# Anticipating Security Threats to a Future Internet

FORWARD Consortium: H. Bos (VU Amsterdam), E. Jonsson (Chalmers University),
E. Djambazova (IPP-BAS), K. Dimitrov (IPP-BAS), S. Ioannidis (FORTH),
E. Kirda (Institue Eurecom), and C. Kruegel (Technical University Vienna)
Email: herbertb@cs.vu.nl, jonsson@chalmers.se, {ead,kpd}@iccs.bas.bg,
sotiris@ics.forth.gr, kirda@eurecom.fr, chris@seclab.tuwien.ac.at

**Abstract**

One of the most critical problems on today's Internet is the lack of security. This gives rise to a plethora of different ways in which the confidentiality, integrity, and availability of data is compromised, and it provides a fertile breeding ground for a thriving underground economy. Thus, when designing a future Internet, it is clear that security must be a first-class design consideration.

To be able to design security for a future Internet, it is first necessary to obtain a thorough understanding of the threats and adversaries that the system must defend against. As a first step toward this understanding, we introduce a number of emerging security threats that need to be considered. These threats were identified by the three working groups that are active in the context of the EU FP7 project FORWARD, and they illuminate different aspects of the threat landscape.

**Keywords**: Emerging Threats, Requirements Analysis, Future Internet

## 1 Introduction

George Santayana famously pointed out that "those who cannot remember the past are condemned to repeat it" [23]. Clearly, this advice is relevant when confronted with the task of inventing and designing a future Internet. In particular, the quote indicates that one should study the current Internet, building upon solutions that proved successful and avoiding those that turned out to be problematic.

One of the most critical problems for the Internet today is the lack of security. In fact, the rise and prevalence of malware [9], online fraud [5], denial of service attacks [17], and spam [20] give clear testimony that the current Internet infrastructure provides no solid basis upon which effective security solutions can be built. The reason for this problem lies in the original design of many core Internet protocols and network architectures, which were created without security considerations in mind. Combined with the known difficulties of retrofitting legacy systems with security, users are left with patchwork solutions that counter individual threats, but fall short when attackers come up with yet another way to break the defenses. Thus, we firmly believe that a future Internet must have security built in from day one. This belief is also shared by the creators of the recent US NSF initiative GENI, which funds research to develop clean-slate redesigns of the Internet [19].

While it is widely recognized that the future Internet requires built-in security mechanisms, the functioning of these mechanisms is less clear. In addition, and equally problematic, the adversarial model is not well understood. That is, while the current problems are known, it is not obvious which threats a future Internet must be armed against. However, identifying the adversarial model and anticipating emerging threats is the first step that is necessary to build a secure, future Internet. Only when the community has a solid understanding of the threats that the future Internet might face, appropriate countermeasures can be designed.

In this paper, we take a first step towards establishing an adversarial model for the future Internet. To this end, we introduce a number of emerging threats that should be taken into account when developing the designs of a future Internet. These threats and challenges were identified through the combined efforts of three working groups that operate within the framework of the EU F7 project FORWARD. The three project working groups look at different aspects of the ways in which we expect that a future Internet will be used, and threats that endanger these use cases. Thus, the threats that are identified by the groups are general developments that any future Internet design must likely address.

The first working group focuses on *malware & fraud*. This reason for selecting this problem area stems from the belief that the future Internet will continue to be used for e-commerce and financial transactions, even increasingly so. Combined with the fact that we witness the formation of a thriving underground economy, we expect significant opportunities for malicious code and scam operations that abuse victims for financial profit. The second working group focuses on threats that are due to the transformation of regular networks into *smart environments*. With this, we mean the increasing number of small devices that have increasing computing power. These small devices will soon be present everywhere and connected permanently to the Internet. Thus, a future Internet design has to address threats that are caused by the explosion of ubiquitous, small devices. Finally, we predict that the Internet will be increasingly used to control *critical systems*. While this often refers to the control of critical infrastructure and industrial plants, this is only part of the picture. The reason is that also software systems can carry out mission-critical tasks, and loss or disruption of connectivity or software failures can cause severe economic damage or threaten lives. Thus, a future Internet has to address the threats that come along with the increased reliance on a functioning, underlying network infrastructure.

In the following sections, we discuss key threats that have been identified by the three working groups and argue that they are relevant for the design of a secure, future Internet.

# 2 Threats: Malware & Fraud

In recent years, we have witnessed a dramatic change in the goals and modes of operation of malicious hackers. As hackers realized the potential monetary gains associated with Internet fraud, there has been a shift from "hacking for fun" (or bragging rights and celebrity within and outside the hacker community) to "hacking for profit." This shift has been leveraged and supported by more traditional crime organizations, which eventually realized the potential of the Internet for their endeavors.

The integration of sophisticated computer attacks with well-established fraud mechanisms devised by organized crime has resulted in an underground economy that trades compromised hosts, personal information, and services in a way similar to other legitimate economies. This expanding underground economy makes it possible to significantly increase the scale of the frauds carried out on the Internet and allows criminals to reach millions of potential victims. Also, criminals are taking full advantage of sophisticated mechanisms, such as the service bots used on IRC channels to automatically verify stolen credit card numbers, the use of e-casinos to launder money, and the use of fast-flux networks to create attack-resilient services.

Many of the problems related to malware and fraud are rooted in the poor security of the end hosts and the gullibility of end users. However, some of the attacks take full advantage of the poor security of the Internet infrastructure and its protocols. In those cases, as discussed below, a future Internet design can play an important role to mitigate the threats.

## 2.1 Network Infrastructure Security

The network infrastructure (such as switches and routers) and network protocols (such as IP and routing protocols) provide the means for remote hosts to exchange packets. Unfortunately, these protocols and the infrastructure offer basically no security. As a result, packets between hosts can be easily intercepted and altered. Incorrect routes can be injected, a practice that is frequently used to send untraceable spam or to cause denial of service attacks. Moreover, there is no accountability, allowing a sender to spoof the origin (source address) of packets.

Of course, there are many different techniques that address parts of the problem. For example, cryptographic solutions can help to maintain the confidentiality and integrity of application data, security enhancements to BGP (such as sBGP) were proposed to defend against the injection of fake routes, and filtering at the network perimeter can help against packet spoofing. However, the problem is that these solutions are only add-on patches that address a small aspect of the problem. In addition, many solutions are not widely-deployed. For example, sBGP is barely used and the routing infrastructure remains vulnerable. Often, the reason for not upgrading is the high cost of changes to the infrastructure, and little incentive because abuse and fraud often do not affect the provider of the infrastructure but end users.

Higher-level protocols must be able to rely on a secure and robust underlying architecture. Thus, we claim that it is imperative to build accountability into a future Internet. Accountability must ensure that packets and actions can be reliably connected to the originating source. In this fashion, it is possible to identify culprits, block their activity, and also use recorded information as evidence for legal actions. This is an important step to prevent criminal activity that can rely on the anonymity and untraceability that today's infrastructure provides.

## 2.2 Domain Registrars and Fast-Flux

In addition to the basic Internet routing and network infrastructure, many Internet applications rely on the existence and correctness of the domain name service (DNS). DNS is a protocol that resolves human-readable names (URIs) to IP addresses. While crucial for the proper functioning of the network, DNS is unfortunately very insecure. This has led to a number of problems in the past, where attackers altered the mapping between a domain name and the corresponding IP address(es), redirecting legitimate traffic to malicious hosts or launching man-in-the-middle attacks. Such cache poisoning attacks were less frequent over the last years, but experienced a dramatic renaissance with the discovery of the Kaminsky bug[1].

Besides the security problems of DNS itself, the protocol is also heavily abused for fraud and botnet operations. One particular problem is fast-flux, and the role that domain registrars play in this scheme. Fast-flux is a DNS technique that is often used by criminals to hide scam and malware delivery sites behind an changing network of compromised hosts that act as proxies. More precisely, with fast-flux, criminals register a domain name and frequently change the mapping between this domain name and multiple IP addresses. These IP addresses typically belong to bots that act as proxies. When a victim resolves the address that belongs to a malicious domain name, he will contact one of the bots. This bot will then forward the traffic to the actual server that is controlled by the criminal (this server is often referred to as the mothership). By inserting an additional layer of proxies between the victim and the mothership, it becomes much harder to locate a malicious host and take it down. Instead, only the IPs of compromised machines (bots) are visible, and these IPs can be quickly changed when a bot-infected machine is cleaned up. In addition to hosting fast-flux networks, criminals also frequently register hundreds or thousands of throw-away domains that are used to advertise scam sites in spam mails. In both cases, domain registrars play an important role. With fast-flux, attackers have to frequently change the IP addresses that belong to a domain name, an activity that is highly unusual and suspicious in normal operations. When criminals manage to register many domains, it is also evident that registrars do not put sufficient checks into place to prevent abuse.

Clearly, the Internet requires a lookup service that allows hosts to securely locate remote resources. However, the current DNS approach is not working. Thus, the future Internet requires a robust and secure lookup service that is resilient to the injection of invalid mappings. This would prevent cache poisoning attacks and fast-flux networks, since it would no longer be possible to map a domain name to a host that is neither authorized nor aware of this mapping. Moreover, the mechanism to register names has to be overhauled.

## 2.3 Novel Communication Protocols

Spam originally referred to the canned meat product sold by the Hormel Foods Corporation. Since then, many other uses of the term have emerged. With respect to security, the term spam denotes unsolicited messages sent to a large number of Internet users with the aim of luring them to specific web sites. These sites are often used by the spammers to sell products. These products may be illegal, or difficult to obtain via normal means (e.g., a popular product that is often endorsed by spam messages is Viagra – a pill that needs to be prescribed by a doctor). However, in some cases, miscreants also attempt to install malicious applications on the computers of the victims that visit the web site.

Although Internet spam is typically identified as being unsolicited, "bulk," or "junk" e-mail messages, spammers have been finding new ways of delivering spam messages to Internet users. By diversifying the delivery vectors with the aim of reaching as many Internet users as possible, spammers wish to increase

[1] A bug in the design of DNS, which was discovered by Dan Kaminsky and which allows attackers to inject incorrect mappings into the cache of DNS resolvers.

their chances of luring visitors to the web sites that they advertise. Another reason for the diversification and discovery of novel spam delivery protocols is that e-mail-based spam filters have improved considerably over the last years. Hence, the chance of an e-mail-based spam reaching a user is lower now than spam that is being sent over other protocols (such as instant messaging). In fact, spammers have been increasingly making use of instant messaging protocols such as Skype and MSN to send unsolicited chat messages to users. Such messages are usually more difficult for the users to block. Instant messaging products usually do not provide sophisticated filtering mechanisms to identify and eliminate spam messages. For example, it is quite common for ICQ users to receive chat messages in Russian that contain URLs of possibly malicious web sites.

In the future, it will be very important for all communication protocols used on the Internet to enable and support filtering and authentication mechanisms. In fact, spam messages are only possible because the sender cannot be authenticated and identified in many protocols (i.e., the sender can often be forged).

Novel communication protocols are also interesting with respect to security because of the threat of botnets. The term botnet refers to a collection of software robots, or bots, that run autonomously. The term is often associated with malicious software. A miscreant, often called “botmaster” or “botherder”, controls a set of bot-infected machines remotely. The power of botnets lies in the fact that bots are usually running on personal computers of unsuspecting victims. Hence, to an outsider, the Internet IPs of these machines typically appear legitimate and difficult to tag as being malicious.

The majority of botnets are controlled via IRC (Internet Relay Chat). That is, the bots log on to specific IRC channels that have been designated by the botmaster and wait and listen for incoming commands. Although IRC is simple and efficient to use from the attacker’s point of view, it is also easy to disrupt and take down. As a result, miscreants have been building botnets that rely on different protocols to implement their command and control infrastructure. For example, P2P protocols such as Gnutella have become popular. A botnet that uses P2P as its command and communication infrastructure is more resilient and more difficult to take down.

In the future, botnets will probably also use other forms of command and control infrastructures that are more stealthy and efficient. For example, there is a chance that future botnets may communicate over social networks, blending the control communication with legitimate user communication. Thus, designers of future communication protocols must be aware of the fact that their protocols will be misused to deliver unsolicited messages to users, and that they can be misused to control botnets. The challenge is to introduce mechanisms to prevent these types of misuse and to add accountability that allows one to identify, block, and remove people (or hosts) that abuse these networks.

# 3 Threats: Smart Environments

The nature of devices connected to the Internet is changing. In addition to traditional computer equipment, smartphones, PDAs, security cameras, and different sorts of other sensors now latch onto a common infrastructure. As a consequence, new security threats arise. In this section, we discuss threats related to smart phones and ubiquitous sensors, all accessing and using a common Internet infrastructure.

## 3.1 Mobile, Smart, and Powerless

Smartphones are mobile phones with PC-like capabilities. In addition to more traditional telephony stacks, calendars, games and address books, they may run any application the user loads onto them. An increasing number of hardware vendors release ever more powerful models, running a variety of applications on a diverse set of operating systems. The application domain of smartphones currently ranges from high-end business markets (targeted, for instance, by RIM’s Blackberry) to consumer and entertainment markets (as targeted by the Apple iPhone, HTC’s implementation of Android, and Nokia 5800 series). In practice, smartphones are used for email, web browsing, centralized calendaring, navigation, music, etc. In addition, phones are frequently used for commercial transactions. Apple and other companies allow applications, music and videos to be purchased online. Payment for goods and services via mobile phone is already provided by Upaid Systems and Black Lab Mobile. In the meantime, companies like Verrus Mobile Technologies, RingGo, Easy Park, NOW! Innovations, Park-Line, mPark and ParkMagic all offer payment-for-parking schemes. Others focus on mass-transit. For instance, Mobile Suica already allows passengers to use their mobile phones to pay for transport on the East Japan Railway Company, the largest passenger railway company in the world.

It is clear that the domain is widening and that it involves money. Analysts predict that in the near future, smartphones will be the primary interface to the Internet and, indeed, the digital world in general. The tsunami of applications engulfing what was previously a *"dumb"* device (a phone) implies that bugs and vulnerabilities to attacks are on the rise as well. Vulnerabilities provide opportunities for attackers, while the increasing importance of smartphones, and the real money involved in the interactions, provide an incentive. Vulnerabilities in the past have allowed attackers to completely take over mobile phones via Bluetooth. Examples included phones from various vendors, such as the Nokia 6310, the Sony Ericsson T68, and the Motorola v80. The process, known as bluebugging, exploited a bug in the Bluetooth implementations. While these are fairly old phones, more recent models, such as the Apple iPhone have also shown to be susceptible to remote exploits

So what is new? Surely, we have seen all of this before in the world of PCs and so the threat of the future is the threat of today? Unfortunately not. Yes, smartphones are just like PCs in processing capacity, range of applications, and vulnerabilities to attacks. But they are very unlike PCs in other respects, notably power and physical location. These two aspects matter when it comes to security.

First, unlike normal PCs, smartphones run on battery power. Power in mobile phones is an extremely scarce resource. For instance, one of the main points of criticism against Apple's iPhone 3G is its short battery life [18]. Software developers bend over backward to make core code run fast on phones, because every cycle consumes power, and every Joule is precious. As a consequence, many of the security solutions that work for desktop PCs do not suit smartphones, simply because they are too heavyweight. File scanning, taint analysis, system call monitoring all consume battery power. Battery life sells phones, and consumer hate recharging. The likely result is that both vendors and consumers will trade security for battery life.

Second, unlike traditional computers, phones go everywhere we go. Attacks may come from sources that are extremely local (e.g., via bluetooth). A person with a laptop or another smartphone that happens to be in the same room could be the source of an attack. That means that security solutions based on in-network scanning are insufficient: they will never even see the bytes that are used to take over the phone, steal information, and plunder the bank account. Worse, phones are small devices, and we do not always keep an eye on them. We may leave them on the beach when we go for a swim, slip them in a coat or shopping bag, forget them on our desks, etc. Theft of a phone is much easier than theft of a desktop PC or even a laptop. Moreover, attackers could "borrow" the phone, copy data from it, install backdoors, etc. This is an important difference compared with the PC you have sitting on your desk.

What would it buy an attacker to steal a phone (and perhaps return it later)? Almost always, having physical access to a device opens up a wide range of attack options, for example, hardware attacks [6]. Attackers may use hardware debugging equipment to snoop on data traveling from and to memory, read or write keys, etc. Direct loss of private data may be an immediate result. However, another and perhaps more insidious threat is when the phone is returned to the owner with a backdoor that allow attackers to gather information for a long period of time.

Is this practical? Let us have another look at the example of bluebugging; the faulty implementation that made the earliest bluetooth phone vulnerable to remote exploits was fixed fairly quickly. However, phones could still be compromised. The only thing that was needed was that the bluebugger talked 'the victim into handing over the phone, which the bluebugger manipulates to set up a backdoor attack and then hands back' [14].

We stress that the trends are not working in our favor. On one hand, mobile phones are an increasingly attractive target for attackers. On the other hand, because of power limitations and physical exposure to hostile environments, phones are inherently more difficult to protect than traditional computers. In a future Internet, it is imperative that solutions are found to protect mobile devices that carry valuable data. Existing paradigms, based on in-network scanning and/or traditional anti-virus software cannot be simply ported to mobile phones.

## 3.2 What were you doing in the Red Light District?

Last month's parking sensors indicate that your car was parked on a Saturday in the Red Light District in Amsterdam. Your mobile phone was seen wandering up and down the seedier streets of Amsterdam for a while before stopping in one place for a good 20 minutes. When finally roaming again, it made a beeline to the Warmoesstraat where several bluetooth devices interacted with it during a 60-minutes interval at a location known to accommodate a (Dutch-style) coffee shop.

This is not (yet another) story about violation of privacy, or "Big Brother is watching you having a day of R&R in Holland." Although the threat of privacy violation cannot be overstated, it is hardly new and we will not address it in this paper. While it may have been painful for you to be spotted in compromising locations on a day that you claimed to be going for a walk on the beach, we worry about something else entirely: were you really there? *Or, to what extent can we trust sensor data connected on to the future Internet?*

**Buggy devices yield wrong information.** A common example of a buggy sensor is the barcode reader in the supermarket that double-charges a product. Similarly, supposedly-disabled RFID tags on clothing or other products frequently set off alarms when a client exits the shop, often creating embarrassment. It has happened to most of us, and these are the simplest examples of what we call buggy devices. While barcode readers result in modest overcharging, other types of buggy sensors are more serious. In some places, automated parking sensors are used to identify your car, and cars are tracked on motorways to pay-as-you-go. In several countries, public transport is accessed using a smart card or phone, directly linkable to you. With the increase of the number of sensors increases the probability of false reports.

Similarly, as Internet Protocol (IP) telephony becomes ever more popular, buggy telephones (or buggy call protocols) may dial arbitrary numbers without the user's consent. Many companies base their revenue on received phone calls as they have special contracts with telephone companies. A buggy device may end up benefiting the company and harming an unaware user. Apart from the financial aspect of the issue, a more important aspect arises. If the buggy telephone can call anywhere, it means it can place calls to people you might not want or should call. For example, an employee in a large company owns a buggy telephone. What will happen if his company finds out that he calls phone numbers belonging to a competitor? It is unclear whether he will be able to convince their employers that something went wrong.

**Sensors or databases can be hacked.** Perhaps the most significant threat is that data can be falsified. Most information is stored in centrally-controlled databases. This is a fact that can be hardly changed, as distributed control costs both in terms of resources and manpower. Take, for example, your telephone company. Whether you are at home or roaming away from it, all your phone calls are logged by your provider in a huge, centralized repository. A sophisticated hacker or a person with knowledge of a company's internals can alter all the information about your phone calls. Such actions can have significant repercussions in one's life. For example, one may be framed by being linked via telephone records to criminals. Mobile tracking, shopping activities, car parking, they all can be manipulated to create a virtual clone of yourself with unknown implications. In the digital world, where everything is connected via the Internet, planting of "evidence" is both easier to do and harder to detect.

Apart from relatively harmless issues, the implications of sensor and database hacking can be very harmful, even lethal. We have sensors and databases in hospitals, in banks, in organizations, in police and justice departments. Given the proper motive, a person with access to such information can easily incriminate or even kill someone. An example is that of medical records. A patient enters a hospital after a serious car accident, far away from their home. The hospital doctors consult the e-record to check the patient's medical history. The original medical record indicates a blood allergy to specific drugs. An altered medical record hides this information. Wrong treatment to the patient can be lethal, and it is very probable to happen as the medical history can be totally changed.

**Legal implications.** In our opinion, the untrustworthiness of sensor data creates a legal void that needs to be filled urgently. All easy solutions are wrong. Admitting sensor data as reliable proof makes little sense if the data may be unreliable, and especially if the data can be altered. But clearly, there *is* a link between the sensor data and reality in most cases, so we cannot dismiss sensor-based evidence either. Current legislation already looks at some of these issues (e.g., the legal status of footage from a surveillance camera), but the scale at which a multitude of sensors will track persons and objects in the future is such that re-thinking legal implications is important. The issues that need to be taken into account ranges from evidence based on individual sensors to collections of sensors, and incorporates both agreements between sensors and anomalies. Arguably the most important question that needs to be answered is how people who are accused on the basis of sensor data can defend themselves. Here, a future Internet design is again asked to provide data provenance, a mechanism that allows one to securely track the origin of data, as well as modifications to this data. While this might not prevent all misuse, it allows one to go back and track those responsible for malicious activity.

# 4 Threats: Critical Systems

The Internet is a communication environment that has become an essential part of our everyday life, in the same fashion as the electricity or the telephony network have become essential over the last one hundred years. The more products and services we access through it, the more dependent we become on its functionality and availability. Indeed, the "functionality" of the Internet has outgrew its initial goal, to transfer information between distant sites; we now expect it to transfer trust and to operate in new critical areas.

The complexity and interdependence of modern inter-networks make identifying threats that impede their proper operation extremely difficult. Covering all possible individual threats would be impractical in the context of this paper, so we will focus on the probable areas of occurrence of threats and outline the most sensitive points of emerging threats that have the potential to endanger critical systems (CS). We believe that covering such threat areas will assist researchers, manufacturers, and the community as a whole, to focus their efforts.

By design, the Internet is not suited for critical applications, since it was built to provide a best-effort packet relay service. Having outgrown its intended purpose, the Internet is now being used by critical applications, and, as a result, it has itself turned into a critical system. The popularity and mass expansion of the Internet encourages its use even in critical applications where it was not previously used. The Internet technologies work along with the specialized technologies for process control and if this tendency reaches the safety-critical systems in their main functionality, it could be a serious threat.

Another area of potential threats is the connection of the Internet to critical infrastructures (CIs). Many CIs (e.g., banks, power stations, industrial complexes, telephony networks, etc.) use the Internet for their communication needs. The problems of Internet's insecurity and vulnerability increase when considering the scale, complexity, connectivity, and interdependency of such critical infrastructures.

## 4.1 Internet and Critical Applications

Penetration of the Internet in the economic sectors of our society poses considerable risks. So, what are the emerging threats from such use? Answering this type of questions will assist us in understanding the relationship between the Internet and critical applications.

Traditionally, in the fields of hazardous industrial processes and safety-critical systems for process control, specialized real-time and fault-tolerant computer systems and communications are used with guaranteed dependability and safety.

On the other hand, according to the integrated vision on dependability and security [3], any undesired event for a system (external or internal) can be regarded as a threat. For example, if an off-the-shelf system is put in a critical application, there is high probability that a fault occurrence may lead to system failure with unpredictable consequences. Fault-tolerant systems preserve their dependabilty and security even when unreliable components and subsystems are used for their design. Unfortunately, such guarantees are not present in the Internet. In particular, the use of off-the-shelf components in the Internet is a common practice. Hence, when used in critical applications, problems can arise.

A possible chain of threat-causing events (threat pathology) could be as follows: (i) The use of the Internet into a critical application induces a gap between the application requirements and the capabilities of the involved Internet components, (ii) this deficiency effectively decreases the robustness of the components and (iii) in turn leads to increased risk and vulnerabilities. Therefore, the increasing and *uncontrolled* use of the Internet in critical systems can be regarded as an area of emerging threats in itself.

To reduce costs and time for design the use of Commercial Off-The-Shelf (COTS) systems and components in critical applications is attractive and will continue. To avoid the above threats, their deficits have to be compensated for.

There are some projects (e.g., [1]) where COTS components are applied to design distributed computer-controlled systems. They are organized using redundancy and design diversity to make the system dependable and secure. Some of the issues addressed in DEAR-COTS are the use of emerging information technologies to cope with heterogeneity issues while providing a dependable user-friendly man-machine interface.

Another direction is to apply COTS and open source standards along with the standards for process control systems. The organizations from industry that develop commercial interface standards work with some military programs to include real-time and fault-tolerance requirements to critical systems [8].

Although it is not applied directly to control critical industrial processes, there are systems that use the Internet technologies and Ethernet to facilitate the access of operators or other maintenance personnel to the control system. That is especially convenient for distributed process control systems where the nodes are distant and their surveillance is done remotely. Process control systems operate with their own protocols and communication channels that are not compliant with the Internet. The control process is independent of any outside activity. The control system is decoupled from other general-purpose systems and networks. The access to the controlled object is left to the controlling system alone. There is a trend, however, to use the flexibility and connectivity of the Internet in so called Internet-based control systems [24], [16]. Special attention must be paid to security and safety of such systems [13].

## 4.2 Internet and Critical Infrastructures

In an overview of incident and vulnerability trends [2], CERT/CC enumerates the most probable threats in the Internet after 2003. Although that prediction may seem old (by Internet standards), it is still relevant for today's state-of-practice. That is, the basic effects of an attack, the attacking methods and opportunities, the major threats are almost the same. The only changes have occurred in their scale. It is interesting to analyze some of the challenges that CERT outlined that relate to critical systems, since these problems persist and are amplified. In particular, some of the conclusions in that overview are now reality. The complexity of the Internet, protocols, and applications is increasing along with our reliance on them. Moreover, critical infrastructures increasingly rely upon the Internet for operation. The attacks become more severe and can lead to monetary/financial loss or even loss or endangerment of human life. The adversary increases the impact by targeting the infrastructure. The opportunities for intrusions grow. The problem is that the fundamental security design difficulties cannot be addressed quickly, and despite the security tools available, the complexity and growth of the network can outpace our capability for defense.

One of the critical issues in the critical infrastructures is the interdependencies among the infrastructures. In [22], the role of ICT technologies in CIs is defined with the term cyber interdependency. An infrastructure has a cyber interdependency if its state depends on information transmitted through the information infrastructure. Virtually all modern infrastructures are dependent on the security of information infrastructure. Given the extensive cyber interdependencies, one has to prepare for the amplification of the security risks.

Being part of the communications infrastructure, the Internet has the same typical vulnerabilities and is prone to similar threats. There are approaches to improve communications' robustness and availability (and those of the Internet in particular) [25]. Using the Eight Ingredients Framework of Communications Infrastructure [21], the vulnerabilities of future networks were studied systematically to determine the vulnerabilities of each of the eight ingredients. The approach relies on vulnerability analysis, since it is recognised that intrinsic weaknesses of communications infrastructure are a finite number and can be identified by the professionals in order to eliminate or mitigate their effects. Combining vulnerability and threat analysis will help improving CI security. Although it is argued that threat analysis is ineffective when the knowledge of possible threats is not certain, the identification of threats helps anticipating challenges in areas of concern that may need more research and development activity.

## 4.3 Discussion and suggestions

As the use of the Internet in critical applications and CIs in the near future is unavoidable, the following main suggestions could be considered:

1. Make a full decoupling of highly-critical systems (hazardous industrial processes, mission-critical tasks) and the Internet
2. Control the introduction of the Internet in other critical areas (legislatively, technically and organizationally).

Some measures related to the second suggestion may be: (a) surveillance, regulation, and coordination between different sectors of CIs in the cases when they are planning the use of the Internet, e.g. like these in

[11]; (b) application of diversity approach when using COTS components [4]; (c) use of compact and trusted applications base; (d) use of integral approach to security (e.g., [12]).

As stated previously, the emerging and future threats are hard to identify, so this may be an unfeasible task. There are many threats that still remain and will be unidentified. At the same time, we can say much more about the stochasic nature of security challenges. The times of occurrence of the challenges that could affect normal operation will rapidly and arbitrarily differ and shall be likelihood uncorellated. Moreover, new challenges will emerge (e.g. new application traffic loads, forms of distributed denial of service (DDoS) attacks, deployment environments, and networking technologies). As a consequence, the affected information infrastructures and delivered network services will change unpredictably. This makes unusable a set of scenarios for resilience prepared in advance and imposes the use of a dynamically reconfigurable and extensible infrastructure with context awareness capabilities. Another challenge is that often a sufficiently sophisticated DDoS attack is indistinguishable from legitimate but enormous traffic (e.g. flash crowd events) [15].

Resilience approach (RA) [4], [10] is a feasible, emergent, and integral approach that can be used for managing the emerging threats. To be successfully implemented, systems and networks have to be designed and built with RA in mind and used in compliance with RA concepts. The main idea is creating a new kind of Information Society Technologies, the resilient technologies [4], that will have and demonstrate an emergent behaviour to successfully withstand and cope with the emergent and arbitrary behaviour of the challenges to normal operations.

Another direction to anticipate the challenges to the future Internet is to develop new architectures that can foster innovation, enhance security and accountability, and accommodate competing interests. Postmodern Internet Architecture [7] aims at developing internetwork layer and auxiliary functionality that operate in policy space as opposed to the current Internet functionality. The overall goal of the project is to make a larger portion of the network design space accessible without sacrificing the economy of scale offered by the nowadays unified Internet.

# 5 Conclusions

In this paper, we have outlined different, emerging threats that the Internet faces today. While we have not proposed any concrete, technical solution to a problem, we believe that it is important to study and understand these threats as a first step to be able to build a better future Internet. As with all engineering projects, requirements analysis is the crucial first step. By analyzing emerging threats from different angles, we hope to provide such a requirements analysis. This can help to ensure that a future Internet does not repeat the mistakes that were made in the design of the current protocols and infrastructure.

# References

[1] DEAR-COTS Project Homepage. `http://dear-cots.di.fc.ul.pt`, 2001.

[2] CERT Report. `http://www.cert.org/annual_rpts/cert_rpt_03.html`, 2003.

[3] A. Avizienis, J.-C. Laprie, B. Randell, and C. Landwehr. Basic Concepts and Taxonomy of Dependable and Secure Computing. In *IEEE Transactions On Dependable And Secure Computing, Vol. 1, No. 1*, pages 11–33, 2004.

[4] A European Network of Excellence. Deliverable D12 Resilience-Building Technologies: State of Knowledge. `http://www.resist-noe.org/events/events.html`, 2006.

[5] D. Anderson, C. Fleizach, S. Savage, and G. Voelker. Spamscatter: Characterizing Internet Scam Hosting Infrastructure. In *Usenix Security Symposium*, 2007.

[6] R. Anderson and M. Kuhn. Low cost attacks on tamper resistant devices. In *Proceedings of the 5th International Workshop on Security Protocols*, pages 125–136, London, UK, 1998. Springer-Verlag.

[7] Bobby Bhattacharjee, Ken Calvert, Jim Griffioen, Neil Spring, and James Sterbenz. Separating Forwarding and Routing, (Postmodern Internet Architecture Project). `http://www.ietf.org/proceedings/07jul/slides/RRG-12/rrg-12.ppt`, 2007.

[8] C. J. Walter, N. Suri, and T. Monaghan. Evaluating COTS Standards for Design of Dependable Systems. In *Proc. of the 2000 Int. Conf. on Dependable Systems and Networks*, pages 87–96, 2000.

[9] D. Dagon, G. Gu, C. Lee, and W. Lee. A Taxonomy of Botnet Structures. In *Annual Computer Security Applications Conference (ACSAC)*, 2007.

[10] David Hutchison, James P.G. Sterbenz. ResiliNets: Multilevel Resilient and Survivable Networking Initiative. `http://www.comp.lancs.ac.uk/resilinets`, 2006.

[11] Homeland Security Advisory Council. Report Of The Critical Infrastructure Task Force. `http://www.dhs.gov/xlibrary/assets/HSAC_CITF_Report_v2.pdf`, 2006.

[12] L. Horacek. Protection on Demand, Information Security that Works for you, The IBM Approach to Security Protection from the Core to the Perimeter. IDC Security Roadshow Sofia, April-12-2007, 2007.

[13] L. Yang and Yang S.H. A framework of security and safety checking for internet-based control systems. In *Int. Journal of Information and Computer Security, Vol. 1, No. 1/2*, pages 185–200, 2007.

[14] G. Legg. The Bluejacking, Bluesnarfing, Bluebugging Blues: Bluetooth Faces Perception of Vulnerability. TechOnline `http://www.wirelessnetdesignline.com/showArticle.jhtml?articleID=192200279`, April 2005.

[15] Linlin Xie, Paul Smith, Mark Banfield, Helmut Leopold, James P.G. Sterbenz, and David Hutchison. Towards Resilient Networks using Programmable Networking Technologies. `http://www.ittc.ku.edu/resilinets/papers/Xie-Smith-Banfield-Leopold-Sterbenz-Hutchison-2005.pdf`, 2006.

[16] M. Coccoli and A. Boccalatte. Future Directions of Internet-based Control Systems. In *Journal of Computing and Information Technology - CIT 10, 2*, pages 115–124, 2002.

[17] D. Moore, G. Voelker, and S. Savage. Inferring Internet Denial of Service Activity. In *Usenix Security Symposium*, 2001.

[18] W. Mossberg. Newer, Faster, Cheaper: iPhone 3G. Wall Street Journal, July 2008.

[19] National Science Foundation (NSF). Press Release 07-057: Three Wishes for a Future Internet? `http://www.nsf.gov/news/news_summ.jsp?cntn_id=109589`, 2007.

[20] A. Ramachandran and N. Feamster. Understanding the Network-level Behavior of Spammers. In *ACM SIGCOMM*, 2006.

[21] Network Reliability and Interoperability Council VI. Homeland Security Physical Security (Focus Group 1A) Final Report, Issue 3. `www.nric.org/fg/nricvifg.html`, Dec. 2003.

[22] S. Rinaldi and J. Peerenboom and T. Kelly. Identifying, Understanding, and Analyzing Critical Infrastructure Interdependencies. In *IEEE Control Systems Magazine*, pages 11–25, 2001.

[23] G. Santayana. *Life of Reason, Reason in Common Sense.* Scribner's, 1905.

[24] S.H. Yang, X. Chen, J.L. Alty. Design issues and implementation of internet-based process control systems. In *Control Engineering Practice, Volume 11, Number 6*, pages 709–720, 2003.

[25] Alcatel-Lucent Technologies. The ARECI Study, Final Report, Availability and Robustness of Electronic Communications Infrastructures. PSC Europe PSCE/RD/024, March 2007.