduate Student	Graduate Student	50%	GA tech

5.4 Experience of the Greek and foreign teams on the proposed subject

Both the Greek and the international teams have significant experience in the field of Internet Technologies in general and network monitoring in particular as can be seen from the attached CVs of the key researchers. ICS-FORTH has traditionally been a pioneer of Internet Technologies in Greece. For example, ICS-FORTH was the fist Greek node to connect to the Internet. ICS-FORTH also created the first Greek ISP: FORTHnet. Continuing its tradition of being a pioneer of the Internet in Greece, ICS-FORTH conducts state-of-the-art research and development in Internet Technologies. FORTHnet, the first and one of the largest ISPs in Greece, has also significant experience in the area of Internet systems and technologies. FORTHnet has also acquired significant experience in network monitoring through projects related to traffic accounting and billing. Finally, Georgia Tech has significant experience in network monitoring and especially in distributed large-scale monitoring experiments.

9. MAIN STAGES OF THE PROJECT – TIME SCHEDULE

According to the phases of 6.4

Duration (in months) **Stages** (briefly)	**6**	**12**	**18**	**24**
(G.s.) Phase 1: Requirements Analysis (F.s.)	←----------→ ←----------→			
(G.s.) Phase 2: Design (F.s.)	←---- ←----	----------→ ----------→		
(G.s.) Phase 3: Implementation – Integration (F.s.)		←-------	----------→	
(G.s.) Phase 4: Deployment – Evaluation (F.s.)			←------ ←------	----------→ ----------→
(G.s.) Phase 5: Commercial Evaluation (F.s.)			←------	----------→
Intermediate Report	(G.s …..... (F.s.)	*		
FINAL REPORT	(G.s) …..... (F.s.)			□

(G.s): Greek side
(F.s.): Foreign side
Use the symbols:
↔ Phase duration
✷ Intermediate Report

5. Brief description of the requested equipment and consumables
(in separate page)

13. (στα ελληνικά, για τις ανάγκες της απόφασης που θα εκδώσει η Γ.Γ.Ε.Τ., στην περίπτωση έγκρισης της πρότασης)

13^{α} Αντικείμενο της πρότασης

Το αντικείμενο αυτής της πρότασης είναι ο σχεδιασμός και η υλοποίηση συστήματος για την συντονισμένη και αυτόματη αναγνώριση επιθέσεων στην ασφάλεια υπολογιστικών συστημάτων μέσω του Διαδικτύου. Τέτοιες επιθέσεις, οι οποίες στην πολύχρωμη ορολογία της πληροφορικής ονομάζονται, ιοί (viruses), σκουλήκια (internet worms), και κατορθώματα (exploits), εκμεταλλεύονται αδυναμίες ή και σφάλματα υπαρχόντων προγραμμάτων με σκοπό την διείσδυση στα υπολογιστικά συστήματα στα οποία εκτελούνται τα προγράμματα αυτά. Το παρόν έργο έχει σαν στόχο την διεξαγωγή έρευνας, καθώς και την ανάπτυξη και αξιολόγηση πρωτότυπου, για την αυτόματη ανίχνευση τέτοιων επιθέσεων όσο το δυνατόν ενωρίτερα, πριν προλάβουν να κάνουν σημαντικές καταστροφές.

13^{β} Σύνθεση των ερευνητικών ομάδων

Ονομα	Τίτλος
Ευάγγελος Μαρκάτος	Αναπληρωτής Καθηγητής Επιστήμης Υπολογιστών
Αγγελος Μπίλας	Αναπληρωτής Καθηγητής Επιστήμης Υπολογιστών
Μανώλης Κατεβαίνης	Καθηγητής Επιστήμης Υπολογιστών
Κωσταντίνος Δόβρολης	Assistant Professor of Computer Science
Βασίλης Σπιταδάκης	Τεχνικός Διυεθυντής
Μανώλης Πετσαγουράκης	Τεχνικός Διυεθυντής
Μαρία Μανασάκη	Επιστήμων Υπολογιστών
Κώστας Ξινίδης	Μεταπτυχιακός Φοιτητής
Κώστας Πλολυχρονάκης	Προπτυχιακός Φοιτητής
Σπύρος Αντωνάτος	Προπτυχιακός Φοιτητής
Νέος μεταπτυχιακός Φοιτητής	Μεταπτυχιακός Φοιτητής
Νέος μεταπτυχιακός Φοιτητής	Μεταπτυχιακός Φοιτητής

13ν Χρονοδιάγραμμα της πρότασης

Σύμφωνα με τα στάδια του 6.4

Διάρκεια (σε μήνες) / **Στάδια** (συνοπτικά)	**6**	**12**	**18**	**24**
(Ελ.πλ.) Ανάλυση Απαιτήσεων (Αλ.πλ.)	←----------→ ←----------→			
(Ελ.πλ.) Σχεδιασμός του Συστήματος (Αλ.πλ.)	←---- ←----	----------→ ----------→		
(Ελ.πλ.) Υλοποίηση – Ολοκλήρωση (Αλ.πλ.)		←------- ←-------	----------→ ----------→	
(Ελ.πλ.) Εγκατάσταση – Ανάλυση Επιδόσεων (Αλ.πλ.)			←------ ←------	----------→ ----------→
(Ελ.πλ.) Μελέτη εμπορικής εκμετάλλευσης (Αλ.πλ.)			←------ ←------	----------→ ----------→
Ενδιάμεση Εκθεση	(G.s …..... (F.s.)	*		
ΤΕΛΙΚΗ ΕΚΘΕΣΗ	(G.s) …..... (F.s.)			□

(Ελ.πλ.): Ελληνική πλευρά
(Αλ.πλ.): Αλλοδαπή πλευρά
Χρήση των συμβόλων:
↔ διάρκεια σταδίου
✷ Ενδιάμεση έκθεση
□ ΤΕΛΙΚΗ ΕΚΘΕΣΗ