SEVENTH FRAMEWORK PROGRAMME

THE

RED BOOK

A Roadmap for Systems Security Research

Managing Threats and Vulnerabilities in the Future Internet

SEVENTH FRAMEWORK PROGRAMME

Information & Communication Technologies
Trustworthy ICT

NETWORK OF EXCELLENCE
Grant Agreement No. 257007

A European Network of Excellence in Managing Threats and Vulnerabilities
in the Future Internet: *Europe for the World*

The Red Book:
A Roadmap for Systems Security Research

Abstract: The Red Book presents a roadmap in the area of systems security, as prepared by the SysSec consortium and its constituency in the first half of 2013.

Contractual Date of Delivery	August 2013
Actual Date of Delivery	August 2013
Dissemination Level	Public
Editor	Evangelos Markatos, Davide Balzarotti
Contributors	All SysSec partners
Quality Assurance	M. Almgren, E. Athanasopoulos, H. Bos, D. Balzarotti, L. Cavallaro, S. Ioannidis, M. Lindorfer, F. Maggi, E. Markatos, F. Moradi, C. Platzer, I. Polakis, M. Polychronakis, A. Slowinska, P. Tsigas, S. Zanero

The SysSec consortium consists of:

FORTH-ICS	Coordinator	Greece
Politecnico Di Milano	Principal Contractor	Italy
Vrije Universiteit Amsterdam	Principal Contractor	The Netherlands
Institut Eurécom	Principal Contractor	France
IICT-BAS	Principal Contractor	Bulgaria
Technical University of Vienna	Principal Contractor	Austria
Chalmers University	Principal Contractor	Sweden
TUBITAK-BILGEM	Principal Contractor	Turkey

The research leading to these results has received funding from the European Union Seventh Framework Programme (FP7/2007-2013) under Grant Agreement Number 257007. This work would not have been possible without the contributions of the SysSec Working Groups, the SysSec Advisory Board, and the broader SysSec community in general. We deeply thank them all.

www.syssec-project.eu

SYSSEC TASK FORCE for the ROADMAP on SYSTEMS SECURITY RESEARCH

CO-CHAIRS

Evangelos Markatos
SysSec Project Manager
Foundation for Research and Technology - Hellas

Davide Balzarotti
SysSec WP4 Leader
Eurecom

MEMBERS

Elias Athanasopoulos
Columbia University
Federico Maggi
Politecnico di Milano
Asia Slowinska
Vrije Universiteit

Lorenzo Cavallaro
Royal Holloway University of London
Michalis Polychronakis
Columbia University and FORTH
Iason Polakis
FORTH and University of Crete

Magnus Almgren
Chalmers
Sotiris Ioannidis
FORTH
Philippas Tsigas
Chalmers

Herbert Bos
Vrije Universiteit
Christian Platzer
TUV
Stefano Zanero
Politecnino di Milano

CONTRIBUTORS

Dennis Andriesse
Vrije Universiteit
Farnaz Moradi
Chalmers University
Simin Nadjm-Tehrani
Linköping University

Martina Lindorfer
TU Vienna
Zlatogor Minchev
Bulgarian Academy of Sciences
Christian Rossow
Vrije Universiteit

Preface

After the completion of its second year of operation, the SysSec Network of Excellence produced this "Red Book of Cybersecurity" to serve as a Roadmap in the area of Systems Security. To realize this book, SysSec put together a "Task Force" of top-level young researchers in the area steered by the advice of SysSec WorkPackage Leaders. The Task Force had vibrant consultations (i) with the Working Groups of SysSec, (ii) with the Associated members of SysSec, and (iii) with the broader Systems Security Community. Capturing their feedback in an on-line questionnaire and in forward-looking "what if" questions, the Task Force was able to distill their knowledge, their concerns, and their vision for the future.

The result of this consultation has been captured in this Red Book which we hope will serve as a *Road Map* of Systems Security Research and as an *advisory document for policy makers* and researchers who would like to have an impact on the Security of the Future Internet.

How to Read this Book

Policy Makers may want to focus on Chapter 1 at page 3 which provides a short Executive Summary of the book and on Chapter 14 in page 103 which describes Grand Challenge Research Problems in the area which can be solved only with the collaboration of several Research Organizations and the support of leading funding Agencies. Related work may be found in the second part of the book in page 107, which provides a good overview of other Research Roadmaps from Europe and from the States.

Young Researchers who are interested in doing a Ph.D. in systems security should read the first part of the book, and especially the final section of each chapter, which describes problems that are appropriate to be solved within the context of a Ph.D. thesis.

Experienced Researchers may want to focus on the first part of the book, which provides an in-depth treatment of various research problems and in Chapter 14 in page 103, which describes Grand Challenge Research Problems in the area.

Journalists may want to focus on sections *.2 and *.3 of the first part, which paint a picture of the average and worst-case consequences of the emerging threats studied.

All should read Chapter 2 in page 7, which lists the identified threats, assets and security domains.

Contents

1 Executive Summary

Based on published results, it is considered larger than the black market of marijuana, heroin, and cocaine combined [13]. Its size was recently estimated to exceed one trillion dollars [243]. It adversely affected more then 88% of Europeans last year [53]. What is it? It is the Global Market of Cyber Crime. As we embraced the convenience and effectiveness of the Internet into our lives, homes, retirement plans, and even wallets, we also opened the door to a new breed of attackers determined to gain profit from this wonderful new cyberworld. Motivated by fun, profit, and even political motives, cyberattackers have now impacted, or threaten to impact, most realms of our lives.

Understanding the dangers we have subjected ourselves to and predicting the threats that are going to materialize, is one of the major tasks of the SysSec Network of Excellence. A four-year project, SysSec has mobilized the top cybersecurity researchers in Europe and challenged them to *think ahead*, *think disruptively*, and finally *predict* what should be the important emerging research areas in cyber security and privacy. This book summarizes the *Emerging Threats* identified during the third year of the project and proposes *Grand Challenges* that, if addressed, will significantly boost the safety and security of the Internet for the years to come.

1.1 Emerging Threats

SysSec, along with its constituency, has identified a number of research issues on which we should focus our efforts. The issues are organized in two groups: *Threats*, which correspond to dangers that may exploit vulnerabilities and cause harm, and *Domains*, which correspond to emerging application areas made possible (i) by advancements in technology, and (ii) by major shifts in society.

The major threats identified are:

> Malware, Botnets, Insider Threats, Targeted Attacks - Advanced Persistent Threats, Web Vulnerabilities, Software Vulnerabilities, SPAM, Malicious Hardware, Data Breaches, Social Engineering - Phishing, Passive/Active Eavesdropping, On-line behavior tracking, and Spoofing - Impersonation.

The major domains identified are:

> Social Networks, On-line Games, e-commerce, e-banking, Sensors - Drones, Embedded Systems, SmartEnvironments, Legacy Systems, Critical Infrastructures, Mobile Systems, Wireless Networks, Implantable Devices, and The Cloud.

The Important Ones

We have asked our constituency to select the threats and domains that they feel are most important of all. The three most important threats selected were:

- Malware
- Targeted Attacks
- Social Engineering - Phishing

The three most important domains selected were:

- Mobile Devices
- Social Networks
- Critical Infrastructures

1.2 Grand Challenges

In addition to emerging threats, SysSec has identified a few *grand challenge* problems. Solving them will be a major step towards creating a trusted and safe cyberspace. These challenges include:

- **No Device Should Be Compromisable:** Develop the necessary hardware and software support to make it impossible for attackers to compromise a computer or communication device for that matter, including smartphones and tablets.
- **Give Users Control Over Their Data:** Provide the necessary mechanisms so that users
  1. will be able to *know which data they have created* (such as text, photos, videos, cookies, web requests, etc.),
  2. will be able to *know what data they have given to third parties* (such as text, photos, cookies, web requests, IP addresses, etc.)
  3. will have the capability to *refuse disclosure of some data* (such as cookies and IP addresses) and still expect a decent level of service,

4. will have the capability to *delete their own data which they have created* (both from the local storage as well as from the cloud), and
5. will, under an appropriate legal framework, have the ability to ask past recipients of their data to erase them as well.

- **Provide Private Moments in Public Places:** Enable users to have private communication in the public areas of the cyberspace. Consider the following analogy: The fact that people are having dinner in a public restaurant does not mean that their conversation could be recorded by the manager of the restaurant, and later made available without their explicit consent. Similarly, the fact that people are communicating in the cyberspace does not imply that parts of their communication can be recorded and used later through means outside their control. We propose to develop mechanisms that will enable people to have a reasonable expectation of privacy in what can be considered a public venue in the cyberspace.

- **Develop Compromise-Tolerant Systems:** Provide adequate security levels even if components of the system have been compromised. It is reasonable to expect that not all attacks will be detected and successfully mitigated. Human errors, software errors, hardware errors, and insufficient protection mechanisms will allow some attacks to go through successfully. This implies that some systems, or components of systems will be compromised, and this may go undetected for a long period of time. Given such an environment, we should develop systems that will be able to provide decent security guarantees even if some of their components are compromised.

2 Introduction

Cyberspace penetration in our everyday lives has reached unprecedented levels. This has left us facing the challenge of understanding, monitoring, and mitigating the security and privacy implications of this wonderfully surprising and inspiringly interesting new medium. In this chapter we describe this challenge from four different points of view: (i) the new *threats* that cyberspace has made possible, (ii) the *assets* that we care about, (iii) the *domains* that have risen, and (iv) the *horizontal research directions* which need to be supported.

2.1 The Cybersecurity Landscape

Throughout this book we treat the notion of security along four dimensions:

- **Threats - Vulnerabilities**. Vulnerabilities and threats are usually artifacts in the on-line world that the attackers may exploit in order to cause harm to their victims. For example, an attacker may exploit a buffer overflow in order to compromise a computer and use it to send SPAM. In the physical world, threats/vulnerabilities could include an open window in a house, an unlocked door, etc. Several threat definitions include the attackers themselves (such as in "*insider threats*" or in "*advanced persistent threats*") in the category of threats as well. We plan to use the same approach.

- **Assets**. Assets are resources that entities (such as people and organizations) hold on to and value. Assets may include money, data, human rights, etc. Cyberspace may impact the same assets as the physical world, but probably in entirely new ways. For example, although privacy has been an asset in the physical world for several years, in cyberspace it may take on a whole new spin, as (i) the data gathered, (ii) the entities gathering such data, and (iii) the potential uses of such gathered data are of unprecedented scale.

- **Domains**. Attackers may stage their attack in a particular domain setting. For example, the domain of social networks could be used by attackers

who want to get at the personal data of particular people. A domain may be thought of as a restriction to (some of) the threats in such a way that they can be studied and/or better solved. Solutions developed for one domain may not necessarily hold on other domain. For example, the domain of implantable devices may restrict the types of DoS attacks which can be made on such devices and may require solutions that can not be applied in large scale servers. The purpose of these domains is not usually to provide security, although without security their provision could be useless. For example, the purpose of the domain of on-line banking is to provide banking services, much as the traditional banking sector has been doing in the physical world for hundreds of years.

- **Horizontal Research Areas**. Finally, we have identified *Horizontal Research* areas - that is, areas that apply to several, if not most, aspects of security. For example, *measuring* security is a horizontal research area whose results may apply to many individual domains and threats.

One may choose to approach the problem of security and privacy from any of the above first dimensions. For example, one might start with a threat, such as a buffer overflow, and explain the types of attacks that can be made possible, the types of assets that can be compromised, and the kind of domains in which such attacks would materialize. As another example, one might start with the assets that seem important and explain how the different domains may set the stage for an attack on these assets and how an attacker may exploit domain-specific vulnerabilities or use threats to materialize such attacks.

We feel, however, that in the recent history of cybersecurity and privacy all above dimensions have been used, so that each individual problem is described from the most convenient and easiest-to-understand dimension. In this work we follow a similar approach and categorize the important aspects of cybersecurity and privacy along these dimensions, so that we are able to illustrate the concepts from the best point of view.

2.2 Mapping the Threats We Fear

The evolution of cyberspace, which triggered an explosion in innovation and novel applications, has offered cyberattackers a wide variety of threats and vulnerabilities that can be used to compromise people's security and privacy. Such threats include:

- **Malware** has been traditionally used as the main vehicle to carry malicious activities for several decades now. Initially materialized as "computer viruses" and originally spread through "floppy disks," malware is still going strong, compromising computers at the speed of the Internet.

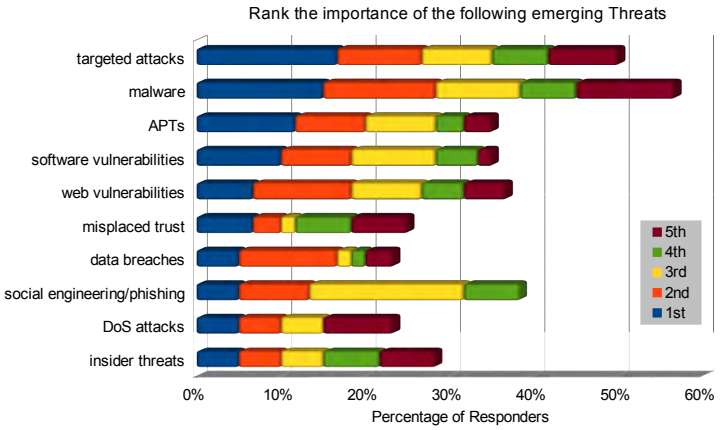

- **Botnets**. To make sure that compromised computers have significant fire-power, attackers organize them into networks, called in the colorful language of computers, *botnets*. Botnets have developed mechanisms to evade detection, to "survive" in case of attack, and to dynamically organize even if several of their members are taken down.

- **Insider Threats**. We should never underestimate the damage which can be done by insiders who have access to large amounts of data and restricted digital systems. Such insiders in the past had been responsible for major data leaks and significant financial fraud. As more and more data are accumulated on-line we expect the threat that such insiders can pose to be a major cause of concern.

- **Targeted Attacks - Advanced Persistent Threats**. Over the past few years, we have seen an increasing number of attacks aimed at sectors of the Industry, such as SCADA systems, or at countries themselves. Such attacks have used state of the art malware and remained undetected, for several weeks after their initial infection. Backed by significant human resources and generously financed, such asymmetric threats represent a major challenge for researchers.

- **Web Vulnerabilities**. The proliferation of Web2.0 coupled with a wide variety of emerging user activities ranging from web banking to on-line gaming gave rise to a set of vulnerabilities which can be exploited only through web browsers.

- The Traditional Stronghold of the attackers, **Software Vulnerabilities**, are being used to exploit systems and will probably continue to be so used in the near future.

- Although considered rather old-fashioned, **DoS attacks** have started to re-appear as the easiest form of on-line pressure.
- **SPAM**. The domain of unwanted messages, although appearing to be in decline, may redefine itself through targeted domain-specific SPAM that is difficult to detect and may appear very similar to the interesting messages received by the intended recipients.
- **Malicious Hardware**. Attackers have been traditionally targeting vulnerable software. However, it has become apparent that attacks on hardware may be more effective, and more difficult to detect. Although such attacks require access to the hardware design/development process, recent examples have demonstrated that such access is not impossible.
- **Data Breaches**. Recently, an increasing percentage of activities are performed (and recorded by various stakeholders) on-line. Some of them are of great importance, such as on-line banking and interaction with the local State or Government, and may represent a high-value target for potential attackers. Unfortunately, opting out of such data collection is not a possibility for ordinary citizens. The data will be collected and will be stored in accordance with the relevant laws. Potential leaks of such data may put the population of entire countries at major risk.
- **Social Engineering - Phishing**. Social engineering has been one of the oldest methods used by attackers and will probably continue to be popular in the future. As technology advances faster than people can understand[1], attackers have the opportunity to exploit the little-understood trust relationships of the ever-changing environment.
- One of the oldest attacks, **Passive/Active Eavesdropping**, popularized by *man-in-the-middle-attacks*, still seems to be very popular. The widespread use of wireless communications, the recent popularity of proxy-based infrastructures, and the availability of technology to retain data for several months by most ISPs, make eavesdropping easier than ever.
- **On-line behavior tracking**. Highly-desired by online advertisers and popularized by the ubiquitous *cookies*, behavior tracking has reached the point where it can track what people read, where they go, what they buy, and even when they change the TV channel.
- **Spoofing - Impersonation.** Several transactions on the Internet do not require strong authentication. For example, even the sender of an IP packet may easily be spoofed. This spoofing, or even impersonation

[1] A notion colorfully termed *future shock* in "The For Ever War" [205].

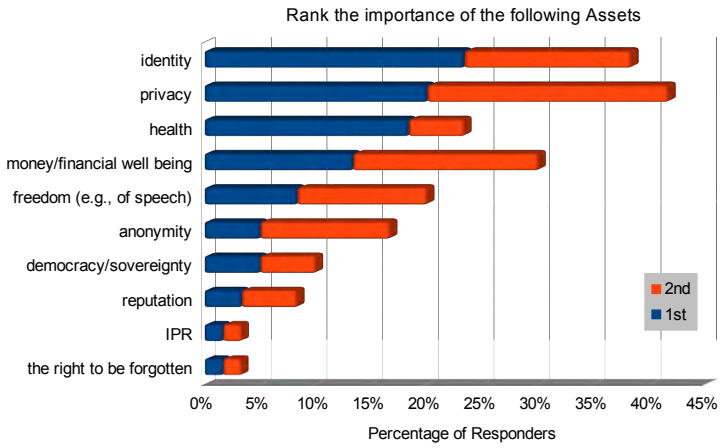

may happen at several different communication levels and may lead to significant damage to both individuals and organizations alike.

2.3 Listing the Assets We Value

We have identified several assets that could be the target of attackers. Assets of particular importance to organizations include money and data. Assets of particular importance to people may also include health, life, human rights, and so on.[2] Thus, we have listed the following Assets we feel important:

- **Life**. The dearest of all an individual's assets, life may be the target of cyberattackers. Indeed, attacks on medical systems, transportation systems, or systems dealing with emergency response may easily lead to massive loss of life.

- **Health**. The increasing use of IT in healthcare may also increase the possibilities of attackers to hurt the health of individuals.

- **The Environment**. Only recently receiving proper attention, the environment is absolutely necessary for the survival and advance of human beings. Threats to the environment, possibly through large-scale pollution and raging fires triggered by cyberattacks, may have devastating consequences for the affected communities.

- **Privacy**. Recognized as a Human Right by the United Nations, privacy will probably be challenged the most in the cyberspace of the near

[2] We chose to separate and explicitly mention "The Right to be Forgotten" as a single entity in order to understand how the general setting impacts this new Right.

future, where each and every action people take on-line will probably be recorded in some database beyond their control.

- **Freedom - Freedom of speech**. Recognized as a Fundamental Human Right, freedom is greatly challenged by the high-end technology in general, and surveillance mechanisms in particular.

- **Democracy, Sovereignty**. People who live in democratic societies usually take it for granted and may underestimate the threats that it may be subject to. However, high technology, including digital storage and communications can be used to undermine democracy and freedoms enjoyed in a democratic society.

- **Identity**. Although our identity in the physical world is well defined, and hardly needs to be proven during the course of a normal day, especially in small-scale environments, such as villages and towns, our identity in the cyberspace is almost entirely based on digital credentials (such as passwords), which can be lost, stolen, sold, and abused much like any piece of information. This opens the door to a wide variety of attacks that can lead to identity theft.

- **The Right to be Forgotten.** The European Commission recently proposed the implementation of the "Right to be Forgotten" in the Cyberspace, that is, the "right to ask service providers to delete the personal information that has been collected by data brokers under a users' consent" [330]. Similar in principle to the "right of oblivion" (in the French law), the Right to be Forgotten will empower people to take control of their digital image and reputation in the cyberspace.

- **Anonymity.** The widespread use of digital technologies has probably impacted anonymity more than any other aspect of our lives. For example, what used to be an anonymous stroll in the shopping mall, a leisure browsing of today's paper, or a casual watch of the television, has been transformed into a fully identified interaction with a network of content providers, a three-level-deep complex hierarchy of advertisers, and a variety of social networks all trying to record each and every aspect of every user's action.

- **Money.** Attackers have been traditionally motivated by the financial rewards of cybercrime, and will probably continue to be so.

2.4 Understanding the Domains of the Game

Evolutions in Science, Technology and Society create new domains that did not exist before. Domains may be defined by common technology platforms

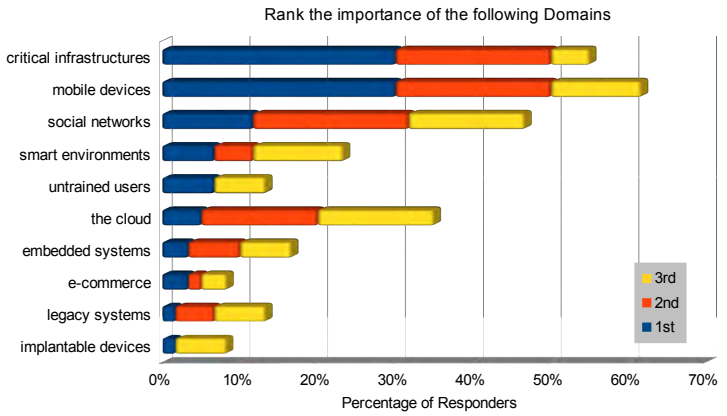

(such as mobile networks), by common societal developments (such as online social networks) and/or common applications (such as e-banking). Such domains may be subject to several threats and may impact several of our assets. However, studying them as an entire domain gives us the opportunity to identify common problems and opportunities.

- **Social Networks.** The advent of on-line Social Networks has redefined our understanding of privacy online. At the same time, the changing role of social networks to become the *de-facto* authentication and personalization means for the web will create interesting security and privacy challenges.

- **On-line Games**. On-line games and virtual worlds present at least two interesting opportunities for cyberattackers: (i) users spend a lot of their time playing games, and (ii) rewards awarded in on-line games can be monetized in the real word.

- **e-commerce**. As more people choose to make their purchases on-line this may increase their exposure to attacks, identity theft, and financial loss.

- **e-banking**. Over the past years most banking-related transactions involve communication networks in one way or another. Balance inquiries, on-line bill payments, debit card payments, on-line purchases, and a large number of other banking transactions involve digital communication networks.

- **Sensors - Drones**. We expect that in the near future we will have a very large number of sensors in our living environment. Ranging from low-lying RFID readers to high-flying unmanned airplanes (known as

drones), such sensors will record a wealth of information and will be an attractive target for attackers.

- **Embedded Systems**. It is expected that cars and objects of everyday use will have a large number of processors that will give them (more) autonomous operation and thus will make them subject to an increasing number of attacks.

- **SmartEnvironments**. The automation and "smart" operation promised by such environments will give more opportunities for attacks, as well as privacy concerns.

- **Legacy Systems**. Although modern systems are implemented in environments that discourage software vulnerabilities, a large part of our software was written several years ago when security was not the main concern.

- **Critical Infrastructures**. These may turn out to be one of the largest challenges faced by cyber security researchers. When Critical Infrastructures were isolated from the Internet, it was extremely difficult to attack them, especially without help from an insider. However, as Critical Infrastructures are being connected to the rest of the cyberspace, they present a high-value and more easily reachable target for attackers.

- **Mobile Systems**. The widespread use of mobile phones and the recent emergence of location-aware smart-phones has given rise to new interesting attacks on the security and privacy of users. Compromising a mobile phone is no longer about dialing a few high-premium numbers and charging the user extra roaming costs. It is about eaves-dropping on all the user's conversations; it is about "following" each and every footstep of the user; it is about having access to the most personal aspects of the users' lives.

- **Wireless Networks**. It has been said that children born in 2012 will not understand why we need "wires" to communicate. This is so true. Most of our communications today are wireless giving attackers the opportunity to jam them, to intercept them, to monitor them and (why not?) to modify them.

- **Implantable Devices**. As IT is integrated into medical care, it gives attackers more opportunities to compromise security and privacy. Implantable devices, for example, on which the patient's life depends, have been shown to be subject to battery draining and other attacks, threatening the lives of the individual patients.

- **The Cloud**. An increasing amount of data and computing operations is currently off-loaded to large-scale remote servers, collectively called "the cloud." These servers, which provide reliable, easily-accessible long-term storage and high-capacity computing capabilities, are being used both by individuals and by organizations to store their data and get some extra computing, if needed. Since cloud servers are a shared resource outside the end user's direct influence, they can easily be a major security/privacy concern.

2.5 Horizontal Research Directions

Although each topic (including attacks, vulnerabilities, assets and domains) includes underlying active research directions, there also exist *horizontal* research directions that can apply to most, if not all, of them. Such horizontal directions may include:

- **Usable security**. In order to be adopted, any security solution has to be easy to use, if not completely transparent, to the end user.
- **Authentication and Authorization**. An integral part in providing security solutions is the ability to authenticate the user (or even both ends) of a communication or transaction. If one of the end points can be spoofed, most security solutions will provide no protection at all.
- **Measuring security**. It has been said that security is more of an *Art* rather than an exact *Science*. This is partly due to the fact that Security can not be accurately measured. Imagine, for example, what would it mean for a system to be 99% secure? How about 99.9% secure? Would that be good? Would that be enough? Defining and measuring (even aspects of) Security is going to be a challenging, but an important area.

2.6 What If?

To make sure that we introduce some "disruptive thinking" to this process, we formulated and asked long-term "*what if*?" questions. Such questions aim to introduce provocative long-term investigations that will lead to fundamentally new thinking with respect to security and privacy. Thus, instead of focusing on small evolutionary improvements in traditional areas of research, we open the door to disruptive revolutionary advance that will create a agenda not for the next two, nor for the next five, but for the next ten to twenty years.

To give an example of such questions from various realms of science and engineering one would ask: "*What if we run out of oil? How will we be able to*

cover our energy needs?" Or: "*What if antibiotics do not work anymore? How will we be able to fight infections and diseases*?" Or: "*What if climate change results in an average sea level rise of two meters in the next few years? How will this impact our lives*?" In this spirit we set out to define a few ambitious questions in the area of security and privacy; questions that will make people think creatively; questions that will create a disruptive approach to security and an open mind to change.

- **What if your device can be made attacker-proof?** Indeed, assume that all computers, laptops, smartphones, etc. can be constructed in such a way that they are **malware-proof**. Let us assume that attackers can no longer install any kind of malicious program on any device. Would that end our security and privacy concerns? If not, what would be the main challenges in such an era?

- **What if 50% of computers out there are compromised?** Indeed, assume that a decent percentage, say 50%, of all computers (servers, desktops, smartphones, laptops, etc.) out there are compromised. How would this impact our sense of security and privacy on the Internet? How would this impact our everyday use of the Internet? How would this impact Internet-based financial growth and innovation?

- **What if you do not own your computing/communication devices anymore?** Let us assume a model where users do not own their devices. Users are allowed to use the devices but (i) they do not have full privileges to install whatever they want on the device, and (ii) at the end of the use period they have to return it, possibly to receive a better one with more features and capabilities. This model of use would be similar to the way we use rented cars today, the way we use company laptops, or even the way we use a hotel room. What would be the impact of such a use model and what will be the security and privacy concerns?

- **What if there is no money in cyber crime?** Let us assume a world where attackers can make very little money, if any at all. Imagine a world where SPAM does not pay, where click-fraud can be easily filtered out, and in general, a world where any shady activity does not pay off. Imagine a world where cyber crime simply does not make money for cyber criminals. What impact would this have on security and privacy?

- **What if the Internet shuts downs for a day or two?** Let us assume that sometime in the future the entire Internet shuts down for a day or two. Let us assume that all communications that are made possible by the Internet will just not be there anymore. How would this impact our lives? What kinds of activities will just not be possible? Furthermore, assume a

world where the threat of this outage is being taken for real by people and organizations, much like the threat of an earthquake or the threat of a tsunami. What impact would such a threat have on our well-being and in our financial lives?

- **What would you like to happen to your data when you pass away?** Assume a world where people keep lots, if not all, of their activities on-line. Assume that most of their photographs are on-line, most of their correspondence is on-line, most of their videos are on-line. Summer holiday pictures, falling-in-love letters, the first-day at school video, the picture of the tooth given to the tooth-fairy; all are on-line. To survive the occasional disk crash and the inevitable hardware upgrade, people would probably store their data in large-scale data centers, currently going by the name "the cloud." What options would we like to give people with respect to their data collection when they pass away? Will people be able to delete it? Will they be able to leave it as an inheritance to their children, much like they leave their family photo albums today? Will they be able to donate it to humankind, possibly for research? What security and privacy challenges would such a world create?

Part I: Threats Identified

3 In Search of Lost Anonymity

Dawn of a new era: as an increasing number of our activities take place over the Internet, a larger percentage of our actions are recorded every day somewhere on-line. Indeed, most of the news articles that we browse, most of the books we read, most of the videos we watch, and most of the things we purchase are recorded somewhere on-line. To make matters worse, even the activities that *do not take place on the Internet are recorded on-line*. For example, with the increasing penetration of smart-phones, most of the places we visit, most of the foods we eat, and most of the people we see are recorded on-line. Sophisticated artificial intelligence algorithms can usually infer the most personal details of our life: where we are, where we sleep, who we are in love with.

> *The cyberspace is an unforgiving medium: it has a lot of capacity to remember, but has no capability to forget.*

Take for example, the recently announced case where a major retail store managed to find out that a teenage girl was pregnant before her parents knew [28]. As surreal as it might seem, the same retail store managed to perfect its algorithms to the level where it *is able to know that women are pregnant even before they know it themselves* [27]. It is not hard to imagine that using algorithms based on artificial intelligence and correlating such findings with smartphone location-based data, such retail stores or data aggregators will soon be able to correctly guess the name of the father of the child as well!

3.1 Who Is Going to Be Affected?

Everyone is going to be affected to some degree or other. Private life will be lost. People will have to learn to lead their lives in the public domain and deal with it. Sadly, lost privacy is not something that can be "found," like a lost wallet, "re-issued" like a lost credit card, or "insured" like lost or stolen goods. Once privacy is lost, it is lost forever. Once an event is out there, it can not be retracted. The cyberspace is an unforgiving medium: it has a lot of capacity to remember, but has no capability to forget.

3.2 What Is Expected to Happen?

Although this information will probably be lawfully collected with the consent (or at least under the tolerance) of the users involved, it should be expected that parts of it will fall into the wrong hands either: (i) through legitimate ways, such as company acquisitions; (ii) through cracks in the system, such as imperfect security mechanisms; or (iii) through illegitimate ways, including extortion and theft. It will not be surprising, for example, if such information about candidates running for office is suddenly leaked just before election day. Similar leaks about company executives may happen just before important deals, such as company mergers, are due to be closed. Unfortunately, such misuse of the information will probably divert public focus from the important issues at hand, such as the company merger, to surprising details about the candidates' or the directors' past. And this is the real danger: losing one's perspective on the important things in life in the turmoil of everyday trivia; *losing sight of events that make history for the sake of details that make headlines*.

3.3 What Is the Worst That Can Happen?

The effects of this publicly available private-life information go beyond the tabloid gossip. We are not talking about juicy details of the lives of celebrities. We are talking about a generation of young people who will not be able to grow up normally. A generation of people who will not be able to have a secret. A generation that will not be able to enjoy the trust of sharing a secret. A generation that will not be able to build relationships solidified by the act of keeping a secret. A generation that will not be able to know the healing effect of being left alone, a generation that will be forced to heal its wounds, if possible at all, in the scrutiny of the public eye. A generation of people who will never learn that mistakes can be stepping stones towards success, but instead will live in fear of having their feet jerked from under them at any moment.

How do we expect the next generation to take risks when all their actions, all their failures, and all their mistakes—which sometimes may be spectacular—will be on-line for public ridicule? How do we expect them to fall in love when all their attempts, and unavoidable failures, will be there for public scrutiny?

Do we build a zoo, put our children in a cage, and invite everyone to watch? Is this the healthy environment we are preparing for the next generation?

3.4 State of the Art

Tracking web browsing within a domain and across domains has long been an issue with respect to user privacy [221]. Third-party domains appear as embedded components in a large number of distinct web sites. They are thus in a position to compile and correlate a user's browsing activity during visits to two or more of these sites. This practice has evolved from unsophisticated approaches, e.g., the use of HTTP cookies, to more elaborate techniques [160, 227], making it difficult for the average user to completely evade them. Even the *private* or *incognito* mode offered by modern browsers may not be enough for users to escape tracking [78].

Do Not Track [15, 16] is a browser technology which enables users to signal, via an HTTP header, that they do not wish to be tracked by websites they do not explicitly visit. Unfortunately, there are no guarantees that such a request will be honored by the receiving site.

We should not lose the sight of events that make history in the cloud of issues that fuel everyday headlines.

Krishnamurthy et al. [242] studied privacy leaks in online social networking services (SNS). They identified the presence of embedded content from third-party domains, such as advertisement providers, in the interactions of a user with the SNS itself, stressing that the combination with personal information inside an SNS could pose a significant threat to user privacy.

There has been significant work on the interplay between SNS and privacy. For example, there has been some focus on protecting privacy in SNS against third-party applications installed in a user's profile within the social network [163, 177, 352]. Facecloak [263] shields a user's personal information from an SNS, and any third-party interaction, by providing fake information to the SNS and storing actual, sensitive information in an encrypted form on a separate server. The authors in FlyByNight [262] propose the use of public key cryptography among friends in a SNS so as to protect their information from a curious social provider and potential data leaks.

Recent work has focused on how to support personalized advertisements without revealing the user's personal information to the providing party. Adnostic [378] offers targeted advertising while preserving the user's privacy by having the web browser profile users, through monitoring of their browsing history, and inferring their interests. It then downloads diverse content from the advertising server and selects which part of it to display to the user. Similarly, RePriv [185] enables the browser to mine a user's web behavior to infer

guidelines for content personalization, which are ultimately communicated to interested sites.

A series of browser add-ons exist [17,40] that block social plugins from the web pages a user visits by removing them or preventing them from loading, in a manner similar to the way Adblock [4] stops advertisements. However, they come at the cost of full loss of functionality as social plugins are completely removed from a page. Note that some of these add-ons are poorly implemented and naively remove the social plugins only after they have appeared on a page, meaning that the corresponding HTTP request containing user-identifying information has already been issued towards the server.

ShareMeNot [48] is a Firefox add-on that strips the cookies from a series of HTTP requests the web browser issues to load social plugins. As a result, no user-identifying information is sent to the social network until the user explicitly interacts with the social plugin. The downside of this approach is that users are deprived of any personalized information offered by the plugin, e.g., the number and names of any of their friends that might have already interacted with it. In other words, users view these social plugins as if they were logged out from the respective SNS (or browsing in "incognito" mode).

3.5 Research Gaps

It is true that privacy and anonymization can not be implemented using technical approaches alone. Legal and policy support is absolutely necessary. However in this report we will focus on the technical challenges. We envision research in the area along the following dimensions:

3.5.1 Prevention

Prevent information from being given away. Make sure that web sites and applications operate with the minimum information required. Demonstrate technologies that perform the required functionality with the minimum information possible. Develop anonymized versions of oneself. Develop systems, such as browsers, that transparently supply the appropriately anonymized version with the minimum possible information. For example, do not give a user's full ID to a web site that just needs to verify the visitor's age.

3.5.2 Monitoring

Monitor for information leakage at all possible levels. Develop honey-profiles (honeypots) to demonstrate and track information leakage. Reverse intrusion detection systems can also be used for continuous monitoring.

3.5.3 Deletion

Develop approaches to (selectively) delete one's data. As a simplest case, consider the *right to be forgotten* [330]. Advance research on how (or if) this can be technically implemented. Focus on how to selectively delete only aspects of one's profile.

3.5.4 Anonymization

Develop mechanisms to anonymize and share data in anonymized form. Then, data collectors and aggregators would be required to process data only in anonymized or encrypted forms.

3.6 Example problems

Tangible example problems might include:

Honey-profiles. Inspired by honeypots, honeynets, and honeytraps used to study attackers and spammers, one may use honeypot-like profiles to track the leakage of information. Honey-profiles could be created and supplied to web sites and applications. These profiles will be furnished with distinctive characteristics, such as specific interests, nicknames, email addresses, etc., that could be recognized in a subsequent feedback loop.

Provide personal data in a provably *k*-anonymous form. Several applications ask permission to have access to personal data, such as the users' age, in order to make sure that the users are over 18. In this research one might change the model and force applications to ask data from *sets* of users. The set will provide data (such as the names of all users in the set and a proof that all of them are over 18), without revealing, however, which individual user of the set is using the current application.

Privacy in an eponymous world . There exist cases where the user's identity can not be hidden, such as when the user performs an operation (such as a web search) while logged in. Explore whether there exist approaches to hide a user's real interests even in this eponymous world. For example, consider users who support one political party but would not like their political views to be registered. For this reason they might visit the web pages of several, or even all, political parties, spending about the same time at each of them. Tools that automate and extend such procedures might significantly confuse classification algorithms. Thus, while users will not be able to hide the fact that they visited a particular party's web site, they will also volunteer the fact that they visited all parties' web sites, making it difficult to classify their political views.

4 Software Vulnerabilities

Extending its definition in the physical world, in computer security a *vulnerability* is a weakness or flaw in one or more software components that can be exploited to compromise the integrity, confidentiality, or availability of a system and its information resources [217]. Besides software, vulnerabilities may exist in other aspects of a system, including protocol design, hardware, system configuration, and operational procedures. After many years of security research and engineering, software vulnerabilities remain one of the primary methods of reducing a system's information assurance.

The massive complexity of modern software is one of the main reasons for the existence of flaws that can lead to system compromise. Vendors also often give security design a secondary priority in favor of rich features, time to market, performance, and overall cost. At the same time, the incessant hunt for new vulnerabilities by malicious hackers, criminals, spies, and even nation states, has resulted in the continuous discovery of new vulnerabilities and in major advances in exploitation techniques.

Common types of software flaws that can lead to vulnerabilities that could be exploited by a malicious adversary include:

Memory errors: buffer overflows, dynamic memory errors (dangling pointers, double or invalid frees, null pointer dereferences), uninitialized variables.

Input validation errors: code or command injection, SQL injection, uncontrolled format strings, cross-site scripting (XSS), directory traversal.

Race conditions: simultaneous access, time-of-check-to-time-of-use (TOCTOU) bugs.

Privilege-confusion: cross-site request forgery (CSRF), clickjacking.

In 2011, the MITRE corporation, an American not-for-profit organization, through its Common Weakness Enumeration (CWE) effort, a community-developed dictionary of software weakness types, compiled a list of the most widespread and critical errors that can lead to serious software vulnerabilities [66]. Organized into three categories, the top 25 most frequently exploited software flaws are:

Insecure Interaction Between Components

- CWE-89 Improper Neutralization of Special Elements used in an SQL Command (“SQL Injection”)
- CWE-78 Improper Neutralization of Special Elements used in an OS Command (“OS Command Injection”)
- CWE-79 Improper Neutralization of Input During Web Page Generation (“Cross-site Scripting”)
- CWE-434 Unrestricted Upload of File with Dangerous Type
- CWE-352 Cross-Site Request Forgery (CSRF)
- CWE-601 URL Redirection to Untrusted Site (“Open Redirect”)

Risky Resource Management

- CWE-120 Buffer Copy without Checking Size of Input (“Classic Buffer Overflow”)
- CWE-22 Improper Limitation of a Pathname to a Restricted Directory (“Path Traversal”)
- CWE-494 Download of Code Without Integrity Check
- CWE-829 Inclusion of Functionality from Untrusted Control Sphere
- CWE-676 Use of Potentially Dangerous Function
- CWE-131 Incorrect Calculation of Buffer Size
- CWE-134 Uncontrolled Format String
- CWE-190 Integer Overflow or Wraparound

Porous Defenses

- CWE-306 Missing Authentication for Critical Function
- CWE-862 Missing Authorization
- CWE-798 Use of Hard-coded Credentials
- CWE-311 Missing Encryption of Sensitive Data
- CWE-807 Reliance on Untrusted Inputs in a Security Decision
- CWE-250 Execution with Unnecessary Privileges
- CWE-863 Incorrect Authorization
- CWE-732 Incorrect Permission Assignment for Critical Resource
- CWE-327 Use of a Broken or Risky Cryptographic Algorithm
- CWE-307 Improper Restriction of Excessive Authentication Attempts
- CWE-759 Use of a One-Way Hash without a Salt

The exploitation of widespread vulnerabilities like the above can lead to security and privacy breaches in essentially any domain of our digital

infrastructure. A thorough discussion of the most commonly used attacks and exploitation techniques is provided in the SysSec Deliverable D7.1: *Review of the State-of-the-Art in Cyberattacks* [373].

4.1 What Is the Problem?

Despite significant advances in software protection and attack mitigation techniques, *exploitable* vulnerabilities are continuously being discovered even in the latest versions of widely used applications, programming libraries, operating systems, online services, embedded software, and other programs. For instance, the exploitation of memory corruption vulnerabilities in server and client applications has been one of the most prevalent means of system compromise and malware infection. Recent prominent examples include in-the-wild exploits against Internet Explorer [33], Adobe Flash Player [6], and Adobe Reader [5, 101], all capable of successfully bypassing the data execution prevention (DEP) and address space layout randomization (ASLR) protections of Windows [279], even on the most recent and fully updated (at the time of public notice) systems.

As secure programming, software protections, and exploit mitigation mechanisms have become more widely employed, successful compromise might require the combined exploitation of multiple vulnerabilities on the same system. A recent exploit against the Chrome browser required a chain of six different vulnerabilities to successfully break out of the Chrome sandbox and achieve arbitrary remote code execution [304].

Besides arbitrary code execution, other possible outcomes with less freedom of choice for the attacker, but probably of equal severity, include disclosure or modification of private data, privilege escalation, logic errors, denial of service, and other unauthorized actions. Indicatively, a memory corruption vulnerability may allow the modification of critical application data, including user identity, configuration, user input, and decision-making information [124].

Privilege escalation attacks are an important threat in multi-user environments or multi-tenant cloud services, as they can allow less-privileged users to gain root access and compromise other users and the system itself. The proliferation of mobile operating systems such as iOS and Android, in which third-party applications run with lower privileges, has made privilege escalation attacks particular relevant, as they can allow malicious applications to gain unrestricted access to a user's device.

The shift towards web services and cloud-based applications has also given rise to a multitude of web-specific attacks that exploit vulnerabilities anywhere between a client's browser and a server's back-end systems. The exploitation of SQL injection, XSS, CSRF, and other more subtle types of application flaws in web services can lead to the disclosure of massive amounts of private

information. Numerous security breaches in high-profile online services have resulted in unauthorized access to whole databases with millions of entries consisting of usernames, passwords, financial information, and other private data [26, 50, 59, 61].

4.2 Who Is Going to Be Affected?

An ever increasing part of our business, social, and personal life involves online services and software running on personal devices and equipment that we use or depend on. Therefore, the risk of exploitable software vulnerabilities can affect all of us. Even people or organizations who do not engage in online interactions or who even do not use any computing devices at all can be affected, as software controls major parts of critical infrastructures. For instance, the massive outbreak of the Conficker worm at the beginning of 2009 resulted in more than 10 million infected machines worldwide [319]. Among the millions of infected computers were machines in the Houses of Parliament in London [286] and the Manchester City Council [254]. Just in the latter case, the infection cost an estimated £1.5 million in total. After a security breach at LinkedIn, a popular social networking service for people in professional occupations, the (weakly) encrypted passwords of more than 6.4 million users were exposed [26].

4.3 What Is Expected to Happen?

Although decades of research and development in secure programming and software protections have materialized in most of the widely used operating systems and applications, experience has shown that the rate of discovery of software vulnerabilities keeps increasing. Given the professionalism and determination of criminals and other threat agents, and the ever increasing complexity and interdependence of current software systems, it is expected that software vulnerabilities will not be eradicated anytime soon. At the same time, the increasing sophistication of recent exploits [101, 304] is an indication that the detection and mitigation of future threats will become harder as a result of the more prevalent use of evasion techniques and stealthy attacks.

4.4 What Is the Worst That Can Happen?

Besides compromising the security and privacy of our digital interactions, software vulnerabilities can put at risk other parts of our daily activities, or even our lives. In the same way a worm subverted industrial systems within Iran's nuclear facilities [250], an extremist group could attempt to compromise

parts of critical infrastructures, such as power grids and traffic control systems, perhaps causing severe damage and potentially mass casualties. Threats against critical infrastructures are further discussed in Chapter 6.

Smaller-scale hostile acts could also be facilitated by the prevalence of software-controlled devices and equipment. Implantable medical devices [206] and cars [123] are two prominent examples.

4.5 State of the Art

After decades of research and engineering aimed at dependable and secure computing [88] with the broader aim of minimizing the undesirable effects of software bugs, and consequently the potential threats stemming from the exploitation of software vulnerabilities, there is a vast amount of literature on the subject [137, 170, 215, 258, 315, 340, 382]. In this section we briefly summarize different broad areas of techniques that contribute towards lowering the risk of software vulnerabilities, especially in terms of their potential to be successfully exploited. A more focused discussion of solutions directed against memory corruption vulnerabilities is provided in Chapter 9.5.

Numerous techniques seek to provide a proactive defense against future threats by eliminating or minimizing certain classes of vulnerabilities, or preventing the manifestation of certain exploitation methods. Broad areas include programming language security features, code analysis techniques, confinement mechanisms, and diversification. Besides best security practices and defensive programming, software hardening techniques include: static source code analysis for finding and eliminating certain classes of programming flaws [387]; augmenting programs with runtime protections using compiler extensions [117, 136, 306], static binary instrumentation [297, 321], dynamic binary instrumentation [125, 235, 299, 325], or library interposition [96]; software fault isolation and sandboxing [196, 238]; and control flow integrity [76].

At the operating system level, many different techniques aim to hinder the exploitation of software vulnerabilities, including non-executable pages [151, 311], ASLR and code diversification [102, 102, 132, 183, 279, 308, 310], and instruction-set randomization [233, 233]. As additional protections usually incur significant runtime overhead, CPUs are constantly enhanced with security features that facilitate the implementation of more lightweight solutions [184, 194, 245].

In the field of web services and cloud-based applications, enhancements and improvements in numerous system aspects, from the browser to the server, aim to improve the security of online interactions. Indicatively, different areas of focus include fundamental design choices of the web platform [98], specific shortcomings of browser implementations [97, 257], cross-site scripting (XSS) [201, 223, 293, 336, 345, 371], and more subtle complexities of web appli-

cations [91, 106, 342]. Also, many academic efforts aim at applying security concepts from operating systems to the web platform [198, 332, 369, 389, 390].

4.6 Research Gaps

As software vulnerabilities are the primary source of security breaches, the use of software hardening and exploit mitigation techniques is very important, as they can offer instant and effective protection against current and future threats. However, the runtime overhead that many of these mechanisms impose, and the deployment complexity of others, often prevent their widespread adoption. Furthermore, not all software developers harden the software they write, source code is not available for commercial applications, and determined attackers have been constantly bypassing security measures [82, 101, 114, 159, 304].

For instance, although address space randomization is an effective countermeasure against code-reuse attacks, its effectiveness is hindered by code segments left in static locations [186, 224], while, depending on the randomization entropy, it might be possible to circumvent it using brute-force guessing [347]. Even if all the code segments of a process are fully randomized, vulnerabilities that allow the leakage of memory contents can enable the calculation of the base address of a DLL at runtime [256, 346, 386]. The above is indicative of a constantly recurring pattern of exploit mitigations that require significant effort to be deployed and adopted, only to be bypassed by a more sophisticated or alternative exploitation technique later on.

Another important issue is the reliance on "remedy" methods that remove specific software vulnerabilities after they have been exposed (e.g., through software patches and automated operating system updates), or provide some mitigation by disinfecting infected hosts (e.g., through the use of virus scanners). Although such methods are very important for keeping systems in good health and up to date with the latest security fixes, they cannot protect against attacks that exploit previously unknown vulnerabilities. In such cases, the relevant patches or signatures provided by vendors to fix new security flaws or detect new malware usually come late, as their generation relies on significant, time-consuming human intervention [350]. Furthermore, administrators are sometimes reluctant to use automated patch installation systems, as they first need to verify through extensive testing that the new patches will not introduce any stability problems [333].

4.7 Example Problems

A few indicative issues for which existing solutions still do not provide a satisfactory level of protection include:

Memory corruption vulnerabilities: despite numerous approaches, from programming language and compiler level improvements to operating sys-

tem mitigations, exploitable memory errors are still being found even in the latest versions of widely used applications. The rise of mobile operating systems such as iOS and Android, in which third-party applications run with lower privileges, has also made kernel-level memory corruption vulnerabilities more relevant than ever—an area that has received less attention compared to user-level applications because of the different threat model that usually applies on personal computers compared to mobile devices.

Data exposure vulnerabilities: the fairly new ecosystem of feature-rich web services and cloud-based applications, with the numerous components and interactions that are involved, continuously exposes subtle flaws in the languages, APIs, protocols, and client or server software used. Although so far a great deal of attention has been paid to preventing a user's machine being taken over, as user data are shifted to the cloud, vulnerabilities in any of the above stages can lead to the exposure of confidential information, ranging from browser cookies to private documents, with equally harmful consequences.

5 Social Networks

Framing the familiar notion of "social networks" in the Digial Era, most people tend to think about *online social networks* such as Facebook, Google+, and Twitter. As the popularity and use of such online social networks has increased, the attackers have also started considering how to use them for nefarious activities. They have become new platforms for conducting malicious activities and desirable targets for launching attacks.

However, such online social networks are just a subset of all the social networks that actually exist. Any set of people and their internal social relationships, representing their interaction, collaboration, or other sort of influence between them, can be modeled as a general social network. These relationships are formed by exchanging emails, making phone calls, co-authoring a scientific article, or a range of other "normal" activities that will build up a network. Such information is now collected and organized to gain insights into people's lives, but this is also a venue that attackers will use. They can either attack the properties of the network by, for example, introducing false nodes, or gain enough information to attack the individual users.

The explosive growth rate of social networks has created the first *digital generation*, consisting of people of all ages and backgrounds. People are creating their digital counterparts for interacting with other users, for both recreational and professional reasons, and may disclose a vast amount of personal data in an attempt to utilize these new services to the fullest. As the social network is a representation of social interaction, it also indirectly shows the trust between different individuals. However, the lack of technical literacy among the majority of users has resulted in a naive approach where the caution demonstrated in the social interactions of the physical world has disappeared. Users are vulnerable to a series of dangers, ranging from identity theft to monetary loss, and lack the critical approach that develops over time and is passed on through generations. As users tend to show a great amount of trust to online communication and interactions, adversaries aim to sneak into a victim's circle of trust through impersonation. As people trust their friends, the cyber criminal can then perform a range of attacks that may not be possible, or effective, as a "stranger."

5.1 Who Is Going to Be Affected?

As we bring our social interactions online, they can be tracked and recorded. Thus, it is not only the specific users of online social networks, such as Facebook or Google+, that will be affected, but anyone who participates in a society that is dependent on information and communication technology. The generation growing up now is the first to have been exposed from birth. As such technology is continuously being integrated into our lives, it is to be expected that more information will be gathered and more people will be affected in the future.

5.2 What Is Expected to Happen?

We are already seeing attackers targeting the online social networks, such as Facebook, Google+ and Twitter. By befriending strangers, cyber criminals can harvest private user data and access their contact lists to perform identity theft, clone user profiles, lure them to malicious websites, and send targeted spam and phishing messages.

5.2.1 Privacy

Users of online social networks tend to share private information, such as education, occupation, relationship status, current location, and personal habits. In the wrong hands, this information can be used to launch sophisticated and targeted attacks against people. Even for individuals who are not users of online social networks, information about their social interactions can still be inferred from public data, such as co-authorship information from DBLP [14]. The problems caused by breaching privacy are described in more detail in Chapter 3.

5.2.2 Spam

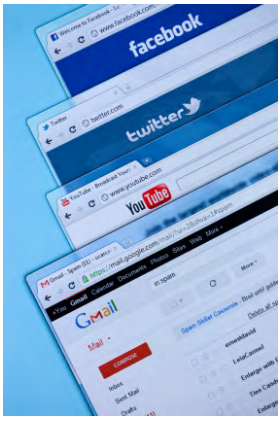

Today, email is not any more the only means for spreading spam, as spammers now use multiple content-sharing platforms, such as online social networks, to increase their success rate. The information provided by users in their profiles, such as education, profession, and relationship status, together with their real email address, provides spammers with a great opportunity to personalize their marketing activities and improve the efficiency of the spam campaigns. Moreover, if a spam email contains personal information, such as the name of the receiver, content-based spam detection tools assign lower spam rates to it and it may therefore

evade detection. Hence, new filtering techniques are required to counter this new type of spam. Third-party applications can also exploit vulnerabilities in users' browsers, conduct cross-site scripting attacks, compromise their machines, form a botnet to launch malicious activities such as DoS attacks, propagate malware, or send spam email. One example is the Koobface botnet [93], which abuses social network sites. It starts by sending a spam message containing a malicious link to a video which, once clicked, redirects the victims to a website where they are asked to install an executable file. This file then infects their machine with the Koobface malware.

5.2.3 Sybil Attack

In addition to the propagation of spam and malware, multiple fake identities in social networks can be used to out-vote honest users, influence online ratings, and manipulate search results [225]. Attackers can either compromise existing accounts or generate fake (Sybil) accounts. The compromised accounts have already established friendships with other users of the social network and are trusted by them. Sybil accounts, however, need to establish friendships and gain trust before launching attacks, such as sending spam. Selling fraudulent or compromised social network accounts is even starting to overtake stealing credit card numbers [275, 334].

5.2.4 Authentication

In order to mitigate attacks from compromised accounts, mechanisms requiring more than a password have been introduced, such as Social Authentication (SA) in Facebook [173]. These types of mechanism require a user to provide two distinct pieces of evidence in order to be authenticated. For example, in SA, users must provide a password and recognize pictures randomly chosen from their friends' pictures. Unfortunately, this type of authentication is vulnerable to advances in face recognition techniques [104, 318]. Different approaches to authentication and authorization, as well as general problems that exist are covered in Chapter 7.

5.2.5 Third Parties

Third-party applications, which are widely deployed in online social networks, can also perform malicious activities, for example exploit vulnerabilities in users' browsers, conduct cross-site scripting attacks, compromise their machines, form a botnet to launch attacks such as denial of service [87], propagate malware, and send spam.

Moreover, malicious third-party applications that access private user data tend to store the information, or send it to advertising and Internet tracking companies, thus violating user privacy. Unfortunately, any data harvested

by such an application are then beyond the control of the social network site. Although online social networks such as Facebook have introduced coarse-grained access-control mechanisms for third-party applications, there is a need for more fine-grained mechanisms [163].

Third-party websites can also use the social plugins provided by social network sites such as Facebook [174] in order to personalize their content, allow users to write feedback for their sites, share the page content with their friends in the social networks, or even be authenticated by a social login plugin. Unfortunately, these plugins also allow third-party websites to access private user data, and allow the social network sites to track user activities outside their platform [239].

5.3 What Is the Worst That Can Happen?

In the wrong hands, online social networks can be used to disseminate wrong information, to perform political censorship [375], to bias public opinion [289], and influence users [281]. New types of attack are also emerging in social networks. In a reverse social engineering attack, an adversary aims at tricking victims into contacting the fake/compromised accounts that are under the attacker's control, instead of contacting the victims directly [218].

Information about people's social interactions can be exploited as a side-channel for different types of attacks. It has been shown [358] that it is possible to use the public social network data to conduct efficient de-anonymization attacks against mobility data. One example is that by using the co-authorship information from DBLP, the authors could generate a social network of conference attendees, and then leverage it to de-anonymize 80% of the nodes in the Infocom 2006 Bluetooth contact traces. In addition, highly sensitive personal information can be inferred from online social networks, even if the user does not explicitly like specific posts or pages. Kosinski et al. [241] have shown that information related to users' sexual orientation or political views can be predicted from other activities with a high accuracy.

5.4 State of the Art

As mentioned in Chapter 3, a considerable amount of work has been devoted to the privacy of social network sites. More examples include Persona, which uses attribute-based encryption and allows users to dictate policies regarding who may view their information [89], and Safebook, a decentralized and privacy-preserving online social network application [142].

Multiple fake identity (*Sybils*) attacks on social networks have been used for forwarding spam and malware, out-voting honest users, and manipulating

online ratings. A variety of solutions, such as SybilGuard [411], SybilLimit [410], SybilInfer [145], SumUp [379], and Canal [384], have been proposed in the literature, which detect Sybils by identifying tightly connected communities of Sybil nodes [385]. However, the assumptions that these solutions make about the structural properties of social networks and the communities of Sybils do not hold for all Sybils in real social networks [285, 408].

Identifying spammers in social networks has also received considerable attention. Different methods have been proposed to automatically identify the accounts used by spammers [364], and to identify more criminal accounts via the study of social relationships from a number of known malicious accounts [406]. Numerous studies of the different types of social networks and the structural properties of social graphs have provided us with an understanding of the structural properties of social networks that correspond to normal social behavior [252, 282, 298]. This knowledge can be used to identify anomalies associated with malicious activities, such as sending spam emails. The distinguishing characteristics of spam and legitimate emails extracted from the structural properties of email networks have led to the introduction of new methods for detecting spammers [111, 197, 247, 288, 380].

Social links that correspond to interpersonal trust relationships have provided a means to populate white lists of legitimate email senders in Reliable Email (RE) [188], to thwart unwanted communications in OSTRA [283], and to mitigate trust-aware collaborative spam in SocialFilter [353]. Social network analysis has also been used for other applications, such as fraud detection [273]. Studies of the social structural properties of human communications have even led to the introduction of new approaches to network intrusion detection using network flow data [156]. Using this method, intruders are identified based on their anti-social behavior in entering communities to which they do not belong. The study of the community structure of social graphs has also allowed the discovery of algorithms that can be used to separate legitimate from unwanted email communications by clustering them into distinct communities [287].

5.5 Research Gaps

The significant growth in the use of social networks as platforms for information dissemination makes it challenging to identify the trustworthiness of the data we consume. As mentioned above, social networks can be used for spreading propaganda and misinformation. Assessing the correctness of propagated information can be very challenging, especially in the setting of anonymity. Therefore, estimation of the trust-worthiness of information sources remains an interesting research gap.

Social networks provide us with a wealth of real-time content that includes significant data concerning ongoing events, such as political elections or natural

disasters. This real-time information is expanding out of all proportion to our ability to process it. This is not just a data mining problem; rather, processing data streams for identification of malicious content and fraudulent sources is the main challenge.

There have been numerous studies of graph mining, which has led to a variety of techniques for mining relational data. These techniques can be used to characterize the structural properties of social graphs that correspond to normal social behavior. This in turn can lead to methods for detecting anomalies, such as possible attacks, that do not conform to the expected social behavior. Although studies of snapshots of social networks can be used for security research, tracking changes over time can give a better insight into the dynamic nature of the network and attackers' behavior. Therefore, coping with the dynamicity of data is another big challenge.

In studies of social interactions between people, there is a trade-off between security and privacy. Data collection and the processing of user activity logs can lead to detection of compromised or malicious accounts in social networks; however, this has to be done in a privacy-preserving manner, otherwise it is not practical.

5.6 Example Problems

With respect to the above research gaps, the following example problems can be defined.

How can we arrive at a measure of confidence in the truthfulness of information that is disseminated through social networks? This is a challenging problem, particularly in cases where the true identity and the trustworthiness of the source of information are not known, due to anonymity, and in cases where there are no other ways to verify the content.

How can we collectively study the information collected from different sources in real time? There is a need for an engine for organizing real-time streaming data gathered from a variety of social sensing platforms, including social networks [388]. How can we effectively parallelize and distribute the data stream processing and introduce methods for identifying cyber criminals based on the aggregated data?

How can we utilize data mining techniques for discriminating between honest and malicious identities? The well-studied techniques for graph mining can be deployed as a tool for combating cyber criminals. Although the structural properties of social graphs have already been used against Sybil attacks and spam, much more can be done.

6 Critical Infrastructure Security

GREATER IN SIZE THAN ANYTHING build so far, current *critical infrastructures* (CI) refer to systems or assets that are vital in modern society and economy. Water supply, electricity, transportation, financial services, health care and telecommunication are the most common examples of CIs. CIs are regulated by different rules and laws, and operated diversely from country to country. In addition, CIs are influenced by non-technological factors such as politics or culture. According to the EU Directive 2008/114/EC [63], a CI is

> "an asset [...] which is essential for the maintenance of vital societal functions, health, safety, security, economic or social well-being of people, and the disruption or destruction of which would have a significant impact [...] as a result of the failure to maintain those functions."

Thanks to the evolution of information and telecommunication technology, controlling CIs remotely (e.g., over the Internet) is feasible and, more importantly, convenient. Therefore, CI actors (e.g., industries and governments) have been progressively incorporating IT systems to consolidate the operation of CIs, up to the point that CIs and IT systems have converged. The term cyber-physical system (CPS) is commonly used in this context to refer to the integration of a physical (critical) system with a cyber (Internet-connected) system, which is typically an industrial control system (ICS). In the remainder of this section, we will use the term CI to refer to the critical infrastructure as a part of the physical environment, and the term CPS to refer to the systems that comprise and interconnect these infrastructures, thus including IT components (i.e., the ICSs).

Security issues arise because two previously isolated worlds, the Internet and the CI systems, are now interconnected. When early CIs were created, neither security nor misuse of the interconnected control system were considered. As a matter of fact, Internet technology is itself an underlying, critical asset of modern CIs, because the ICSs that control them are often distributed (over remote, Internet-connected locations).

This section highlights the most relevant security problems and the state of the art of CPSs, with a particular emphasis on the ICS part.

6.1 What Is the Problem?

From the above premises, it is clear that well-known challenging threats such as malware, botnets, or denial of service attacks, which have been compromising the security of Internet-connected devices, are likely to become threats for CIs as well. In contrast to traditional Internet-connected devices, CIs can take tangible actions in the physical environment, thus posing serious safety risks, along with the possibility of production loss, equipment damage and information theft. The first incident on a SCADA system dates back to 1982, when a trojan supposedly infected the ICS that controlled the so-called "Siberian Pipeline" and caused an explosion equivalent to 3 kilotons of TNT [278]. Further exacerbating this scenario, today's SCADA-controlled systems are widespread, given the market traction of smart grids and smart buildings, and thus more appealing to offenders [361,392,395]. Although SCADA implementations can vary from vendor to vendor, the specifications of the control protocols (e.g., PLC) are publicly available [32] and the devices can be acquired by anyone who has sufficient funding. In addition, the control software runs on general purpose OSs (e.g., Windows), and devices were originally deployed in isolated environments where network connectivity was not considered. Needless to say, SCADA software comes with several serious vulnerabilities [47], most of them caused by buffer overflow and input validation bugs, which culminated in experts describing SCADA security as "laughable" [51]. Unfortunately, these vulnerable ICS are publicly accessible over the Internet. One such center of exploits is called SHODAN [49], a search engine tailored at finding and exposing online embedded devices such as webcams, routers, power plants or even wind turbines. Unsurprisingly, "scada" is the most searched term on SHODAN. How well these exploits perform in real-world scenarios, however, is hard to estimate.

According to the information that CERTS and governments collected, offenders increasingly targeted critical infrastructures of countries: The Industrial Control Systems-Cyber Emergency Response Team (ICS-CERT) responded to 198 incidents against CIs in 2012, 52% more than the previous year. The two most impacted sectors in 2012 are energy (41% of reported incidents) and water (15%) [45]. There are debates within the research community about the accuracy of the answers collected in a recent survey conducted by SANS among industries and organizations that adopt SCADA and process-control systems [74]. Despite such debates, the survey corroborates the anecdotal belief that SCADA and ICS adopters are aware of the security risks. Roughly 50% of the participants reported that they were taking countermeasures that included patching, access control and log analysis. Unfortunately, the PLC layer appears to be a weak spot, where it is often difficult to deploy proper monitoring mechanisms.

6.2 Who Is Going to Be Affected?

Several sectors are theoretically exposed to the aforementioned threats. Basically, every adopter of network-connected process-control systems is likely to be affected. Public health, energy production, telecommunication and public water supply are just a few examples of systems that will be under threat unless they deploy adequate countermeasures.

Furthermore, an aggravating factor is, that today's CIs are getting larger. With the increasing adoption of smart grids, virtually everyone, even individuals, is part of the CPS ecosystem. Cyber attacks are therefore likely to affect everyone. Even when ordinary people are not directly affected by failures of modern CPSs, they are still susceptible to cascade effects. Since sectors adopting ICSs are also influenced by cultural, political or economical factors, the impact of an attack is more widespread than in an isolated system that uses ICSs, for instance, to control production.

6.3 What Is Expected to Happen?

The critical nature of CIs renders them intriguing targets with disastrous consequences, including loss of human lives. In some respects, predictions from 3 years ago can already be observed in the wild. Yet, it appears that the actors behind the weekly reported threats are probing without causing deliberate damage. For instance, the Stuxnet [176] infection of 2009–2010, which influenced thousands of devices, reached very sensitive targets. A recent report [272] describes that earlier versions of the sophisticated cyber weapon contained other known versions of the malicious code that were reportedly unleashed by the US and Israel several years ago, in an attempt to sabotage Iran's nuclear program. This indicates that Stuxnet was active about two years before the main incident. It also implies that none of the two campaigns of Stuxnet (2007 and 2009–2010) had a serious impact on Iran's nuclear facilities, the avowed main target of the attack. Even though Stuxnet essentially failed, an important fact remains: Stuxnet was developed (by offices of nation states, as recently confirmed officially [69,70], although the US government has never admitted using cyber weapons) with careful planning and the use of product-specific 0-day vulnerabilities, and it had the potential and the opportunity to cause serious damage on a national level.

The widespread belief that standard protection tools (e.g., VPNs, firewalls, etc.) would suffice to secure network-connected SCADA equipment is just a myth. In fact, Stuxnet reached its targets from an infected USB drive. It then used other exploits and local-network probing techniques to find and infect other targets within the production environment. This attack vector is impossible to restrict with network-based access control alone. Instead, a full-blown security infrastructure, including access and account policies would

be needed, something that is not even supported by most SCADA systems and their backbones.

Subsequent milestones were Duqu (2011) and Flame (2012), both designed with intelligence gathering purposes, although Flame is more opportunistic as it spreads also to mobile devices and uses ambient sensors (e.g., microphone) to steal information. These are two examples of the second most important application of cyber weapons: espionage. Due to the similarity of some code fragments of Duqu, Flame and the variants of Stuxnet, it is not unrealistic to conclude that Duqu was designed to be the precursor of the next Stuxnet [127], to gather intelligence about CI targets.

Whether Flame will be the precursor of the often predicted "year of cyber attacks (2013)," remains to be seen. As mentioned in Section 6.7, recent industrial research efforts are moving toward this direction by deploying honeypot ICSs to collect object evidence of attacks, which would be of help in answering these questions.

6.4 What Is the Worst That Can Happen?

The discovery of Stuxnet, and the related events, concretely showed to the security the potential impact of attacks against CIs; this significantly increased the concerns and interest of the community. Today, vulnerabilities and attacks against CIs continue to be discovered every week in the wild. A recent case is dated February 23, 2013, when the US Department of Homeland Security (DHS) reported that in a cyberattack against 23 natural gas pipeline operators, crucial information was stolen [71]. Although the DHS report, not yet disclosed to the public, does not mention the sources of the espionage, the digital signatures of the attacks have been identified by independent researchers as belonging to a particular group recently linked to China's military (although China has denied the allegations) [72]. Unfortunately, these attacks will continue to spread. This is corroborated by the increased amount of incidents reported and, more importantly, by the recent cyber-espionage cases, which are likely to be the precursor of more targeted and sophisticated attacks.

In addition to (intentional) attacks and unintentional incidents in ICSs, which both impact the physical world, we believe that unintentionally caused failures are also bound to happen. Instability, natural and artificial faults [259] or unexpected conditions in the physical systems, which eventually translate into "signals," processed by ICSs to take proper control actions, can

retroactively lead to unexpected conditions in the ICS software, which could ultimately lead to failure loops with devastating consequences.

These premises allow us to draw a global picture of what could happen in the future if the current menaces continue their evolution. The word "cyberwar" [119] appears frequently in the majority of recent threat reports and news subsections. This word should be used with care, because, as of March 2013, there is no strong evidence as to whether the aforementioned threats have translated into concrete, planned attacks, as opposed to "testing" performed by the attackers (or governments). On the other hand, the future scenario is frightening as it includes disasters caused by viruses like Stuxnet that infect critical control systems, causing such events as traffic accidents, train or plane collisions, nuclear power plans meltdowns or explosions. Needless to say, such attacks may end up with a massive loss of life and an exacerbate the global financial crisis. Ultimately, the economy is also a critical system, with strong impact on the physical world, which is highly dependent on computers. Once attackers have gained control of a CI, they can operate it at their will.

6.5 State of the Art

Recent EU-funded research projects concerning the security of CIs are CRISALIS (`http://www.crisalis-project.eu/`), which focuses on practical aspects of detection of vulnerabilities and attacks, and SESAME (`https://www.sesame-project.eu/`), with the same focus, although more oriented toward observing the CIs from the physical side (mainly on smart grids). We also refer the reader to recent work on attack assessment [381], analysis [376] (on espionage attack triage), survey and challenges of smart grid security [392] and critique [314]. Recent reference books worth mentioning are [146, 328].

With system security of CIs being a young research field, a few *notable publications*—reviewed in the remainder of this section—appeared in the last two years at leading conferences. Most of the literature about detection or protection methods focuses on SCADA protocols or on smart grids.

6.5.1 Anomaly Detection of SCADA Events and Protocols

[203, 204] address the detection of *process-related threats* in ICS used in CIs. These threats take place when an attacker impersonates a user to perform actions that appear legitimate although they are intended to disrupt the industrial process. They tested their approach on 101,025 *log entries to detect anomalous patterns* of user actions. This preliminary case study suggests that the approach is effective. One year later the same authors extended their work beyond log analysis and are concentrating on binary protocols, including those adopted by SCADA implementations (e.g., MODBUS). The motivation behind [68] is that several complex and high-impact attacks specifically targeting binary protocols

were reported. They apply anomaly-based algorithms using n-gram analysis for payload inspection. This approach is basically very general, because it requires zero knowledge of the underlying protocol specification. They present a thorough analysis and evaluation of several detection algorithms that apply variants of n-gram analysis to real-life environments. They tested the approach on a ICS dataset collected of a real-world plant over 30 days of observation. These data are used as background training traffic. They also generated network traces of exploits against such binary protocols.[1] They finally conclude that, in the presence of data with high variety, high detection rates and low false positive rates are unfeasible at the same time.

6.5.2 Attacking and Protecting PLCs

Along a similar line, although with a different purpose, the authors of [274] focus on PLCs, which drive the behavior of ICSs. They work under the assumption of an adversary with no knowledge of the PLC's and the objective of verifying whether such an attacker can cause damage against a control system using the PLC. Their system, called SABOT, automatically maps the control instructions in a PLC to an adversary-provided specification of the target control system's behavior. This recovers sufficient semantics of the PLC's internal layout to instantiate arbitrary malicious controller code. They show that SABOT is successful in practice, although this only serves to amplify already existing concerns. The authors suggest that the perimeter security of ICS should be improved to ensure that PLCs and corporate networks are air-gapped. In addition, they propose that future PLCs should incorporate security mechanisms such as control-logic obfuscation.

A recent publication [324] focuses on a static analysis of smart grid device software and firmware with the goal of detecting vulnerabilities automatically. The approach has been implemented as part of Oak Ridge National Laboratory's (ORNL's) existing testbed for smart meters.

6.5.3 Domestic Smart Meters

Given the relatively low barrier of access to domestic smart meters for research purposes, there have been several strong contribution to this topic.

[338] focuses on smart meters, where previous research has demonstrated that they allow to draw inferences about the activities of their users. Modern smart meters rely on wireless communication for remotely collecting usage data from electricity, gas, and water meters. This motivated the researchers to conduct a security and privacy analysis of wireless smart meters, which they found lacked basic security measures to ensure privacy, data integrity

[1] Examples of these exploits, although not necessarily the same used by the authors, are available at `http://scadahacker.blogspot.com`.

and authenticity. The meters that they examined broadcast their energy usage data over insecure networks every 30 seconds, although these broadcasts should only be received when the utility company performs their legitimate reads. The authors showed that this issue allows monitoring of energy usage from hundreds of homes in a neighborhood with modest technical effort, and demonstrated how these data allow the identification of unoccupied residences or people's routines. The authors conclude by recommending security remedies, including a solution based on defensive jamming that can be deployed more easily than upgrading the meters themselves. A more interesting defensive mechanism is proposed in [407]. The key concept is to use battery-based load hiding, where a battery is inserted as a power supply "buffer" between the (insecure) smart meter and home devices at strategic times, in order to hide appliance loads from smart meters. Although this approach has been proposed in the past, the authors demonstrated that it is susceptible to attacks that recover precise load change information. Their proposed approach differs fundamentally from previous work because it maximizes the error between the load demanded by a home and the external load seen by a smart meter, thus rendering precise load change recovery attacks difficult. Along a similar line, in [240] the authors propose a battery-recharging algorithm that renders the meter reading probabilistically independent of the actual power usage. In addition, the approach relies on stochastic dynamic programming to charges and discharges the battery in the optimal way to maximize savings in the energy cost.

With modern automated smart-meter reading and billing systems, electricity theft is also an issue that costs billions of dollars per year in many countries. In [270] the authors propose the first threat model for detecting electricity theft, and a learning-based statistical technique that combines this threat model with an outlier-detection algorithm to detect unexpected usage profiles. They evaluated their approach using real metering data and showed that electricity thieves indeed exhibit a recognizable profile.

Recently, smart meter security has also been tackled from an anomaly detection point of view. In [327] the authors studied a smart meter technology equipped with a trusted platform for the storage and communication of metering data. Despite these security features, the authors acknowledge the need for an embedded real-time anomaly detector that protects both the cyber and physical domains [383] from data manipulation, smart meter recalibration, reset and sleep attacks.

6.6 Research Gaps

Given the inherent interdisciplinarity of the CPSs ecosystem, research on security aspects of CPSs also requires a deeper knowledge of many different re-

search areas. Moreover, the high complexity and the high deployment costs of CPSs make scientific research very expensive, with a high access barrier. For instance, conducting experiments on security protection tools for power-grid ICSs in real-world conditions may be impossible. In contrast, obtaining samples of advanced malware families for experiments in the wild is straightforward. This, however, is changing, as some simulation platforms [191, 192, 290, 326, 359]—or, better, testbeds [31, 35, 36]—are being built by governments agencies (also in Europe [412]) to support research and (military) training. The main research targets that arise are to determine how accurately these systems can simulate the true operations of CIs and, more importantly, to test countermeasures under realistic conditions.

The causes of the threats against CIs are unknown or very uncertain. Apart from the many speculations, there is no strong evidence to confirm that attackers are nation states, secret services or actual cybercriminals with malicious purposes. The cause of this is twofold. Real-world attacks against CIs found the organization unprepared; thus, few or no data were collected that could be used to reconstruct the scenario. Even where data are available, attacks such as Stuxnet are extremely complex, such that they would require data collected from a multitude of (distributed) sources and actors. Clearly, this was not possible. This lack of data impacts the research community, which is left with malware samples, many guesses, and little strong evidence. This raises the research question regarding how to collect and disseminate such data through scientific repositories such as those proposed by previous consortia [62].

6.7 Example Problems

From the above analysis of the state of the art and research gaps, we can formulate the following research problems.

Designing and deploying honeypot systems in real-world ICSs to collect evidence and create datasets for experiments. SCADA honeynets have been proposed in the past [46], although the variety of SCADA implementations make it difficult to decode the TCP-encapsulated protocols. In addition, collecting data on the PLC layer is challenging. Many industries are admittedly leaving this layer unmonitored because of the difficulties in data collection. After the answer to the first part of this research problem, which consists of the design and implementation of a honeypot for the major vendors, the second part concerns the creation of a legal framework that regulates their deployment, operation and use for data collection purposes.

This research line has already drawn some attention. Indeed, some efforts have been made towards "finding out who is really attacking

your ICS infrastructure." The most notable example is described in a recent industry research paper by Trendmicro [398], who deployed a SCADA/ICS honeypot system that included dummy web servers mimicking the control panel of a water pressure station as well as real PLC devices exposed on the Internet with default login credentials, which act as traps by imitating the activities of a real production system. Thirty-nine attempts to access or alter unauthorized resources of the honeypot were discovered during less than a month of observation. The report mentions that "China accounted for the majority of the attack attempts at 35%, followed by the United States at 19% and Lao at 12%."

Evaluating the accuracy of current modeling and simulation tools and, possibly, design better simulation tools. There are plenty of SCADA/ICS/CI simulation tools, created to fill the gap that many researchers face when they need real devices to test their security mechanisms. It is unclear, however, how accurate these systems are and how much they adhere to the reality. Each study in this field has obviously justified the proposed approach. What is missing is a systematization effort, toward the creation of a framework that can be used to evaluate existing and future simulators. This framework will have to take into account the characteristics of real-world attacks: How well is a simulation tool able to emulate the behavior of a real-world attack such as Stuxnet?

Information correlation and attack scenario reconstruction. Intrusion detection research is one of the main consumers of the data collected by honeypot systems. In particular, as it happened in the past when intrusion detection research was rampant, the correlation of various sources of information is one of the most challenging research problems. In ICS/SCADA systems this problem is more difficult, due to the inherent interdisciplinarity of the area and to the variety of protocols and vendors involved. For instance, one of the questions that need answering is to what extent attacks perpetrated (and detected) on the TCP side of a SCADA network are visible also on the PLC side and, if that is the case, to what extent these are correlated.

7 Authentication and Authorization

Having achieved personalization at an unprecedented scale, current services offer specialized content according to their users' preferences. Using a web site or a mobile application, in the majority of cases, requires creating a user account and authenticating with the provided service at a later time. Authentication is carried out by providing the correct credentials, usually expressed in the form of a username and a text-based password. Password-based authentication is the de facto method of access control in web services as it is cheap and simple in principle. However, the way users choose and services manage passwords may expose them to attacks [130]. A simple password is more convenient for users to remember; however, a simple password or its permutation is also more likely to be included in a word dictionary used in guessing attacks. Even if users select complex passwords, security pitfalls in the way services manage authentication credentials could lead to leaks, often on a massive scale [26, 50, 59, 61].

A common mitigation of such leaks is storing the output of one-way hash functions instead of the password itself, although this is not a practice followed by everyone [29]. Nevertheless, modern hardware enables powerful password-cracking platforms [60, 234] that can reveal the input that generated a given password digest. Aided by a dictionary and following certain assumptions that optimize the process, such systems feed a large number of possible inputs to the hash function in a rapid fashion. Furthermore, passwords can be also obtained by malware and social engineering attacks such as phishing [153]. This problem is only exacerbated by the fact that users reuse passwords across services [189] which means that domino-like attacks could be carried out [52, 56, 126].

While researchers have argued that passwords are not by any means the most valuable asset in cybercrime [182], password theft can cause annoyance, financial damages, data loss, and loss of privacy [199, 212]. It comes as no surprise that there is a strong push to replace passwords [23, 42, 129]. Some mechanisms that offer an alternative to textual pass-

words include public-key mechanisms, such as Microsoft's CardSpace [122] and TLS client certificates [155], graphical passwords [103], and many more; unfortunately, none of the proposed alternatives has proven sufficiently enticing [108]. Two-factor authentication [77] is the most common way to complement password-based systems by requiring an additional password acquired through a secondary independent channel. Currently, high-value services, such as online banking services and e-mail providers, have deployed such solutions in the form of either hardware tokens or smart-phone applications. Besides the obvious overhead of such a system in terms of both cost and effort, a survey has shown that it can push users to choose weaker passwords [399]. Moreover, it does not scale as the services increase. Single-sign-on services, such as OpenID and Facebook Connect [277], offer the option of maintaining a single online identity protected by a single password, though which users may access third-party services. However, they present a single point of failure, do not change the users' habit of selecting weak passwords, may carry privacy-related risks, and can also suffer vulnerabilities themselves [391].

7.1 What Is the Problem?

Password-based authentication has changed little in the many decades it has been in use, and today it is more popular than ever, with countless web applications using passwords to authenticate their users. In brief, on first registering with a service, a user selects the username and password that will be used for authentication. The application stores the username in plain text, while it attaches a random prefix to the password, usually known as a *salt*, hashes the outcome using a cryptographic hash function such as SHA1 or SHA2, stores the hash output along with the salt in the database, and discards the plain-text password. The salt is prefixed to ensure that, even if a password is shared by multiple users, a different hash will be generated and stored in the database, and identical passwords cannot be identified. Most web services require that authenticating users send their username and password in plain text to the service, and authentication is performed by using the stored salt and transmitted password to reproduce a hash, and compare it with the password in store. Users could also authenticate without sending their plain-text password [248]; however, such mechanisms are less prevalent and a hash is still stored on the server. We should also note that there are cases where passwords are simply stored verbatim in the database [43].

Passwords can be stolen, either in their plain-text form, or hashed. Assuming that the device used to enter the password (e.g., a PC or smartphone) has not been compromised, and, thus, is not running malware than can capture user input, passwords can be obtained by monitoring unencrypted communications [81], tricking users to divulge them voluntarily (e.g., through phishing or

social engineering) [153], or, as we have frequently observed recently, through leaking the password database of services [26, 59]. Hardware advances have overcome the irreversibility of hash functions.

Moreover, users frequently reuse the same password at multiple web sites. In fact, Florencio et al. [181] studied the password habits of more than half a million users, and found that on average a password is shared with six other sites. Password reuse imposes a significant problem because it implies that an attacker cracking a single password can gain access to multiple sites and services for which the user holds an account. Password reuse also acts as a counterincentive for sites to use hash functions like `bcrypt`, as an attacker could target a less secure site—for instance one that saves passwords in plain—compromise its database, and use the obtained passwords in other services.

7.2 Who Is Going to Be Affected?

Users of all systems that implement text-based password authentication experience risks. If a system is compromised, then passwords may be leaked. We stress here that an attacker can expose a system's passwords without fully compromising it. For example, a successful SQL injection can reveal all passwords stored in a web site's database. Furthermore, users that recycle the same password in many different services potentially receive the security offered by the service providing the least security guarantees. If a weak service is compromised, then all the user information stored in services where the victim has registered with the same password is at risk.

7.3 What Is Expected to Happen?

An attacker who somehow obtains hashed passwords from the database of a web service has all the data needed to attempt to crack them. He holds the password hashes and knows the function as well as any salt used to generate them.

With this information in hand, an attacker can employ various methodologies to crack the passwords in the obtained database [234]. The simplest approach is by brute force. He can try every possible combination of valid characters, generate a hash, and check it against the values in the database. Obviously, this approach requires abundant processing cycles and time. Alternatively, the attacker can also use a pre-constructed dictionary containing potential passwords (offline dictionary attack). Using dictionaries can greatly speedup the cracking process, especially assuming that most users do not use strong passwords [107].

The parallelism of modern GPUs can greatly improve the speed of password cracking. The complexity of the hash function used greatly affects the amount of time required to crack a password hash. Recently, GPUs were able to crack

eight character-long, NTLM-encrypted passwords in just about five hours [60]. Yet another approach involves using cloud resources to quickly crack various types of hashes [9]. If a salt is not used, cracking can be accelerated further by employing rainbow tables [374].

Once passwords are cracked, the attacker can do any of the following:

- Access all information stored in the service by other users.
- Steal the identity of a victim or impersonate them.
- Incriminate the user by carrying out questionable activities using their profile.
- Escalate to more valuable assets. For example, accessing a victim's web e-mail might be sufficient for compromising their e-banking account.

7.4 What Is the Worst That Can Happen?

A cautionary tale that shows what can happen when multiple services are interconnected, all using weak authentication mechanisms, is the *epic hack* of Mat Honan, a reporter working for Wired.com. In his own words [212]:

> In the space of one hour, my entire digital life was destroyed. First my Google account was taken over, then deleted. Next my Twitter account was compromised, and used as a platform to broadcast racist and homophobic messages. And worst of all, my AppleID account was broken into, and my hackers used it to remotely erase all of the data on my iPhone, iPad, and MacBook.

However, although Mat Honan experienced the dramatic consequences of password theft, and eventually of losing control over his digital assets, he managed to sustain his quality levels along many critical dimensions. First, he didn't suffer from critical financial loss. Second, the health of his life was not in danger, and, third, he didn't face heavy incrimination, connecting him with illegal actions and eventually making him face, wrongly, the consequences of violating the law. Mat Honan may managed to escape from such serious dangers because he was already famous, and he was more of a victim of a bad prank, than a targeted attack.

Nevertheless, password theft and, eventually, identity theft can lead to very bad consequences along all the above three dimensions. First, it is reported that financial loss from identity theft is increasing [19]. Note that the report states that year by year less people are victims, but the loss is greater. Second, although there are no reported cases of identity theft that can lead to life-threatening situations, a lot of medical data are stored on-line, and password theft can lead to privacy leaks associated with the health's condition of a victim.

Finally, identity theft can lead to serious incrimination, since today a user's social profile and all activities connected with it can be strong evidence for certain violations. For example, during the last Olympic Games many athletes were expelled from the games for tweeting racially charged content [54].

7.5 State of the Art

7.5.1 Delegated Authentication and Authorization

Communication and resource sharing between web services is a desired ability that benefits both users and services. The most straightforward way for service A to exchange information about a user with service B is for that user to provide his credentials for the latter. This carries security risks ranging from the unrestricted access service A receives to the compromise of the user's credentials as he shares them with more and more services. For that matter, delegated authentication and authorization methods have been developed as an alternative to the users' providing their actual credentials to a service.

The OAuth 2.0 authorization framework [57] enables a third party to request access to a credential-restricted resource from its owner and receive that access without knowledge of the owner's credentials. For that to happen, the resource owner authenticates with the resource server, using his credentials, and obtains an access token which can be used in place of the owner's credentials for the restricted resource. Moreover, the owner is able to limit the token's capabilities so as to set a specific permission scope, lifetime and other attributes for that third party's actions. The OpenID 2.0 authentication standard [41] provides a way for an end user to prove ownership of a claimed identity to a third party. Its intended purpose is for users to log in to web services without registering for a new account as long as they already have a registered identity with an OpenID provider. Users visit a web service and attempt to log in simply by claiming an identity and specifying the OpenID provider that will verify their control over that identity. The web service indirectly, through the users' user-agent, requests and receives an assertion about their ownership of the claimed identity from the OpenID provider.

Facebook Connect [419] is an attempt by the social network to build on top of both OAuth and OpenID to produce an authentication and authorization framework combined with the social information and graph its users form. Other popular web parties such as Google [22] and Twitter [58] are doing the same. BrowserID [291,292] or Mozilla Persona is a single-sign-on mechanism that uses e-mail addresses to represent user identities. Users are able to claim an e-mail address as their identity as long as they can prove ownership. E-mail providers take up the role of providing proof to a web service, in a cryptography-secured manner, that a given user-agent, trying to log in to that service, has also managed to successfully authenticate itself to the e-mail

provider. One of the benefits over OpenID and Facebook Connect is that the identity provider (e-mail provider) does not find out which web service the user is trying to use. On the other hand, while the identity provider does not learn the relaying party, the relaying party learns the user's identity on the identity provider's service; i.e., his e-mail address. PseudoID [152] employs blind cryptographic signatures to eliminate this privacy concern. Moreover, while Facebook Connect and Google Login associate the user with a social profile and may share some of that information with the third-party web service, BrowserID does not. While BrowserID and Facebook Connect seem to eliminate the need for web services to maintain and manage the security credentials for their users, they also present single points of failure that, if abused could result in domino-like security failures. For instance, a user who has enabled Facebook Connect to log in to a plethora of web services, he only needs to manage the Facebook password. However, if the same password is also used for another service that does not support Facebook Connect, a potential leak from either Facebook or that service could allow an attacker access to all the services connected to that user's Facebook identity. Another example, is the case where security flaws in the single-sign-on system enable an attacker to access the victim's account in any of the services supporting such password-less login [391].

7.5.2 Password Cracking

Password cracking is not a new technique [344]. However, up until recently, the use of cryptographic hash functions in the way authentication information was handled by services appeared as an effective defense. This is no longer the case since modern CPUs/GPUs [60] can be combined to form powerful cracking platforms targeting password digests. Many of them, such as CloudCracker, are provided as an off-the-shelf paid service for the average user [9]. Even if strong cryptographic hash functions are used for keeping passwords safe, other properties of the system may be exploited to boost the process of cracking [12]. Research efforts have been directed towards suggesting harder-to-guess passwords and pass-phrases [234,322,396].

7.5.3 Adaptive Hashing

Adaptive cryptographic hash functions, such as `bcrypt` [323] and `scrypt` [313], have been proposed to address the increasing ease with which password hashes can be cracked. These hash functions can adapt to hardware evolution, by deliberately wasting resources - either computational or memory - during a hash validation. By employing such hash functions, a web site can slow down an attacker sufficiently in cracking a particular user password. However, this also requires that the service *invests* additional resources into generating

hashes and evaluating passwords. We should emphasize that, compared to brute-force attacks, these functions have much less effect on dictionary attacks.

7.6 Research Gaps

Text-based passwords are convenient and are already accepted by the majority of users. However, today we need stronger authentication mechanisms.

Rich Authentication. It is well known that for authenticating with a party we use *something we know* (i.e., a password), *something we have* (i.e., a token), or *something we are* (i.e., biometrics). There is an interesting gap between security and usability in current forms of authentication. We have strong authentication mechanisms, but it is hard to use them effectively because they are not convenient. On the other hand, there are certain mechanisms that are already accepted by users, such as text-based passwords or 4-digit PINs. Unfortunately, these provide low security guarantees. One research challenge is to invent new *rich* authentication mechanisms, variants or combinations of the currently existing ones, that provide better security without sacrificing convenience.

Service Decoupling. Services experience heavy interconnection, explicitly or implicitly. It is common practice to use an e-mail service for registering (or resetting the credentials) to another service. Social applications can also interfere with content delivered via third-party networks. For example, a Twitter account may post comments in the user's Facebook feed, if it is so configured. This service coupling provides new and dynamic functionality; however, it is security sensitive. An attacker needs only to compromise one service and may then take over many of the victim's valuable assets just by exploiting this service interconnection [126, 212]. It is challenging, from a research point of view, to identify all this interconnection, create taxonomies with current practices, study the ways current services interconnect with each other, and design new techniques for interconnecting services in a secure fashion.

7.7 Example Problems

Some interesting problems in this area include:

Factors in Authentication. It is common to combine multiple communication channels for providing stronger authentication, something commonly known as 2-factor authentication. As a quick example, consider a user authenticating with a service by giving a password and a code received by SMS. It is debatable which factors are more efficient from the usability perspective, while providing the most security guarantees.

Constrained-input Devices. The personal computer is changing decade by decade. From desktop PCs and laptops, we have evolved to smartphones, which are apparently not rich in input capabilities. In the near future it is quite possible that we will experience even more constrained devices in terms of text input [24]. It is challenging to explore new authentication capabilities for such devices, since it is harder for them to support the text-based password model, and we know in advance that these devices will undergo massive adaptation in the near future.

8 Security of Mobile Devices

In an era of explosive growth in the use of mobile communication devices, the need for a secure and reliable infrastructure has never been more apparent. Although connectivity is provided at large by WLAN networks, true mobility is only possible through the cellular infrastructure. Ericsson forecasts 80% of the world population will have WCDMA/HSPA (3G) coverage and 35% will benefit from an LTE (4G) connectivity by 2016. During the single month of September 2012, 11 million automobile were connected to the mobile network and recent forecasts refer to over 50 Exabytes of data being exchanged by mobile devices and smartphones have been recently reported [172].

To give just one example, with more than 500 million of activations reported in Q3 2012, Android mobile devices are becoming ubiquitous and trends show that the pace is unlikely to slow [267]. Android devices are extremely appealing: powerful, with a functional and easy-to-use user interface for accessing sensitive user and enterprise data, they can easily replace traditional computing devices, especially when information is consumed rather than produced.

Application marketplaces, such as Google Play and the Apple App Store, drive the entire economy of mobile applications. For instance, with more than 600,000 applications installed, Google Play has generated revenues of about 237M USD per year [161]. The prospect of such a fortune, combined with the quite unique Android ecosystem, with its high turnovers and access to sensitive data, has unfortunately also attracted the interests of cybercriminals, with the result that there is an alarming increase in the rate of malware strikes against Android devices. Breaches of users privacy (e.g., access to address books and GPS coordinates) [416], monetization through premium SMS and calls [416], and colluding malware to bypass 2-factor authentication schemes [150, 231] are all real threats rather than a fictional forecast. Recent studies back up such statements, reporting how mobile marketplaces have been abused to host malware or legitimate-seeming applications in which malicious components are embedded [414].

This clearly reflects a shift from an environment in which malware was developed for fun, to the current situation, where malware is distributed for

financial profit. Given the importance of the problem, significant research efforts have been invested in gaining a better understanding of the mobile malware phenomenon. However, given the rate at which mobile malware is growing, it seems we are still a long way from solving the problem, and we must hope we are not already of time.

It is important to point out that understanding is not a mere academic exercise: it is of paramount importance to acquire the knowledge necessary to characterize a specific threat, in order to devise novel, effective, and efficient techniques for detection and mitigation.

8.1 Who Is Going to Be Affected?

The consequences of infected mobile devices will affect all users alike. Smartphones have now become ubiquitous, and they are a constant presence in almost every household. However, we currently lack flexible and efficient policies to regulate private-to-enterprise bring-your-own-device (BYOD) contexts, just to give an example. How can we effectively implement evasion-resistant techniques for information leakage detection? How can we detect, mitigate, or contain unknown malicious behaviors?

8.2 What Is Expected to Happen?

Even in a non-BYOD scenario, the compromise of a smartphone can be catastrophic. Apart from breaches of user privacy (e.g., access to address books and GPS coordinates) [416], monetization through premium SMS and calls [416], and colluding malware to bypass 2-factor authentication schemes [150, 231], as noted above, this may also ultimately turn an infected smartphone into a real mobile bot, with serious consequences (see for instance Chapter 11) [8].

Although the mobile malware harvested recently on a major US cellular provider by the research community over a 3-month period in 2012 appears in a very limited number of devices (3,492 out of over 380 million—less than 0.0009% [253]), forecasts for 2013 are not looking good. According to Lookout 2013 Mobile Threat Predictions, *"[...] people will purchase more than 1.2 billion mobile devices, surpassing PCs as the most common Internet access device in the world. Mobile platforms will continue to expand at breakneck speed, as people are forecast to download over 70 billion mobile apps in 2014."* [261]. Globally, 18 million Android users are expected to face malware infection during 2013, with monetization through premium SMS and calls being the predominant revenue for cybercriminals. Moreover, during the same year mobile spam is expected to increase, turning into a serious threat vector.

8.3 What Is the Worst That Can Happen?

Losses can happen in almost any domain: financial, personal data [212], and intellectual property are the easiest to think of, but other attacks could also be potentially life-threatening. For instance, denial of services to avoid calling emergency numbers, malicious location-based services and leakage of GPS coordinates, which may enable traditional crime activities (as outlined above), may all be concrete attacks rather than fictional artifacts. In addition, any life-affecting device (e.g., cars, NFC-based insulin pumps), if improperly used and *fully* controlled by a smartphone (assuming read/write/exec accesses and no fallback safety checks) may threaten life itself or violate an individuals privacy.

8.4 State of the Art

To contribute to an understanding of the security problems affecting smartphones, La Polla et al. surveyed the related literature in the 2004–2011 period, highlighting threats, vulnerabilities, and attacks [246]. Despite the similarities, there are in fact a number of security-related differences between mobile devices and PCs (e.g., monetization through premium SMS and calls [416]) and they need to be dealt with specifically.

With a few exceptions focused on enhancing mobile OSes with state-of-the-art memory error protections [105], iOS privacy violation detection [162], and a recent detailed analysis of cellular networks [253], current research is mainly concerned with understanding, analyzing, and mitigating Android malware threats.

DroidScope [405] is a framework for creating dynamic analysis tools for Android malware that trades simplicity and efficiency for transparency. As an out-of-the-box approach, it instruments the Android emulator, but it may incur high overhead (for instance, when taint-tracking is enabled). DroidScope employs a 2-level virtual machine introspection (VMI) [187] to gather information about the system (i.e., OS-level and Android-specific behaviors) and exposes hooks and a set of APIs that enable the development of plugins to perform both fine and coarse-grained analyses (e.g., system call, single instruction tracing, and taint tracking). Unfortunately, DroidScope just offers a set of hooks that other analyses can build upon to intercept interesting events and does not perform any behavioral analysis *per se*.

Enck et al. presented TaintDroid [167], a framework to enable dynamic taint analysis of Android applications. TaintDroid's main goal is to track how sensitive information flows between the system and applications, or between applications, in order to automatically identify information leaks. Because of the complexity of Android, TaintDroid relies on different levels of instrumentation to perform its analyses. For example, to propagate taint

information through native methods and IPC, TaintDroid patches JNI call bridges and the Binder IPC library. TaintDroid is effective, as it allows tainting to propagate between many different levels, and efficient, as it does so with a very low overhead. Unfortunately, this comes at the expense of low resiliency and transparency: modifying internal Android components inevitably exposes TaintDroid to a series of detection and evasion techniques [121,341,355].

DroidBox is a dynamic in-the-box Android malware analyzer [372] that uses the custom instrumentation of the Android system and kernel to track a sample's behavior, relying on TaintDroid to perform taint tracking of sensitive information [167]. Building on TaintDroid and instrumenting Android's internal components makes DroidBox prone to the problems of in-the-box analyses: malware can detect and evade the analyses or, worse, even disable them.

Andrubis [7] is an extension to the Anubis dynamic malware analysis system to analyze Android malware [99,220]. According to its web site, it is mainly built on top of both TaintDroid [167] and DroidBox [372] and it thus shares their weaknesses (mainly due to operating "in-the-box").

CopperDroid performs automatic out-of-the-box dynamic behavioral analysis of Android malware [11,331]. To this end, CopperDroid presents a *unified* system call-centric analysis to characterize low-level OS-specific and high-level Android-specific behaviors, including IPC and RPC interactions—of paramount importance on Android. Based on the observation that such behaviors are all eventually achieved through the invocation of system calls, CopperDroid's VM-based dynamic system call-centric analysis is able to faithfully describe the behavior of Android malware whether it is initiated from Java, JNI or native code execution. Based on the observation that Android applications are inherently user-driven and feature a number of implicit but well-defined entry points, CopperDroid furthermore describes the design and implementation of a stimulation approach aimed at disclosing additional malware behaviors. The authors carried out an extensive evaluation of the system to assess its effectiveness on three different Android malware data sets: one of more than 1,200 samples belonging to 49 Android malware families (Android Malware Genome Project); one containing about 400 samples over 13 families (Contagio project); and a final one, previously unanalyzed, comprising more than 1,300 samples, provided by McAfee. Their experiments show that CopperDroid's unified system call-based analysis faithfully describes OS- and Android-specific behaviors, while a proper malware stimulation strategy (e.g., sending SMS, placing calls) successfully discloses additional behaviors in a non-negligible portion of the analyzed malware samples.

Google Bouncer [260], as its name suggests, is a service that "bounces" malicious applications off from the official Google Play (market). Little is known about it, except that it is a QEMU-based dynamic analysis framework.

All the other information come from reverse-engineering attempts [303] and it is thus hard to compare it to any other research-oriented approach.

DroidMOSS [414] relies on signatures for detecting malware in app markets. Similarly, DroidRanger [417] and JuxtApps [207] identify known mobile malware repackaged in different apps. Although quite successful, signature-based techniques limit the detection effectiveness only to known malware (and it is vulnerable to the adoption of reflection, native code, and obfuscation in general).

Enck et al. [168] reported on a study of Android permissions found in a large dataset of Google Play apps, aimed at understanding their security characteristics. Such an understanding is an interesting starting point to bootstrap the design of techniques that are able to enforce security policies [402] and avoid the installation of apps requesting a dangerous combination [169] or an overprivileged set of permissions [178, 312]. Although promising, the peculiarity of Android apps (e.g., a potential combination of Java and native code) can easily elude policy enforcement (when confined to protecting the Java API—as represented by the state-of-the-art) or collude to perform malicious actions while maintaining a legitimate-seeming appearance. This clearly calls for continuing research in this direction.

Aurasium [402] is an app rewriting framework (Java only) that enables dynamic and fine-grained policy enforcement of Android applications. Unfortunately, working at the application level exposes Aurasium to easy detection or evasion attacks by malicious Android applications. For example, regular applications can rely on native code to detect and disable hooks in the global offset table, even without privilege escalation exploits.

SmartDroid [413] makes use of hybrid analyses that statically identify paths leading to suspicious actions (e.g., accessing sensitive data) and dynamically determine UI elements that take the execution flow down paths identified by the static analysis. To this end, the authors instrument both the Android emulator and Android's internal components to infer which UI elements can trigger suspicious behaviors. In addition, they evaluate SmartDroid on a testbed of 7 different malware samples. Unfortunately, SmartDroid is vulnerable to obfuscation and reflection, which make it hard—if not impossible—to statically determine every possible execution path.

Anand et al. propose ACTEve [83], an algorithm that utilizes concolic execution to automatically generate input events for smartphone applications. ACTEve is fully automatic: it does not require a learning phase (such as capture-and-replay approaches) and uses novel techniques to prevent the path-explosion problem. Unfortunately, the average running time of ACTEve falls within the range of *hours*, which makes it ill-suited to automated large scale analyses or practical in-device detection.

Finally, Zhou et al. performed a very detailed and fine-grained static analysis of a large set of real-world Android malware, collected from August 2010 to October 2011 [416]. The authors released this data set as the Android Malware Genome Project [415].

8.5 Research Gaps

Recent research has shown that mobile security is still in its infancy and, although we can borrow and build on traditional (PC) malware research, the problem space is hard and quite diverse, which calls for solutions to be devised specifically for mobile devices. For instance, is sending an SMS a malicious action on its own? Hard to say; potentially, if this happens in the background.

Given the nature of the problem space (i.e., smartphones travel across network boundaries and store—on the same media—potentially sensitive personal and enterprise data), confidentiality (privacy) and integrity seem, once again, to be the most important property to guarantee. However, assuming we can tolerate financial and intellectual property losses to some extent, privacy (e.g., no unauthorized data disclosure) may be the foremost challenge to meet.

It is not hard to imagine mobile malware infections that collude with and abet traditional crimes more effectively. For instance, leaked GPS coordinates or malicious location-based services may indicate that someone has reached or left a specific location, enabling a number of traditional criminal-related activities (e.g., burglary, kidnapping, stalking, terrorist attacks). Information flow has long been studied as a way to track how sensitive data propagate throughout a program's execution, enabling integrity and confidentiality properties to be realized in a number of successful scenarios. Unfortunately, the very same technique has been shown to suffer from a slew of easy-to-deploy attacks and evasions when applied to the analysis analyze or containment of malicious infections [121, 341, 355]. It is clear that alternative and evasion-resistant solutions, which potentially retrofit traditional malware, need to be explored.

An interesting long-term research project would aim at exploring highly scalable technologies for efficient monitoring and analysis of security events that have the potential to compromise mobile devices. In particular, such objectives may be addressed by approaching three different steps within the mobile communication process: 1) data monitoring and analysis—dedicated to monitoring multi-source data, including mobile device security event detection and analysis, and network-based features generated by smartphone communication; 2) data aggregation and correlation—by adopting a synergy of contextual information from the terminals correlated with application behavioral profiles and honeypot events; and 3) management of mobile-related trustworthy indicators—that will use the latest malware classifications and gen-

erated anomaly alerts to feed an accurate mobile security dashboard that helps in the understanding and management of new mobile malware outbreaks.

8.6 Example Problems

Smartphones are a relatively new technology and, although they were originally designed with security principles in mind, they have been shown to be as vulnerable as traditional computing devices.

As mentioned above, effective and evasion-resistant ways to detect or contain information leaks would be of paramount importance, especially in a context where such devices have access to sensitive personal and enterprise data, while crossing the boundaries between a number of different networks. It is clear, in fact, that existing techniques, such as taint tracking, although effective in principle, are ill-suited when it comes to containing the effect of malicious computations; thus, novel solutions must be sought.

Another interesting research direction would be to explore hardware-supported virtualization solutions to physically separate processes in context-dependent scenarios, in order to enforce security policies automatically, without requiring users' approval.

An orthogonal research direction would be to explore instead the possibility of analyzing network traffic from network operators, which offers a unique observation point from which (malicious) behaviors may be inferred and security or containment policies enforced.

9 Legacy Systems

Joined into a single entity through an interconnection of information services, an information system is only as strong as its weakest link. This link may well be an obscure, older library or a program that is vulnerable to attacks. While security issues occur in all software, older programs can be especially vulnerable, as they were not designed with security in mind. And even if they were, the programmers probably did not have knowledge of the latest and greatest exploitation techniques.

9.1 What Is the Problem?

The research community has long recognized the security problem introduced by legacy systems, and multiple interesting solutions were developed. Despite all these efforts, attacks are reported daily [284, 301, 343]. The problems persist in the real world because the adopted solutions prove insufficient, whereas more powerful protection mechanisms are too slow for practical use, break compatibility with other programs, or require source code that is not available for legacy software.

Among the most dangerous software errors that are of primary importance in legacy software are memory corruption attacks in the C/C++ languages [143]. From clients to servers and from big iron to mobile phones—all have fallen victim. For the rest of this section, we will thus mainly focus on this particular vulnerability.

Memory corruption attacks have become more and more sophisticated. Typically, they arrive as data over a regular communication channel, e.g., the network, and trigger pre-existing low-level software vulnerabilities. In a classic exploit, a program receives input from an attacker, and stores it in an undersized buffer, e.g., an array. A buffer overflow (i.e., a particular type of memory corruption attacks) occurs when the program is able to write beyond the end of the buffer. When attackers successfully exploit such flaws, they usually gain control over

the execution of a program (a *control-diverting* attack), or modify some critical data (a *non-control-diverting* attack).

9.2 Who Is Going to Be Affected?

The consequences of memory corruption attacks that cannot be prevented or detected will affect all users. For years, we have witnessed attacks targeting home users, critical networks [86, 255], or even Iranian nuclear facilities [179, 180, 366]. Even if the vulnerable systems do not always run legacy software, the attackers often do not have access to the source code, so they proceed in exactly the same way: they perform an analysis of the binary executable to learn how to exploit a vulnerability.

9.3 What Is Expected to Happen?

Attackers exploit vulnerabilities in legacy software to take control of either a vulnerable process or the whole machine. In the latter case, they can for example easily download and run a password stealing application, a keylogger, or a bot of their choice. The consequences may vary from privacy breaching (e.g., in the case of a keylogger), via serious financial losses, to attacks taking control over critical networks.

9.4 What Is the Worst That Can Happen?

As we said above, exploitable vulnerabilities in legacy software can be used to perform arbitrarily malicious actions. It is difficult to say what other uses will be found for machines controlled by the attackers. In principle, depending on the target, massive loss of life is possible, but not likely.

9.5 State of the Art

In this section, we survey the most important solutions developed as responses to memory corruption attacks that can be deployed even if we do not have an access to the source code of the (possibly) vulnerable application.

Anti-virus software and *network intrusion detection systems (NIDS)* monitor executable files or network traffic, and frequently search for signatures, i.e., patterns distinguishing malicious attacks from benign data. However, polymorphic attacks, zero-day attacks, and data encryption, all render signature-based solutions limited.

Control-flow integrity (CFI) [76, 238] is designed to thwart control-diverting attacks. It ensures that the execution of a program does not stray from a restricted set of possibilities—it dictates that the execution must follow the path of a precomputed control-flow graph (CFG). The graph needs to maintain all possible valid executions of the program lest an innocent process be flagged

as an attack. To determine the CFG, one could employ either static or dynamic analysis. However, none is simple in practice—static analysis has an inherent difficulty in resolving indirect branch targets, while dynamic analysis often covers only a part of a program's execution paths. Several ways of approaching the problem have been proposed: for example a combination of static and dynamic analysis by Xu et al. [401], value set analysis presented by Balakrishnan et al. [90], and a framework proposed by Kinder et al. [237], which combines control and data flow analysis by means of abstract interpretation. The CFI policy is enforced at runtime, and a possible implementation may compare the target address of each control-flow transfer, i.e., each jump, call, or return, to a set of allowed destinations.

CFI does not detect non-control-diverting attacks, but it is a useful and cheap-to-enforce policy, which effectively stops the non-control-diverting ones. The mechanism realized by Abadi et al. [76] employs binary rewriting, and requires neither recompilation nor source-code access. The average performance overhead is 15%, with a maximum of 45%.

Runtime host solutions take advantage of the wealth of information present when a vulnerable application is running to protect against attacks. *Dynamic Taint Analysis (DTA)*, proposed by Denning et al. [149] and later implemented in TaintCheck [299], and a plethora of other systems [110,134,138,209,320,354], is one of the few techniques that protect legacy binaries against memory corruption attacks on control data. The technique is implemented by transparently modifying the runtime environment. In a nutshell, untrusted data from the network is tagged as *tainted*, and its propagation is tracked throughout a program execution. An alert is generated (only) if an exploit takes place, e.g., when the address of a function to be invoked is tainted (this never happens in a benign situation). The technique proves to be reliable and generate few, if any, false positives. However, it can slow down the protected application by an order of magnitude, and in practice, it is limited to non-production machines such as honeypots or malware analysis engines. Furthermore, DTA can usually detect only control-flow diverting attacks, so it does not defend against the non-control-diverting ones.

The above solutions are good at stopping control-flow diversions, but powerless against corruption of non-control data. As a response to this problem, BodyArmour [356,357] is a tool chain to bolt a layer of protection onto existing C binaries to shield them from state-of-the-art memory corruption attacks, including the non-control-diverting ones. It employs dynamic information flow tracking. First, it monitors the execution of a vulnerable application to understand the layout of memory, and unearth buffer locations and sizes. Later, it hardens the application so that buffer overflows are no longer possible. However, this technique is based on dynamic analysis, so it protects only those

parts of the program that were observed in the learning phase. This means that, if a function has not been executed at all, its vulnerabilities will go undetected.

We do not consider detection mechanisms such as *anomaly detection* or *behavior-based approaches*. Although great deal of research has investigated the applications of these techniques to detect attacks, reducing the number of false positives is still the core problem for these systems.

9.6 Research Gaps

We envision research in the area of legacy software in the following dimensions.

9.6.1 Attack Detection

As already stated, BodyArmour [357] is a tool chain to bolt a layer of protection on existing C binaries to shield them from memory corruption attacks. Since it is based on dynamic analysis, it also suffers from coverage issues - we can only analyze what we execute. Lack of coverage may cause BodyArmour to miss arrays and array accesses and thus lead to false negatives. Another popular, yet very different, approach to analyzing a binary is static analysis. Even though the method is less accurate than dynamic analysis, it offers full code coverage. Consequently, it might be interesting to explore a hybrid solution, which would marry BodyArmour to static protection approaches, such as WIT [80].

9.6.2 Search for Vulnerabilities and Crashes

The approaches discussed so far aim at attack detection at production time or in a honeypot. They are effective, but they do not remove the vulnerabilities themselves. Although it is better to crash than to allow exploitation, crashes are undesirable too. Thus, ideally we would like to find as many bugs as possible by means of fuzz testing even before deploying an application. Fuzzers feed programs invalid, unexpected, or random data to see if they crash or exhibit unexpected behavior. Information about inputs exploiting a security vulnerability allows system administrators to filter out offensive inputs, or if necessary, stop the application.

The most effective fuzzing technique today is whitebox fuzzing [118, 128, 195]. By means of symbolic execution, it exercises all possible execution paths through the program and thus uncovers all possible bugs—although it may take years to do so. Since symbolic execution scales poorly, a possible solution might be to focus first on functions/code fragments that *look* more vulnerable than others.

Previous research has shown that software complexity metrics collected from software artifacts are helpful in finding vulnerable code components [190, 300, 349, 418]. However, even though complexity metrics serve as useful in-

dicators, they are too generic, and they suffer from low precision or recall values. Moreover, most of the current approaches operate at the granularity of modules or files, which cannot be applied to legacy software for which we do not have the source code. As observed by Zimmermann et al. [418], we need metrics that exploit the unique characteristics of vulnerabilities, e.g., buffer overflows or integer overruns.

Summarizing, an important research question is how to evaluate the complexity of code fragments in existing binaries, so that we can focus the effective yet expensive symbolic execution on code that is more likely to have exploitable vulnerabilities.

9.7 Example Problems

There follows a list of example problems that could be addressed in research projects.

Combining static and dynamic analysis for vulnerability discovery. As we have discussed above, it is crucial to be able to detect attacks in legacy software. Dynamic analysis benefits from the wealth of information available while an application is running, while static analysis offers full code coverage. It might be beneficial to design and build a hybrid solution that exploits both approaches.

Searching for vulnerabilities in a haystack. A thorough analysis of the whole binary is extremely time consuming. The research question is whether we can pinpoint code fragments that have more potential for vulnerabilities/crashes, and should thus be scrutinized first. For example, in the case of buffer overflows, intuition suggests that convoluted pointer computations are hard for a programmer to follow. Thus, we should focus on code with complex pointer arithmetic instructions that accesses arrays first.

10 Usable Security

KEYS, LOCKS, AND CHAINS: security in computer systems is almost as old as computers themselves. Although the enforcement methods may have changed from physical security, which is still important, to a more conceptual and system-inherent form, the basic concept is still the same: How can a computer system be secured, so that only permitted personnel are able to use it? In the old days, it was simply a matter of giving the key to the building where the computer was housed to the right person. Even then, once an employee left the company, the key had to be recovered to ensure security. Neglecting this precaution, would create a potential security risk. Nowadays, the problems are conceptually similar. Instead of physical security keys, users are provided with passwords, PIN codes and access tokens so that they may enter restricted areas or access private information.

10.1 What Is the Problem?

Even though the majority of these problems are essentially the same as 30 years ago, one key factor has changed: complexity. What was previously a single device, is now a multitude of different accounts, Web gates and PIN codes. Additionally, connectivity is at an all-time high and is showing no tendency to slow down. This *Internet of things*, as it is called by experts, is a future for communication and computing devices that has already begun. Unfortunately, this progress not only opens up opportunities for development and technological advance, it also enables miscreants to conduct their daily business on a much broader scale. To mitigate present and future threats, a large research community is constantly developing solutions to block incoming attacks and ultimately prevent users from falling victim to various forms of attack. Commendable as these research results may be, they usually bear several important properties that often hinder their acceptance by normal users. Even protection mechanisms that are already well-established can suffer from the following problems:

- **Simplicity**: Useable security must be simple. A normal user cannot be willing to deal with the task of creating a security policy for accessing the

Internet, for example. Therefore, very complicated methods of securing a device are bound to be rejected by the masses.

- **Transparency**: Even security-aware users can not always deduce how a system works and where the possibilities for attacks arise. A good example here is a registered e-mail address that is used somewhere else without notification to identify a user. There are threats for some users that cannot be anticipated without a deeper knowledge of the underlying system.

- **Restrictiveness**: Most security solutions impose restrictions on their users. Passwords must be entered and memorized, device locks must be removed before using a device, firewalls prohibit unconfined network usage, etc. Users who see their devices as tools to do a job, which simply have to work properly, will gladly sacrifice security for convenience if given the choice. Therefore, the choice of which options to give the end-user for circumventing or re-defining security-critical aspects has to be a well-considered one.

To put it briefly, there are several reasons why a user may choose not to use the security mechanisms provided, preferring to go with a more convenient, unsecured solution instead. It is the researcher's responsibility to keep the target system safe anyway. .

10.2 Who Is Going to Be Affected?

While the scope of the problem is hard to define precisely, the potential victims of this threat are more easily identified. This type of threat specifically influences the everyday user of devices connected to the Internet. They simply cannot cope with the speed with which new technologies hit the market. Even experts such as network administrators, programmers and technically versatile individuals have a difficult time keeping up with new developments, let alone the possible threats they entail. Unfortunately, the ordinary end-user makes up the vast majority of customers dealing with (personal) computers. Therefore, the target community is one of the largest imaginable; it essentially comprises the whole Internet.

10.3 What Is Expected to Happen?

The effects of the previously discussed development are already visible. More and more users fail to take precautions because they negatively impact their workflow. As a result, these systems are prone to various attacks, ranging from stolen passwords and account data to infected machines that do their operator's bidding. If this trend continues, it will be virtually impossible to

create secure systems without isolating them from external influences—as Apple does with the App Store, for instance. Even then, a reluctant user who ignores updates, for example, can still fall victim to various attacks. It is, therefore, the research community's task to consider how a possible solution to a given problem might be easily adopted by the end-user.

10.4 What Is the Worst That Can Happen?

To take this thought even further, let us consider the two most extreme cases. In a hypothetic scenario, security precautions are so complicated to utilize, that no ordinary user is actually using them. Instead, every online action is considered a public process. Consequently, the Internet for ordinary users would lose a lot of potential and decrease to a pure information-retrieval facility. Netbanking, online booking, or even private messaging would not be possible anymore.
In the other extreme case, users would be forced to use all available security mechanisms, for instance by a security policy. The result would essentially be the same. With more restrictions and access control, usability decreases to a point where it is no longer feasible to even use a device. An example could be online banking, that can only be used from a single IP address over a VPN tunnel, where passwords have to be changed with every login and one-time security tokens are sent to a mobile device. While this method would certainly be more secure than a conventional login/password, it would cause most users to fall back to conventional manual wire transfers.
In reality, such an extreme case will hopefully never occur. Still, there is a natural balance between usable security mechanisms and convenience. This balance also exists in corporate systems, where security experts are confronted with the task to decide where to put the bar between user restrictions and security.

10.5 State of the Art

So, what is the situation today? This question can obviously not be answered with a simple *good* or *bad*. In 2005, Gutmann et al. made a very astute observation: a little more complexity is acceptable for a fair offering in value [202]. It ultimately comes down to exactly this question. A security mechanism, if not enforced by the system itself, will only be accepted if it offers a certain degree of added value. Usually, a trade-off between usability and security has to be found. A trade-off, however, is always a compromise [112]. The result is a system that is more complicated than one without security, but still less secure than it could be. Although acceptable, such a condition is certainly not desirable. The optimal case would be a consensus between both categories.

A perfect example of a consensus of security and usability is the now widespread application of mobile transaction-authentication numbers (or mTANs) in netbanking applications. From the consumers' perspective, this method increased usability since it only requires them to have a mobile phone at hand instead of a printed list of transaction numbers that can be lost, incorrectly maintained, invalidated, etc. From the bank's perspective, the security gain from introducing a second channel for each transaction is also enormous. In fact, the damage caused by phishing and man-in-the-middle attacks on online banking sites has been low compared to pre-mTAN systems. Recently, the Zeus Virus [113] and its mobile counterpart Zitmo have successfully proven that no security facility is impenetrable [226]], successfully infiltrating the two-channel security infrastructure represented by the mTAN facility.

Unfortunately, such a consensus is not always possible. Here, passwords are a good example. Most users, even inexperienced ones, know that using the same credentials for different Web sites and accounts is not a good idea. Remembering dozens of different passwords, or even creating an algorithm to derive the password from the target web site is tedious and rarely done. As a result, most users utilize from one to four different passwords for their accounts, resulting in multiple compromised accounts if a single attack is successful.

When iterating though these problems, the reader might think that no progress has been made in terms of usable security. In fact, there have been more or less successful initiatives to create a homogenous security and authentication environment. One of the most prominent examples is OAuth [57], a protocol for secure and even transient authentication among different applications. Even though the protocol has been widely adopted, it still requires developers to adhere to it's standard when developing their solution. And this requirement is not always easy to meet. Besides, once advanced to Version 2.0, the main contributor to the protocol, Eran Hammer, decided to leave the initiative, and even requested that his name be removed from related documents.

In general, researchers tend to ignore the usability aspects of their work, just the way engineers and programmers tend to ignore security [202] as such. The reasons are very similar. It is hard enough to get a project to work properly. Once it is operational, the engineer is told by his supervisor to include more features, while the Ph.D. student is told to find a solution for the next problem. And that happens despite a multitude of attempts to design systems, that bridge this gap between security and usability [329].

10.6 Research Gaps

Since the current state-of-the-art for usable security is hard to depict, defining a concrete research gap is equally hard. The underlying problem can, however, be reduced to some more specific issues.

10.6.1 Security Design Principles

Currently, there is no such thing as usability guidelines for security researchers. A newly developed method or prototype needs to function in order to be recognized. As already mentioned above, developers and researchers alike, simply lack the incentive to further their development in this direction. In some cases it does not even make sense. A new method for protocol reverse engineering will certainly not profit from a fancy interface. A new method for mitigating injection attacks on web browsers, however, would profit from a seamless, transparent plugin with a simple option to enable it, instead of a complicated setup procedure, with options that even experts have a hard time understanding. It would therefore be highly profitable to have a basic collection of do's and dont's for user interaction.

10.6.2 Field Study: Usability

Another field worth investigating is the usability of already existing solutions. This problem usually arises when someone decides to create a solution that is partially or completely based on existing solutions. In order not to invent the wheel from the beginning, it makes sense to build upon previously researched and published techniques. The drawback of this method is that the researcher often has to deal with prototype implementations or tailored solutions. A field study of the most recent developments would be of immense help in deciding whether and how it is possible to use these approaches.

10.6.3 Collateral Feasibility Considerations

Finally, the most important thing to do is incorporate usability decisions in the development process right from the beginning. This is something that cannot be done in a thesis but only by raising the awareness of people designing security solutions. Most companies are forced to do this at some level anyway, otherwise they will simply lose customers. The same does not hold true for researchers, as they are not required to incorporate feasibility studies, or rewarded if they do so. By raising community awareness about the importance of usability, however, this picture may change. The ultimate goal is still to protect users from attacks by providing smart solutions that they can use.

10.7 Example Problems

A multitude of examples exist for this specific topic. However, the problem is best depicted by the previously mentioned incident that happened to Matt Honan, a renowned author for the magazine Wired [212]. He fell victim to an attack that ultimately led to his Amazon, Apple, Google and Twitter account being compromised, and his Apple devices (iPhone, iPad and MacBook) being wiped. Interestingly, the only direct fault of the victim was his casual approach to backups, a fault shared by many users. The incident itself was possible because the different accounts were chained together in some way or another. Some had alternate e-mail accounts as backup authentication, others showed the last few digits of a credit card, while some were simply compromised by calling support and asking for a password reset. Inconvenient as the incident may be, it clearly illustrates an example of unusable security along well-placed examples for social engineering.

10.7.1 Authentication

Authentication played a major role in the above attack. One security guideline is not to use two vital e-mail accounts for a two-factor authentication. If one should get compromised, it is quite possible that the other will be lost as well. Instead, the backup e-mail should be used only once, or on a different system (e.g., SMS two factor authentication) altogether. From the user's perspective, however, it makes sense to use the same e-mail over and over again. No ordinary user is able to create e-mail aliases as needed and remember them afterwards. The same is true for passwords. If the same password is used on multiple platforms (e.g., Google, Facebook), it is easy to compromise both once the password is somehow derived. It is quite easy to depict the problem. Solving it, on the other hand, is a completely different story. The sensible thing would be not to allow secondary e-mail accounts as user verification. To be effective, however, such a guideline has to be enforced throughout different platforms, and that is something no one can guarantee. Alternatively, a new method could be devised that ensures that daisy-chaining accounts together is not possible, while users still have the possibility of retrieving their lost or forgotten credentials. This topic could serve as the foundation of a research thesis.

10.7.2 Backup

People are constantly told how important it is to backup their data. But who has terabytes of external storage lying around? And if so, how often does the ordinary user bother to actually create a backup? The answer is certainly: "Not often enough." Devising a safe, cheap but still usable form of backup would be

of value for many users. Again, this narrow field would certainly be suitable for consideration in the course of a thesis.

10.7.3 Sensitive Information

Part of the above incident was caused by sensitive information (in this case credit card numbers) being shown on Amazon's account detail. The last four digits are deemed insensitive by Amazon. However, Apple uses just these four numbers to prove one's identity. This connection is simply impossible to see for an ordinary user. Thus, information propagation through different, secured systems is another field where research could attempt to create transparent solutions.

Summing up, the threat from unusable security may not be a direct, immediate one, but it is there nevertheless. As security researchers, we are therefore obliged to develop our systems, not only with the basic concept in mind, but with a broader view that also considers the users who actually have to deal with it. Several of these shortcomings are intertwined with sociological and psychological aspects, calling for interdisciplinary research to create usable solutions.

11 The Botnet that Would not Die

Linked through malware, botnets are cyber infrastructures consisting of hundreds, thousands, or even millions of hosts that are all under the control of criminals. Botnets are responsible for most of the illegal activities on the Internet today: spam, denial of service attacks, theft of sensitive information (passwords, banking details and intellectual property), and spread of further malware. Moreover, botnets are the workhorses in high-profile incidents such as the state-sponsored Stuxnet attack on the uranium enrichment facility in Iran.

Despite the alarmist headlines in the popular press (sometimes echoed by researchers) about the highly advanced botnets that supposedly threaten the very Internet itself—and everything connected to it—so far it has been relatively simple to take them down.

This is about to change.

11.1 What Is the Problem?

In this chapter, we will argue that botnets are becoming so resilient that very soon, they cannot be taken down using conventional means (e.g., using a sinkhole, or taking down a few servers). In all probability, the only realistic way to take down such botnets will be to resort to what is known as "hacking back:" abusing vulnerabilities in the malware to compromise the botnet from the inside out. However, as doing so involves actively executing code on other people's machine (a criminal offense in most countries), this is something for which our legal system is not at all prepared. It is also hugely unpopular.

In this chapter, we will describe the trends towards more resilience in modern botnets. Moreover, we will back up our arguments with data from real botnets—to convince the reader we are not yet another group of researchers crying wolf.

11.1.1 P2P Botnets

The most common type of architecture for existing botnets is still based on a central Command-and-Control (C&C) server. Consequently, these C&C servers

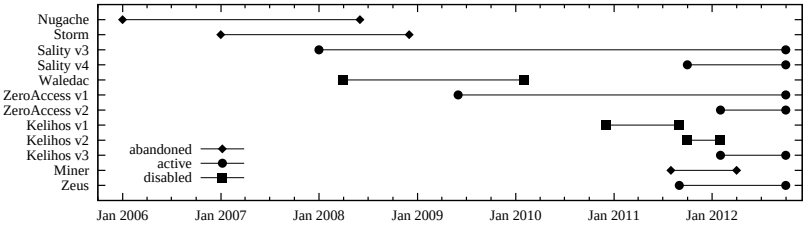

Figure 11.1: Lifetimes of botnet variants. Note that Sality has been up since 2007.

have received much attention from security researchers and law enforcement in takedown attempts [157,363]. In response, botnet controllers (botmasters) have designed and implemented new architectures to make their botnets more resilient. Some botnets use fast-flux DNS, which relies on a large pool of IPs belonging to compromised systems to mask out the actual address of an attacker-controlled mothership that delivers malicious content or runs scam campaigns [296,309].

In addition, attackers have implemented domain generation algorithms (DGAs) to generate pseudo-random domain names used for C&C dynamically (e.g., depending on seed values such as the current date/time and Twitter trends) [85]. For instance, the Zeus DGA currently generates a thousand domains a day.

However resilient such botnets have become, they have not stopped security researchers and law enforcement from taking them down. This is not the case for a new breed of botnets, based on peer-to-peer (P2P) technology, that appear to have been designed with resilience in mind.

In a P2P botnet, bots connect to other bots to exchange C&C traffic, eliminating the need for centralized servers. As a result, P2P botnets cannot be disrupted using the traditional approach of attacking critical centralized infrastructure. Figure 11.1 shows the lifespans of twelve different botnets based on P2P technology. Observe that ZeroAccess has been up since 2009. Incredibly, the Sality botnet which counts about a million nodes has been operational since 2007. In 2007, George W. Bush was still in the White House, nobody had heard about Stuxnet, and Nokia still reigned supreme in the mobile phone market!

To be sure, researchers did manage to take down several P2P botnets in the past. The Storm and Waledac botnets were probably the most famous of these [211,362]. Thus, P2P by itself does not provide resilience. The point is that modern botnets explicitly incorporate resilience in their design, with fallback C&C channels (often based on DGA recovery), heavy encryption, signed

messages, slow peerlist exchanges, and reputation-based peer replacement. Moreover, unlike earlier P2P botnets like Storm and Waledac, almost all new P2P botnets use an unstructured P2P infrastructure—which turns out to be much harder to take down.

11.1.2 Resilience

We will distinguish between different kinds of resilience for P2P networks:

Intelligence gathering resilience: For many attacks on (and mitigation techniques for) P2P botnets, we need to find out who is infected and perhaps even the topology of the P2P network.

Modern botnets are made resilient against intelligence gathering for instance by keeping most bots behind NATs and firewalls (making them impossible to crawl), not using explicit identifiers (making it impossible to estimate the exact number of infected machines, as IP addresses are both reused and renewed), not exchanging peers frequently and not exchanging many peers to begin with (making crawling very slow), etc.

All these techniques are employed even today by botnets like Sality and Zeus. For instance, Sality bots exchange peers once every 40 minutes and Zeus every 30 minutes. Moreover, Sality exchanges only a single peer at a time. A future botnet may be slower still. If you are too slow in mapping out the bots, the churn in the P2P network will have made your numbers obsolete almost immediately. This limits the measures that you can take based on these numbers. For instance, placing many peers in quarantine unnecessarily will probably not be acceptable.

Disruption resilience: well-known attacks to disrupt P2P botnets include *sinkholing* (where all bots are redirected to an attacker-controlled machine called a *sinkhole*), and *command poisoning* (where we distribute spoofed commands to other bots and/or transmit invalid messages). A simple measure against command poisoning is to protect the command using signatures and nonces (e.g., Kelihos and Sality v4 use RSA2048, while ZeroAccess uses RSA1024). Typically, encryption is used to prevent them from being visible in the network.

Moreover, to sinkhole an unstructured P2P botnet, it is crucial that we are able to poison other peers' peer lists. Modern botnets spend considerable effort to prevent this. For instance, by not exchanging many peers (e.g., Sality exchanges only one peer at a time), not replacing peers with a high reputation (Sality), replacing peers in a non-predictable manner (e.g., Sality picks peers at random, rather than recent ones, like Kelihos and Nugache, or peers that are close like Zeus and Storm), by providing a backup C&C channels (such as Zeus's DGA-based channel, or Kelihos'

fast-flux channel), and by occasionally cleaning up the peer-list through the backup channel (Zeus). Several modern, unstructured P2P botnets are no longer susceptible against Sybil and Eclipse attacks.

To make matters even worse, advanced bot software like Zeus is extremely stealthy. The probability of an AV scanner detecting the malware is not very high. Given the wealth of resilience measures already available in active botnets, and the incredibly long lifespans of some of these infrastructures, we anticipate that very soon, there will be botnets that we cannot take down using sinkholing, that will be extremely hard to crawl or measure (by the time you have charted a significant percentage of the botnet, the churn will have made your numbers obsolete), and that are not susceptible to spoofed commands.

11.2 Who Is Going to Be Affected?

The consequences of botnets that cannot be taken down within the boundaries of the law (without breaking into other people's machines), and that cannot even be easily mapped will affect all users. Unless we adopt other, more drastic means, to disrupt the botnet, we will have to take into account a permanent malicious infrastructure.

The botnets can be active continuously or lie dormant until they receive an activation signal. The lifespan of the malicious software is measured in years, even today. There is no indication that this will change.

While individual infected computers will be cleaned, the infrastructure itself will remain intact.

11.3 What Is Expected to Happen?

The beauty of a botnet, from the perspective of a cybercriminal, is that it is so versatile. It can be turned into any kind of cyber weapon: to steal, to disrupt, to damage, to extort, and even (in the case of Stuxnet) to physically destroy.

The primary and most immediate consequences are and will be financial. Every month, the banking sector in several EU member states is losing tens of millions of euros and the money lost in cleaning up infected machines is likewise significant. The most prominent banking trojan today is the highly resilient P2P Zeus. This botnet has been operational for almost two years now and has so far withstood all attempts to take it down.

Other predictable consequences include password stealing, denial of service attacks, ransomware, etc. However, there is no saying what other uses will be found for these malicious infrastructures.

11.4 What Is the Worst That Can Happen?

Botnets can be used to manipulate critical elements of the economy, such as the stock market, with possibly dire consequences. Moreover, they represent a constant threat to other (and practically all) infrastructures in advanced nations. For instance, they may provide hostile elements with a way to disrupt a society's infrastructure in unprecedented ways—e.g., in case of conflict. In principle, massive loss of life is possible, but not likely.

11.5 State of the Art

One of the most relevant studies to date is the analysis of P2P botnets by Rossow et al. [337]. In this paper, the researchers analyze specifically the resilience of modern P2P botnets against various kinds of attack. Moreover, they provide a formal model to describe both the botnets and the attacks.

We have discussed the resilience of several current and past P2P botnets. For a full discussion of each of the botnets, we refer the interested reader to malware analysis reports [113,116,175,251,276,351,400]. We have used results from these works to aid our manual code analyses, although in most cases the P2P resilience was yet undocumented.

As noted earlier, a few examples of enumeration and takedown operations against past P2P botnets exist. For instance, Holz et al. performed an early crawl of the Storm botnet, and also discussed some general resilience aspects of structured P2P botnets [211]. The sinkholing results of Stock et al. against Waledac represent the first successful attack against an unstructured P2P botnet [362]. Sinclair et al. have described the vulnerabilities of Waledac in detail [351]. The attacks against previous variants of Kelihos are also examples of recent sinkholing successes against unstructured P2P botnets [397].

The problem of crawling P2P botnets was first addressed by Kanich et al., based on lessons learned while crawling the Storm botnet [229]. An alternative concept to enumerate infected hosts (included NATed hosts) in structured P2P botnets was proposed by Kang et al. [228]. Their method involves the introduction of many fake nodes (sensors) into the target structured botnet. These sensors find infected hosts by monitoring search requests from bots looking for commands.

In several previous works, graph models have been used to describe network structures. Holme et al. used graph models to study the response of complex networks to several attacks [210]. The first application of random graphs, small world structures, and scale-free networks in the context of botnets was given by Dagon et al. [144]. Davis et al. used graph simulations to analyze the impact of bot disinfections on the communication effectiveness of P2P botnets [147]. Recently Yen and Reiter discussed the role of assorta-

tive mixing in P2P botnets and its consequences for network resilience and recovery [409].

To establish an idea of the threats we may expect from future P2P botnets, several researchers have designed their own theoretical highly resilient P2P botnets [213, 295, 360, 403, 404]. We are currently not aware of any existing P2P botnets based on ideas from these academic proposals.

11.6 Research Gaps

We envision Research in the area along the following dimensions:

11.6.1 Prudent Counter-attacks

Assuming the old ways no longer work in taking down botnets, what are the new ways? Can we alert users that they may be infected without becoming too intrusive? Are there safe ways to penetrate other people's computers and remove infections? This direction of research is not very popular today, as it represents what are known as offensive techniques. Most research departments eschew such research. We believe that we need a better understanding of what the options are.

11.6.2 Legislation

Currently, most countries lack a legal framework for dealing with these new advanced botnets. We have no guidelines as to how and when we can take more invasive measures against resilient malicious infrastructures. Nor is there clarity as to who should do it. And there is even less clarity when it comes to striking back at machines that are located in other countries (assuming you can even tell). We need research into the desirability of such measures, the boundaries for such measures, etc.

11.7 Example Problems

Tangible example problems might include:

Legal boundaries for hacking back. Can we provide clear and intelligible legislation that clarifies under what circumstances the government is allowed to strike back at botnets? Which computers is it allowed to attack—just the ones in its own country or may borders be crossed if need be (and if so, under what circumstances)?

Poisoned fruit. Rather than taking the P2P botnets down, can we disrupt their efficiency sufficiently to make them less interesting for attackers? For instance, can we inject an overwhelming amount of fake data, so that it becomes hard for the bot masters to extract the useful information?

12 Malware

Most users and administrators of computer systems configure their devices by installing software of their choice according to their needs. Often, however, not all software running on a device is vetted by its owner. Malware, short for malicious software, is an umbrella term referring to software that gets installed and operates against a user's will, usually for the benefit of a third party. Categorized depending on properties such as the malware's infection and propagation strategy, stealthiness, and purpose, common types of malware include viruses, worms, spyware, rootkits, keyloggers, backdoors, trojans, ransomware, and others [368].

Viruses usually infect executable files or documents, and require some form of human intervention in order to spread, such as plugging in an infected USB flash drive (or, in older times, inserting a diskette), being tricked into clicking on a malicious URL or attachment, or intentionally installing a malicious program disguised as (or contained in) a legitimate-looking application. In contrast, worms are autonomous, self-replicating programs that spread across the network by exploiting defects in widely-used software running on victim hosts. Other types of malware can be installed as a result of direct unauthorized access to a computing device, manual intrusions (often involving some form of social engineering), or automated exploitation by malicious websites or documents.

Historically, early viruses and worms were usually the outcome of experimentation and curiosity. Most of them were harmless, although they often unintentionally resulted in significant service disruption [165]. Fast forward a couple of decades, when organized cybercriminals develop sophisticated malware with the aim of illegal financial gain, while governments employ malware for gathering intelligence or even tactical operations (as was the case with the Stuxnet worm, discussed in Section 6.3).

12.1 What Is the Problem?

The rise in the number of malware variants continues at a steady pace. Indicatively, McAfee reports a growth in the number of new malware samples of about 8–12 million per quarter for 2012, while as of April 2013 they have more

than 128 million malware samples in their database [73]. Symantec reports that in 2012, one in 291 emails contained some form of malware [75].

At the same time, the increasing professionalism of cyber criminals makes defending against sophisticated malware increasingly hard. Once sophisticated tricks of the most skilled virus authors, advanced evasion techniques like code obfuscation, packing, and polymorphism are now the norm in most instances of malicious code. Using polymorphism, the malware is mutated so that each instance acquires a unique byte pattern, thereby making signature extraction for the whole breed infeasible. As the number of new vulnerabilities and malware variants grows at a frantic pace, detection approaches based on threat signatures, which are employed by most virus scanners and intrusion detection systems, cannot cope with the vast number of new malicious code variants [302].

12.2 Who Is Going to Be Affected?

Any computing device of sufficient capabilities can potentially be infected by malware. Besides personal computers and servers, the traditional targets of malware, mobile phones and tablets have recently started being plagued by malicious applications. Indicatively, McAfee reports that the growth in the number of mobile malware threats almost doubles every quarter, with 95% of the total number of samples in their database arriving in 2012 [73]. Computers and mobile devices, however, are not the only target. Malware can infect routers [141], phones [139], printers [140], gaming consoles [236], cars [123], and essentially any programmable computing device. As discussed in Chapter 6, industrial systems are often exposed to various threats, including malware infection, while malware managed to creep even into the International Space Station [64].

12.3 What Is Expected to Happen?

Practice has shown that malware authors continually try to devise new ways of evading existing detection systems, improve the stealthiness of their malicious code, and expand their reach to as many systems as possible. This is evident in several recent trends, including the proliferation of signed malware, server-side polymorphism, and the significant increase in the number of malware samples for mobile devices and typically less-targeted operating systems, such as Mac OS X [73].

Malware that has been digitally signed using a trusted certificate is capable of infecting even systems with strict configurations that allow the installation of software only from trusted sources. In recent incidents, malware authors managed to steal digital certificates from reputable software companies, which they then used to sign their malware binaries. Server-side polymorphism is

another recent technique that malware authors employ to render signature-based antivirus protection ineffective. By dynamically generating a different instance of the malware binary at the server, each victim is infected by a unique version of the malware that is unlikely to be encountered in the future. By hiding the logic of the polymorphic engine at the server, malware analysts have also a harder time identifying common patterns that could be used for detection. Other techniques used for hindering analysis and detection including anti-debugging tricks, VM detection, dormant functionality triggered by time or other events, memory-resident code, and advanced code obfuscation and metamorphism, are also expected to be used increasingly in future malware strands.

12.4 What Is the Worst That Can Happen?

Depending on the author's intent, malware can potentially have devastating consequences. Typically, upon infection, malware attempts to steal every bit of private information from the victim's device, including credit card numbers, personal files, and access credentials to web-banking, web-mail, or social networking websites. Infected computers also usually become "bots" in the attacker's network of compromised hosts. The threat of such botnets is extensively discussed in Chapter 11. Such networks of infected computers are essentially the infrastructure that allows cybercriminals to conduct a wide range of illegal activities, such as sending spam e-mails, launching denial of service attacks, hosting web sites for phishing, seeding malware, or publishing illegal material, and naturally, for probing and compromising other hosts.

A worrisome prospect is the rise of ransomware [75], which encrypts important user documents or even locks the user's computer completely. To restore their files, users are asked to pay a ransom, in return for the decryption key. Depending on the malware's sophistication and the type of encryption used, cracking the encryption through other means can be infeasible. If the motive is not financial, catastrophic malware can irreversibly erase data (even backup files, if those are reachable through the same internal network), and destroy crucial system components, such as the BIOS or other device firmware, causing severe damage.

12.5 State of the Art

In parallel with the development of cybercrime into a large underground economy driven by financial gain, malicious software has changed deeply. Originally, malicious software was mostly simple self-propagating code crafted primarily in low-level languages and with limited code reuse. Today, malicious software has turned into an industry that provides the tools that cybercriminals use to run their business [335]. Like legitimate software, malware is

equipped with auto-update functionality that allows malware operators to deploy arbitrary code to the infected hosts. New malware versions are frequently developed and deployed; Panda Labs observed 73,000 new malware samples per day in 2011 [67]. Clearly, the majority of these samples are not really new software but rather repacks or incremental updates of previous malware. Malware authors update their code in an endless arms race against security countermeasures such as anti-virus engines and spam filters. Furthermore, to succeed in a crowded, competitive market they innovate, refining the malware to better support cybercriminals' modus operandi or to find new ways to profit at the expense of their victims. Understanding how malware is updated over time by its authors is thus an interesting and challenging research problem with practical applications. Previous work has focused on constructing the phylogeny of malware [208, 230]. However, quantifying the differences between versions can provide an indication of the development effort behind this industry, over the observation period. To provide deeper insight into malicious software and its development, one needs to go a step further and identify how the changes between malware versions relate to the functionality of the malware. This is the main challenge in today's research. By utilizing techniques that combine dynamic and static code analysis to identify the component of a malware binary that is responsible for each behavior observed in a malware execution, the evolution of each component across malware versions can be measured. By comparing subsequent malware versions, code that is shared with previous versions and code that was added or removed can be identified. From the system-level activity, high-level behavior, such as downloading and executing a binary or harvesting email addresses, can be inferred. Aside from refining existing malware and introducing new techniques to evolve it, a certain trend towards propagating malware to new platforms is certainly apparent. However, to date, Windows systems are still the main target of malware attacks. Even so, samples have started to trickle down to other operating systems, such as Android or Mac OS. Given the growth of these markets, more sophisticated forms of malicious code can be expected in the near future.

12.6 Research Gaps

The increasing sophistication of malware has exposed limitations in existing virus scanners and malware detectors. Signature-based approaches cannot keep up with malware variants that employ packing and polymorphism, necessitating more advanced malicious code scanning and analysis techniques. Approaches based on runtime behavioral profiling and detection are a promising step, and behavioral heuristics are supported to some extent by current antivirus systems, but usually come as extra features not enabled by default due to their increased runtime overhead and proneness to false alarms. Al-

ternative software distribution schemes based on whitelisting or strict vetting, such as the one followed by Apple's App Store, reduce the chances of infection, but also limit the options of users to only those programs that have been vetted. Attackers that manage to slip through the vetting process or even steal "clean" developer keys may have increased potential for successful malware distribution.

Automated malware analysis systems face significant challenges due to the increasing rate of new samples that must be analyzed on a daily basis, and the need for more complex analysis for non-trivial samples. Given a certain malware analysis infrastructure, more samples must be processed in a time unit and more cycles must be spent per sample. Factors that drive up the analysis cost include stealthy malware that uses polymorphism and metamorphism, anti-debugging and VM-detection techniques, dormant functionality, and environment-dependent malware.

12.7 Example Problems

Problems caused by malware are extremely common. In fact, most of today's attack scenarios involve some sort of malware as an enabler. Some examples include:

Botnets. They are probably the best example of how malware is used for monetary gain. Botnets rely solely on unsolicited installations of malicious programs on ordinary computers to function. Other than targeted attacks, they aim at infecting ordinary users, who often may not know how to secure their systems properly. Chapter 11 provides detailed information about a Botnet's modus operandi. The basic enabler for such an installation, however, is still Windows-based malware.

Platform independence. Today, malware is still almost exclusively targeted at Microsoft Windows operating systems. With the biggest market share, these systems are more widely distributed and, therefore, more valuable. Although mobile malware and malware for other operating systems, such as Mac OS or Linux, is definitely on the rise, widespread attacks on these platforms have not yet had an impact on the general public—even though infections are often transmitted via Browser exploits or drive-by-downloads, methods that are platform independent.

(In)secure design. The most important problem, however, is the impossibility to design completely secure systems. As a result, there will always be an arms race between systems developers and miscreants that try to infect them with malware. There is, however, a noticeable shift in responsibilities. While open, general-purpose operating systems like Windows, Linux or Mac OS are ultimately the responsibility of the users themselves,

the introduction of App stores alleviates some of this responsibility. While they may restrict parts of the operating system's functionality, their positive effect is undeniable. With app stores, large companies can investigate applications before releasing them publicly [260, 417]. The result is a more secure environment which is less prone to be infected by malware than their open counterparts.

These examples represent just a quick glance at problems connected to malware. It is safe to assume that malware will play an important role in future operating systems. Still, research and companies are pointing in the right direction. They are constantly devising new methods to avert, or at least contain large malware outbreaks.

13 Social Engineering and Phishing

No one should underestimate the impact that the human factor has on security. Any chain is only as strong as its weakest link, and that is also the case with computer security. Consequently, adversaries often employ various techniques of social engineering to bypass or break security mechanisms by manipulating users. An accurate description of social engineering has been given by Kevin Mitnick [1], arguably among the most famous figures in this context:

> "Social engineering is using manipulation, influence and deception to get a person, a trusted insider within an organization, to comply with a request, and the request is usually to release information or to perform some sort of action item that benefits that attacker."

13.1 What Is the Problem?

Social Engineering is not limited to people within an organization. It is also employed against end users in various attack scenarios such as personalized spam and phishing campaigns. The term "phishing" was coined in 1996, when attackers used to refer to a compromised account as phish. At that time, phishers used to trade compromised accounts as a form of electronic currency. Today, phishing has evolved into a more sophisticated threat, with many targets. A typical phishing attack usually entails the adversary posing as someone the user trusts (e.g., a friend, a boss, an administrator, a web service) and requiring them to divulge some information (e.g., a password) or complete an action (e.g. click on a link).

Lots of recent incidents highlight the effectiveness of social engineering attacks. Accordingly, the findings of this year's RSA advanced persistent threat (APT) summit [339] designate social engineering as the number-one threat vector. Interestingly, social engineering is also the main infection mechanism used by Android malware writers to infect new devices: indeed, as detailed in Chapter 8, the Android security model isolates application processes so that applications cannot interfere with each other. As a result, classic infection mechanisms (e.g., drive-by download, memory corruption, code injection)

are no longer effective. Therefore, the malware authors must resort to social-engineering-based techniques for persuading victims to install legitimate-looking applications that hide malicious functionalities.

13.2 Who Is Going to Be Affected?

Most Internet users have come across cases of social engineering, such as in cases of spam emails originating from friends' email addresses [365] that have been compromised. In other cases, the emails originate from different addresses and just masquerade as having been sent by a friend. In both cases, the attackers goal is to exploit the implicit trust users show toward communication from their online contacts.

Nowadays, with the seemingly universal adoption of online social networks, and the abundance of personal information released, users are unwillingly and unknowingly aiding attackers in launching social engineering attacks. Thus, users of such social services are bound to become the main target of personalized spam campaigns, that incorporate user information in order to appear more convincing.

Furthermore, the explosive increase of mobile devices with Internet connectivity (i.e., smartphones) is slowly shifting the focus of malware authors to these devices. Smartphones combine telephone devices capable of "dialing-in" (i.e., have a built-in billing system), with a sophisticated environment capable of executing arbitrary code and, at the same time, offer a full-featured browser access to the Internet. Therefore, smartphones present a large attack surface as their users visit arbitrary sites on the web.

Attacks against high-value assets have been seen and are expected to become even more prominent, as activists resort to digital media for furthering political schemes [269], protesting against lawmaking and opposing oppressive regimes [84].

Chapter 6 of the SysSec Deliverable D7.1: *Review of the State-of-the-Art in Cyberattacks* [373] discusses the state of the art in social-network- and social-engineering-based attacks.

13.3 What Is Expected to Happen?

While the typical phishing activities via email and online social networks will continue to affect Internet users, social engineering is also expected to increase in various areas (most notably, by targeting mobile devices), and in its sophistication and scale (e.g., thanks to automation techniques).

Regarding mobile devices, as detailed in Chapter 8, Android's security design enforces that each installed application must run with a distinct user account. As a consequence, each application process has its own, isolated (virtual) memory space. Albeit simple, this security mechanism prevents a

(vulnerable) application (e.g., a browser) running within a certain process space, from being exploited to execute malicious code. In this security model, the role of social engineering and phishing becomes paramount, because the attacker model changes completely: the attacker is able to compromise a device only by managing to have a malicious application executed with the correct privileges. In the last 2 years, we have indeed witnessed numerous cases of mobile malware campaigns (see also Chapter 8), in which users were fooled using classic social-engineering techniques (e.g., email attachments, malicious applications disguised as legitimate software through official and unofficial marketplaces). In this evolving computing landscape, where powerful, expensive hosts shared by many users are giving way to powerful, inexpensive devices owned by a single user, we expect that social engineering techniques will be utilized more than in the past.

Given the prevalence of portable computers and mobile phones, the offenders have also been exploiting the voice channel in phishing campaigns. Although classic, this luring scheme has recently been gaining more attention [135, 264, 266]. Typically, the phishers contact their victims, via live telephone calls, automated responders, or SMS [148], and attempt to trick them into revealing sensitive information, performing some action (e.g., to unlock or repair their computers), or placing a payment for some expired (bogus) insurance. Albeit simple, this attack scheme is effective where other means fail: indeed, a live conversation enhances the effectiveness of social engineering techniques significantly. This does not happen normally with e-mailing because the e-mails have to be read, which decreases the chances of the attackers successfully luring the victims. A recent study also noticed a variation of this scheme, which is arguably more effective, where victims are lured into *contacting* the phishers through telephone numbers spread via social network's messaging systems (e.g., Twitter [164]).

Another interesting, yet dangerous, expected scenario is the significant spread of attacks that affect mobile payment and banking systems. In these systems, the mobile device, and thus its owner, becomes a much more sensitive target than in the classic PC-based banking operations. First, the usability of mobile apps guarantees a larger user base of payment systems. Second, trojans were among the first type of malicious software to be ported to mobile platforms, including ZitMo and SpitMo [120] (i.e., variants of ZeuS and SpyEye for Android). Notably, this scenario also creates complex attack venues such as the case were mobile devices are used as authorization tokens for financial transactions initiated from, for instance, a traditional web application.

In conclusion, what we expect is a less steeply increasing trend in traditional, email-based phishing than in the past. Traditional phishing will arguably give ground to heterogeneous (e.g., voice, social networks, SMS) phishing campaigns where social engineering will play a significant role, also thanks

to automated social engineering mechanisms [218] that employ smart bots to generate legitimate-looking text-based interactions (e.g., Facebook chat sessions). We can argue that this increased sophistication will culminate in personalized yet large-scale spear phishing attacks, where the attackers will be able to gather intelligence information about potentially any (social network) user, so as to generate tailored, realistic phishing messages [222, 294].

Thanks to the "amplification" effect ensured by the massive user base of today's social networks, future phishing attacks will probably have a large, real-world impact, while retaining the effectiveness and the power of targeted attacks. An interesting example of such an impact was the tremendous drop in stock values due to a news agency's compromised Twitter account, being used to publish a false story about the White House being bombed [131]. Although stocks recovered fast, this case showed the potential implications. Reports estimated that the three-minute plunge briefly wiped out $136.5 billion of the index's value [370].

13.4 What Is the Worst That Can Happen?

Depending on the target, the threat of phishing attacks extends from economic loss, to matters of national security. Several cases of targeted phishing incidents have been reported in the news, with targets including staff employed in critical infrastructure, political figures, military personnel, and common users. As Internet connectivity becomes more prevalent, and more people and services are connected, such attacks are bound to expand and evolve, ranging from economic fraud to organized acts of terror. From cyber-espionage to hostile attacks, all facilitated by social engineering and phishing techniques, the imminent threat at a global scale stresses the need for security researchers to tackle this unresolved problem. To demonstrate the potential impact of such attacks, we refer to recent incidents involving various types of attacks, all incorporating some aspect of social engineering or phishing.

Cases of mobile malware deploying phishing campaigns have already been reported and the potential impact on end users, in terms of stolen money, is very alarming. Eurograbber has infected over 30,000 users across Europe and targets e-banking services. Initially, the victim's computer is infected. During the next access to a banking website, the user is tricked into downloading a virus onto his mobile device as well. After the infection is complete, the attackers are able to intercept the security tokens sent as part of two-factor authentication schemes employed by web banking sites, and authorize transactions that remove money from the victim's account. Articles reported over $47 million being stolen [226], however future attacks that target banks all over the world could result in losses on a much larger scale.

Early in 2012, the most senior military commander of NATO was the victim of an impersonation attack. Attackers created a fake Facebook account with his name, hoping to trick people close to him into divulging personal details or sensitive information [200]. While this threat was prevented, in the future, more elaborate attacks can result in malicious individuals gaining access to critical information. Furthermore, recent reports [37] revealed that during the 2008 USA presidential campaigns of Barack Obama and John McCain, hackers employing phishing techniques were able to gain access to emails and a range of campaign files, from policy position papers to travel plans. As reported by CNN [10], during a security summit US Defense Secretary Chuck Hagel attributed this cyber-attack to the Government of China.

Several real-world examples demonstrate the potential impact of targeted attacks against critical infrastructure. A prominent example is that of Stuxnet, a highly sophisticated piece of malware designed to only target Siemens supervisory control and data acquisition (SCADA) systems that control and monitor industrial processes. Specifically, it targeted Iran's nuclear facilities. Reports speculate that the infections occurred through contaminated USB sticks [30] deliberately left to be found by engineers that worked in these facilities. This social engineering technique known as *baiting*, can assist attackers in bypassing firewalls or, as in this case, gain access to internal networks that are not accessible otherwise. According to the New York Times [38], specialists have attributed the development of this attack to government agencies, namely the USA and Israel, which has also been suggested by the official press of the Iranian government [39]. These cases indicate the effectiveness of social engineering and phishing, as they are even employed by government agencies. Stuxnet targeted uranium enrichment facilities and proved to be successful in shutting some down. What will happen when attackers start targeting facilities that control nuclear warheads?

In another case with potential devastating effects, at least two power distribution companies were the target of social engineering attacks [214]. The companies were called by an individual posing as a representative of a large software company, warning them that their computer had been infected by viruses and requesting them to run certain, potentially vulnerable, services. Luckily, the transmission managers identified the social engineering attacks and did not comply. However, this does not mean that such an attack will not succeed in the future, and one can only imagine the damage that could be caused by malicious adversaries gaining access to such a critical infrastructure as that of a power distribution company.

13.5 State of the Art

Dhamija's [154] is among the most cited works regarding "phishing." Although dating back to 2006, this research was the first that provided empirical evidence about the reasons why phishing attacks work: by analyzing the (ineffectiveness of) standard security indicators, the paper corroborates with objective findings the anecdotal (true) belief that phishing and social engineering work because of the scarce security education of the typical users. Albeit simple, this concept is still the foundation of today's social-engineering-based attacks. Three years later, Bilge et al. in [104] showed that, once an attacker has managed to infiltrate a victim's online social circle, the victim will trust the attacker and blindly follow any link they post, regardless of whether the victim knows the attacker in real life. Throughout the years, phishing and social engineering have evolved to find new ways to exploit trust relationships between human subjects, or between a human subject and an institution or website. A recent example is the abuse of short URLs [265] (e.g., `bit.ly`, `tinyurl.com`), to which users have grown accustomed thanks to Twitter, to spread phishing and other malicious resources on social networks and email campaigns. Unfortunately, many years later, security warnings, which are supposed to help inexperienced users to distinguish between trustworthy and non-trustworthy websites or resources, are still of debatable effectiveness [79].

Effective, personalized phishing and social-engineering-based attacks has been considered a small-scale threat, because collecting sufficient information and launching tailored attacks require time and manual effort. However Balduzzi et al. [92] and Polakis et al. [317] both demonstrated how online social networks can be used as oracles, for mapping users' email addresses to their Facebook profiles. Thus, using the information contained in the profiles, one could construct very convincing personalized spam emails. Furthermore, the authors have shown [109] that *automated* social engineering in social networks is feasible. They introduce the concept of socialbots, automated programs that mimic real online social network users with the goal of infiltrating a victim's social circle. They operated their proof-of-concept "socialbot" on Facebook for eight weeks and showed that current online social networks can be infiltrated with a success rate of up to 80%. Additionally, they show that, depending on users' privacy settings, an infiltration can result in privacy breaches with more users involved. Other work in the past tackled the threat of automated social engineering on social networks. Notably, Irani et al. [218] measured the feasibility of "attracting" victims using honey profiles, to eventually lure them into clicking on some malicious link. This "passive" social engineering approach turned out to be effective and once again showed that humans are often the weakest security link.

Scammers operating in online social networks have been analyzed by Stringhini et al. [364], where the authors observed that the scammers' social network usage patterns are distinctive because of their malicious behavior. This allowed them to design a system to profile and detect likely-malicious accounts with high confidence. In their work, the authors collaborated with Twitter and detected and deleted 15,857 spamming accounts. Three years later, Egele et al. [164] improved previous approaches to adapt them to the new techniques used by the attackers. Indeed, while in the past most of the scamming activity in online social networks used to be carried out through the creation of bogus accounts created, modern scammers have realized that compromising legitimate, real accounts makes their phishing and social engineering activities even more reliable. Egele's approach copes with this aspect using a combination of statistical modeling and anomaly detection to identify accounts that exhibit a sudden change in behavior. They tested their approach on a large-scale dataset of more than 1.4 billion publicly-available Twitter messages and on a dataset of 106 million Facebook messages. Their approach was able to identify compromised accounts on both social networks.

Recently, Onarlioglu et al. [307] performed a large-scale measurement of how real users deal with Internet attacks, including phishing and other social-engineering-based threats. Their findings suggest that non-technical users can often avert relatively simple threats very effectively, although they do so by following their intuition, without actually perceiving the severity of the threat. Another interesting, yet unsurprising, finding is that trick banners that are common in file sharing websites and shortened URLs have high success rates of deceiving non-technical users, thus posing a severe security risk. Non-technical users, and in particular elderly users, have also been targeted through less-sophisticated yet effective means: the so-called "vishing" (i.e., voice phishing) is the practice of defrauding users through telephone calls. We cannot identify when vishing first appeared (probably back in the phreaking era), neither can we state that this threat has disappeared [3, 55]. Albeit not widespread, due to its small scalability, vishing, also known as "phone scam" or "419 scam," has received some attention from researchers. To make this a viable business, modern scammers have begun to take advantage of the customers' familiarity with "new technologies" such as Internet-based telephony, text-messages [20], and automated telephone services. The first detailed description of the vishing phenomenon was by Ollmann [305], who provided brief, clear definitions of the emerging "*-ishing" practices (e.g., smishing, vishing) and pointed out the characteristics of the vishing attack vectors. Maggi [264] was the first to analyze this phenomenon from user-provided reports of suspected vishing activity. The majority of vishing activity registered was targeted against US phone users. By analyzing the content of the transcribed phone conversations, the author found that keywords such as "credit" and "press" (a key) or "account" are

fairly popular. Another interesting finding, which confirms anecdotal beliefs, is that vishers often rely on interactive voice responders to automate their calls. Recently, a study by Isacenkova et al. [219], based on a publicly available dataset of 419 scams, showed that this type of phishing practices, which are sometimes initiated via email, are on the rise. Interestingly, as suggested by Maggi et al. [266] in the past, Isacenkova et al. also found that phone numbers are the cornerstone that allows the different campaigns to be grouped together; their experiments also show that it is possible to identify large groups of scam campaigns probably run by the same criminal organizations. Suspicious phone calls have also been put under the microscope of Fujitsu and Nagoya University, which developed a proprietary technique for creating a model of a scammer's typical voice tone [34]; together with the extraction of keywords characteristics of scams, Fujitsu's system can detect suspicious situations of "overtrust."

13.6 Research Gaps

The effectiveness of social engineering and phishing attacks lies in the fact that users are unsuspecting and tend to trust communication (seemingly) originating from online contacts and sent through inherently "compromisable" media (e.g., email, online social networks). Defending against such attacks requires inter-disciplinary research in two orthogonal dimensions: (i) effective methods for educating users about the attacks, providing them with the basic skills for identifying them, and (ii) developing defense mechanisms for automatically identifying phishing attempts.

A major challenge for both dimensions is the ever increasing spear phishing attacks. From a technical aspect, they are deployed on a much smaller scale, and are thus able to evade the existing infrastructure (e.g., spam-traps) that collects samples for updating spam filters. From a user perspective, the content is crafted to resemble a legitimate communication and includes information and details that are very convincing, and can trick even careful users. Overall, we expect that in the near future attackers will have incorporated and be heavily dependent on social engineering techniques for delivering their attacks; accordingly, researchers will also have to focus on implementing effective countermeasures.

13.7 Example Problems

Even though phishing and social engineering are a relatively cold topic from a research perspective, these issues still lack effective solution. In this section we provide three example research problems, which all revolve around the idea of correlating phishing activities: the goal is to gain insights into how cybercriminals use their resources to carry out phishing and related threats.

This will shed more light on their modus operandi and, hopefully, allow researchers and practitioners to track them.

Intelligence-gathering malware and spear phishing: the prerequisite of spear phishing is that the phisher has gathered enough intelligence to "personalize" the interaction with the victim, with the goal of increasing the chances of the victim falling prey to the scam. On the one hand, the miscreants are arguably collecting such "intelligence" information manually; on the other hand, the abundance of data-stealing malware, both for desktop and mobile platforms, raises the question of whether this malicious software could be the main source of such information. The answer to this question is arguably positive; however, the empirical evidence is lacking, especially if we want to answer more complex questions such as "to what extent can spear phishing be automated?" To answer this and related questions, an accurate data-collection and correlation activity must be carried out.

Phishing and mobile malware infections: to bypass restrictions set by mobile OS, attackers incorporate social engineering to trick users into installing applications. An in-depth study focusing on the correlation between mobile infections and the techniques used to trick users into installing malicious apps, can provide researchers with valuable information that will lead to the implementation of monitoring components that detect such malicious activities and prohibit users from completing them. Apart from identifying the phishing techniques, these components can draw inspiration from traditional anti-phishing defenses such as domain blacklisting, and blacklisting malicious applications that have not been removed from application markets.

Cross-channel phishing correlation: modern phishing is complex and often involves several channels (e.g., email, IM, SMS, phone, online social networks). A holistic approach is thus required to observe and mitigate phishing more effectively. Based on an idea proposed in [266], we propose to capture different aspects of phishing campaigns, with a particular focus on the emerging use of the voice channel. The general approach is to record inbound calls received on honey phone lines, place outbound calls to the same caller identifiers (when available) and also to telephone numbers obtained from different sources. These sources include, for instance, phishing or scam instant messages, suspicious emails (e.g., spam, phishing). Extracted telephone numbers, URLs and popular words can be correlated to recognize campaigns by means of cross-channel relationships between messages.

14 Grand Challenges

One of the most important functions of a Roadmap is the definition of a few *Grand Challenge* problems. These are meant to be used as beacons that may enable longer-term research that should be valid for the next decade. A meaningful long-lasting solution to these challenges would probably require (i) significant technical work, (ii) interdisciplinary collaboration, and (iii) possibly legal support as well.

The challenges which have been identified include:

14.1 No Device Should Be Compromisable

Develop the necessary hardware and software support to make it impossible for attackers to compromise a computer or communication device for that matter, including smartphones and tablets. This challenge may summon support from policy or legal divisions as well.

14.2 Give Users Control Over Their Data

Provide the necessary mechanisms so that users

1. will be able to *know which data they have created* (such as text, photos, videos, cookies, web requests, etc.),

2. will be able to *know what data they have given to third parties* (such as text, photos, cookies, web requests, IP addresses, etc.)

3. will have the capability to *refuse disclosure of some data* (such as cookies and IP addresses) and still expect a decent level of service,

4. will have the capability to *delete their own data which they have created* (both from the local storage as well as from the cloud), and

5. will, under an appropriate legal framework, have the ability to ask past recipients of their data to erase them as well.

14.3 Provide Private Moments in Public Places

Enable users to have private communication in the public areas of the cyberspace. Consider the following analogy: The fact that people are having dinner in a public restaurant does not mean that their conversation could be recorded by the manager of the restaurant, and later made available without their explicit consent. Similarly, the fact that people are communicating in the cyberspace does not imply that parts of their communication can be recorded and used later through means outside their control. We propose to develop mechanisms that will enable people to have a reasonable expectation of privacy in what can be considered a public venue in the cyberspace.

14.4 Develop Compromise-Tolerant Systems

Provide adequate security levels even if components of the system have been compromised. It is reasonable to expect that not all attacks will be detected and successfully mitigated. Human errors, software errors, hardware errors, and insufficient protection mechanisms will allow some attacks to go through successfully. This implies that some systems, or components of systems will be compromised, and this may go undetected for a long period of time. Given such an environment, we should develop systems that will be able to provide decent security guarantees even if some of their components are compromised. Should a bank's accounts be allowed to empty because a teller's computer has been compromised? Should a cloud provider's password file be out in the open because an employee's account has been compromised? Should a user's private life be out in the open because a friend's account in a social network has been compromised?

How shall we design systems that will be able to provide decent levels of privacy and security given that some of their components have been compromised?

Part II: Related Work

15 Cyber Security: A Crisis of Prioritization

Pushing the state of the art, "Cyber Security: A Crisis of Prioritization" is considered one of the seminal works in this area [100, 249]. Ordered by the US President and implemented by the President's Information Technology Advisory Committee, the report suggested that Information Technology Infrastructure is "Critical," treated software as a major vulnerability, suggested that current solutions (such as endless patching) are not adequate, urged for the development of fundamentally new security models and methods, and elevated Cyber Security to the level of *National Importance*.

15.1 Problems - Priorities Identified

The report outlined several Cyber Security Research Priorities, including:

- "Usable and Reliable **Authentication**. Although there exist a lot of useful work on cryptographic protocols we need more research in usable and large-scale authentication which at the same time would decouple authentication from identification in order to address privacy issues."
- "**Secure fundamental Internet protocols** including BGP (Border Gateway Protocol) and DNS (Domain Name Service)."
- "**Secure software engineering and Software assurance**. Research is needed to develop secure programming languages and code that remains secure even when executed in different environments."
- "Provide a **holistic approach to System Security**. That is, the security of an integrated system is much more than just securing its individual components. For example, we need ways to build secure systems both from trusted and untrusted components."
- "Facilitate continuous **Monitoring and Detection** of malicious activities and attacks, including Intrusion Detection, real-time data collection, anomaly detection and appropriate data presentation that will allow operators to better understand incidents in progress."

- "Develop **Mitigation and Recovery methodologies**, to respond to unforeseen events and recover from any resultant damage. This area includes rapid automated discovery of outages and attacks, new architectures to enable rapid recovery, simplify systems to reduce human errors, and provide fault tolerance and graceful degradation."
- "Improve **Cyber Forensics** to more effectively catch criminals and deter criminal activities. To enable Law Enforcement Agencies to identify criminal activities in Cyber Space, we need sophisticated Cyber Forensics tools and mechanisms, such as traceback of network traffic to identify origins of attacks, efficient search of massive data stores to identify stolen information, and identifying attackers based on their behavior."
- "**Model new technologies and provide TestBeds** to experiment with them. Such testbeds and methodologies should scale to millions of nodes, should scale to very large amounts of data and should be designed in such a way as to preserve the confidentiality of data."
- "Some scientific disciplines have developed universally acknowledged metrics and benchmarks which enable researchers measure the effectiveness of their approaches and provably compare their contribution to the state of the art. In this spirit, we need to develop **Security Metrics, Benchmarks and Best Practices** for the Cyber Security field as well."

15.2 Recommendations

The main recommendations of the Report include:

- "NSF budget in this area be increased by $90 million annually."
- "The PITAC recommends that the Federal government intensify its efforts to promote recruitment and retention of cyber security researchers and students at research universities, with a goal of at least doubling the size of the civilian cyber security fundamental research community by the end of the decade."
- "The PITAC recommends that the Federal government strengthen its cyber security technology transfer partnership with the private sector. Specifically, the Federal government should place greater emphasis on the development of metrics, models, datasets, and testbeds so that new products and best practices can be evaluated."
- "PITAC recommends that the Interagency Working Group on Critical Information Infrastructure Protection (CIIP) become the focal point for coordinating Federal cyber security R&D efforts."

16 Forward: Managing Threats in ICT Infrastructures

Quantifying the cyber security priorities in 2008 and 2009, the FORWARD project (`http://www.ict-forward.eu/`), supported by the European Commission, established working groups to (i) discuss best practices, progress and priorities, (ii) set the research agendas to be pursued in Europe, and (iii) identify possible new research areas and threats that need to be addressed.

The main result of the project, the FORWARD Whitebook [133], contained detailed and concrete scenarios of how adversaries could exploit the emerging threats identified by the FORWARD project working groups to carry out their malicious actions. These scenarios illustrated future dangers and provided arguments to policy makers that are needed to support research in critical areas.

16.1 Research Directions Identified

The main research areas identified by FORWARD were grouped into several categories:

- "**Networking**. This area includes (i) attacks against the infrastructure of the Internet, such as against routers and routing algorithms, (ii) denial of service attacks where strategic links or essential backbone nodes are taken out of service, and (iii) wire-tapping attacks where the confidentiality or integrity of traffic is compromised, both on wired and wireless links. In addition to attacks against the Internet infrastructure, attacks may also be directed against end devices including (i) denial of service attacks against servers on the Internet, for example, by exploiting known vulnerabilities in applications or systems, (ii) distributed denial of service attacks, where the Internet infrastructure and the large number of unprotected nodes on the Internet are used to drown selected sites in traffic, and (iii) improper design or improper use of the services that the Internet offers, for example, the design of mission-critical systems that

are accessible from the Internet and possibly in-turn also depend on its services."

- "**Hardware and Virtualization**. This is probably the lowest level in the systems hierarchy where attackers may choose to operate. Although these attacks are usually difficult to deploy, they can remain stealthy for quite some time and thus be very effective. Such attacks may include (i) malicious hardware, and (ii) attacks within the cloud."

- "**Weak Devices**. Capitalizing on their small size and power requirements, such devices have recently enjoyed widespread deployment in the form of lightweight sensors, and RFID. Their deployment in the wild, and their mostly wireless communication abilities make them vulnerable to a wide variety of attacks including (i) information snooping, (ii) inserting false or misleading information, (iii) jamming radio channels, (iv) making nodes run out of battery by never letting them sleep, (v) giving the impression of phantom nodes that do not exist, (vi) giving the impression of connectivity that does not exist, and (vii) making messages go through an attacking node that can selectively drop messages from the system. Mobile phones (and PDAs) also fall under this category of weak devices, and can also be a target for attacks including (i) mobile malware, (ii) eavesdropping, and (iii) DoS Attacks."

- "**Complexity**. Over the past years we have been building increasingly complex systems which, by definition, are more prone to errors and attacks. Since these systems are difficult, if not impossible, to accurately model, they are challenging to test and may lead to several threats including: (i) unforeseen cascading effects, (ii) large-scale deployment of any attack, (iii) vulnerable system parts due to incomplete system maintenance, (iv) dormant functionality hidden in a program, and (v) race conditions and bugs due to multi-threaded/parallel nature of applications."

- "**Data Manipulation**: more people, more data, more value. As more people use the Internet, and as more organizations collect and store data on-line, we are bound to see an increasing number of attacks against (or based on) these data. The attacks may target several dimensions including: (i) erosion of privacy due to ubiquitous sensors, (ii) false sensor data due to fabrication or falsification, (iii) data leaked from social networks, and (iv) data gathered from (or for) on-line games."

- "**Attack Infrastructure**. To launch large-scale attacks, several adversaries develop and deploy distributed offensive platforms (such as botnets), which serve as underground economy support structures serving (and

operating on) advanced malware designed to evade detection and resist capture."

- "**Human Factors**. Humans are usually the weakest link in the security of several systems. Either as insider threats, or as end users, they may be the key element in the success of a cyber attack. Humans interact with security in several aspects including (i) user interfaces, which clearly convey a security (or lack thereof) to the user, (ii) insiders, who may have the access mechanisms needed to compromise a system, (iii) social engineering using all forms of communication, such as email, VoIP phones, and Instant Messaging Systems, and (iv) targeted attacks to individuals or groups of people."
- "**Insufficient Security Requirements**. Some systems, such as legacy systems (sometimes deployed even before the deployment of the commercial Internet), may have security requirements which are not adequate for the current time and scale."

16.2 Recommendations

The Report provided the following Recommendations:

Recommendation 1: "The EC should stimulate efforts that carry out research and development of techniques and systems (tools) for protecting against emerging ICT threats. The priority areas are:

- Protection of systems that are difficult to build, manage, and understand due to their scale and complexity
- Protection against malicious code (malware)
- Protection against threats that compromise users' privacy, particularly those on online social networks"

Recommendation 2: "The EC should support ongoing efforts to monitor developments in the ICT threat landscape. The threat landscape often changes rapidly and unpredictably as new technologies are deployed or new attacks are discovered. One requires an established and prepared entity to quickly react to these changes and assess the threat potential of new developments."

Recommendation 3: "The EC should support awareness initiatives and programs to educate its citizens about online threats and possible preventive actions. The reason is that certain threats cannot be addressed by technical means alone. Instead, defenses rely on proper reactions from informed users."

Recommendation 4: "The EC must recognize that ICT infrastructure is interconnected and deployed on a global scale. Hence, particular emphasis must be put on international collaborations and initiatives."

Recommendation 5: "The EC should (continue to) encourage interdisciplinary work and initiatives to bring together researchers from academia and industry as well as policymakers to cooperate on finding solutions to the threats against ICT infrastructure."

17 Federal Plan for Cyber Security and Information Assurance for Research and Development

Research and Development in Information Assurance are the main issues of the *Federal Plan for Cyber Security and Information Assurance Research and Development* [217], developed in 2006 by the Interagency Working Group (IWG) on Cyber Security and Information Assurance (CSIA), an organization under the United States National Science and Technology Council (NSTC). This Cabinet-level Council is the principal means for the US President to coordinate science and technology across the diverse parts of the Federal research and development enterprise.

The Plan presents baseline information and provides a coordinated interagency framework for addressing critical gaps in cyber security and information assurance capabilities and technologies. The Plan focuses on interagency research and development (R&D) priorities and is intended to complement agency-specific prioritization and R&D planning efforts in cyber security and information assurance. The Plan also describes the key Federal role in supporting R&D to strengthen the overall security of the IT infrastructure through the development of fundamentally more secure next-generation technologies.

17.1 Identified Research and Development Objectives

After a review of current legislative and regulatory policy requirements, analyses of cyber security threats and infrastructure vulnerabilities, and agency mission requirements, the CSIA IWG derived the following strategic Federal objectives for cyber security and information assurance R&D:

- "Support research, development, testing, and evaluation of cyber security and information assurance technologies aimed at preventing, protecting against, detecting, responding to, and recovering from cyber attacks that may have large-scale consequences."

- "Address cyber security and information assurance R&D needs that are unique to critical infrastructures."
- "Develop and accelerate the deployment of new communication protocols that better assure the security of information transmitted over networks."
- "Support the establishment of experimental environments such as test-beds that allow government, academic, and industry researchers to conduct a broad range of cyber security and information assurance development and assessment activities."
- "Provide a foundation for the long-term goal of economically informed, risk-based cyber security and information assurance decision making."
- "Provide novel and next-generation secure IT concepts and architectures through long-term research."
- "Facilitate technology transition and diffusion of Federally funded R&D results into commercial products and services and private-sector use."

17.2 Recommendations

The Plan recommends that cyber security and information assurance be accorded high priority at all levels of the Government and be integral to the design, implementation, and use of all components of the IT infrastructure. A critical observation is that the work that began with the Plan of identifying and prioritizing Federal cyber security and information assurance R&D efforts must be an ongoing process. Continuation of ongoing interagency coordination is needed to focus Federal R&D activities on the most significant threats to critical infrastructures and Federal agency missions and to maximize the gains from these investments.

The specifics of the strategy proposed in this Plan are articulated in a set of findings and recommendations, summarized as follows:

"**Target Federal R&D investments to strategic cyber security and information assurance needs.** Federal cyber security and information assurance R&D managers should reassess the Nation's strategic and longer-term cyber security and information assurance needs to ensure that Federal R&D addresses those needs and complements areas in which the private sector is productively engaged."

"**Focus on threats with the greatest potential impact.** Federal agencies should focus cyber security and information assurance R&D investments on high- impact threats as well as on investigation of innovative approaches to increasing the overall security and information assurance of IT systems."

"**Make cyber security and information assurance R&D both an individual agency and an interagency budget priority.** Agencies should consider cyber security and information assurance R&D policy guidance as they address their mission-related R&D requirements. To achieve the greatest possible benefit from investments throughout the Federal government, cyber security and information assurance R&D should have high priority for individual agencies as well as for coordinated interagency efforts."

"**Support sustained interagency coordination and collaboration on cyber security and information assurance R&D.** Sustained coordination and collaboration among agencies will be required to accomplish the goals identified in this Plan. Agencies should participate in interagency R&D coordination and collaboration on an ongoing basis."

"**Build security in from the beginning.** The Federal cyber security and information assurance R&D portfolio should support fundamental R&D exploring inherently more secure next-generation technologies that will replace today's patching of the current insecure infrastructure."

"**Assess security implications of emerging information technologies.** The Federal government should assess the security implications and the potential impact of R&D results in new information technologies as they emerge in such fields as optical computing, quantum computing, and pervasively embedded computing."

"**Develop a roadmap for Federal cyber security and information assurance R&D.** Agencies should use this Plan's technical priorities and investment analyses to work with the private sector to develop a roadmap of cyber security and information assurance R&D priorities. This effort should emphasize coordinated agency activities that address technical and investment gaps and should accelerate development of strategic capabilities."

"**Develop and apply new metrics to assess cyber security and information assurance.** As part of roadmapping, Federal agencies should develop and implement a multi-agency plan to support the R&D for a new generation of methods and technologies for cost-effectively measuring IT component, network, and system security. These methods should evolve with time."

"**Institute more effective coordination with the private sector.** The Federal government should review private-sector cyber security and information assurance practices and countermeasures to help identify capability gaps in existing technologies, and should engage the private sector in efforts to better understand each other's views on cyber security and information

assurance R&D needs, priorities, and investments. Federal agencies supporting cyber security and information assurance R&D should improve communication and coordination with operators of both Federal and private-sector critical infrastructures with shared interests. Information exchange and outreach activities that accelerate technology transition should be integral parts of Federal cyber security and information assurance R&D activities."

"**Strengthen R&D partnerships, including those with international partners.** The Federal government should foster a broad partnership of government, the IT industry, researchers, and private-sector users to develop, test, and deploy a more secure next-generation Internet. The Federal government should initiate this partnership by holding a national workshop to solicit views and guidance on cyber security and information assurance R&D needs from stakeholders outside of the Federal research community. In addition, impediments to collaborative international R&D should be identified and addressed in order to facilitate joint activities that support the common interests of the United States and international partners."

18 EffectsPlus Trust and Security Research Roadmap

STEERED BY by the Waterford Institute Of Technology (`www.wit.ie`), along with the partners Hewlett-Packard Ltd, SAP AG, ATOS and Università degli Studi di Trento, the EFFECTSPLUS project is a Coordination and support action aimed at the following five objectives:

- "The coordination of Trust and Security, Privacy, and Compliance for Future Internet through the structures and activities of the Future Internet Assembly (FIA) or subsequent initiatives"
- "The Clustering of Trust and Security (T&S) projects"
- "The provision of a channel for feedback and dissemination of R&D results and issues, and bringing together challenges for the strategic research agenda"
- "To analyse results from current and earlier trust and security work (i.e., from calls prior to Call 5), and to identify key areas and key players from new projects (Call 5) for the preparation of clustering and roadmapping activity"
- "To build and support the community of interests in trust and security results; providing the logistics and support for workshops, documents and web-based dissemination"

As part of its activities, the project held a research roadmapping and project clustering event in Brussels on 29–30 March, 2011. The participants at the event were representatives of European FP7 projects in the broad area of Trust and Security, and the objective was to identify core challenges and issues for research to be addressed in the timeframe 2010–2020 (in connection with the "Horizon 2020" strategy), as well as a shared vision of trust and security in the Future Internet.

18.1 Roadmap Structure and Realization

The proposed roadmap[1] is structured in four sections/chapters (Changes, Vision, Challenges, Solutions and Research Needs).

The roadmap was put together through a brainstorming session in two separate focus groups (one on Systems and Networks, and another on Services and Cloud Computing). Brainstorming results were translated into a mindmap, and then into a full writeup.

In the following subsections, we will separately analyze each section of the roadmap.

18.2 Changes

The changes foreseen can be grouped into:

- Changes in the impact of the Internet usage on end users (citizens)
- Changes in the way business is conducted (in different sectors)
- Changes in the broader socio-economic landscape

18.2.1 Changes for End Users/Citizens

End users will find the Internet more and more integral to their lives. This will also increase their awareness of the danger of loss of control over their private information. These conflicting agendas will raise interesting questions about the role of privacy-enhancing technologies and privacy-related research in the future.

There will be increasing interaction between physical and digital life, also thanks to always-on, always-connected, interoperable and smart devices.

18.2.2 Changes in the Business World

The business world is moving towards a scenario where everything happens "as a service," with personalized, heterogeneous services being offered through the cloud (including critical and sensitive infrastructure such as e-Health). The composition and orchestration of services will create issues with misbehaving/malicious components. Virtualized and outsourced infrastructures will become dominant.

This, and the consumer perspectives outlined above, will move the scale of the Internet to billions of nodes and above, increasing traffic and complexity.

Critical Infrastructures will become increasingly Internet-connected, and subject to sophisticated attacks.

[1] Which is unfortunately only available as a 2-year-old draft at the time of writing

18.2.3 Changes in Society and the Wider Economy

We are witnessing a massive growth of cyberthreats and cybercrime (which will reach physical infrastructures thanks to their connections with digital control systems).

Society as a whole will witness increased globalization, accompanied by deperimeterization and the increased digital nomadism of users.

18.3 Vision

The vision for the future of security starts with an improvement of privacy and awareness for users, and their empowerment to take care of their data. Similarly, businesses need to become more risk-aware.

Developers will need tools to build secure applications (as automatically as possible) and securely compose and orchestrate services, satisfying well-defined security properties. This will avoid security issues being a barrier to technology improvements.

An increase in user accountability will need to be carefully counterbalanced by a protection of human rights.

18.4 Challenges

A number of challenges must be met before this vision can be realized. Users and businesses must become able to understand and control their security and privacy posture, also through appropriate security metrics. Building secure and resilient systems must become easier, through appropriate tools and assurance frameworks.

Improved tools to express security policies and certify digital identities, as well as improved handling of system security issues and guaranteeing availability, are also a pressing need.

Several challenges are related with the new developments in the field: cloud computing, the rise of mobile devices, and socioeconomic changes.

18.5 Approaches and Potential Solutions

Potential technical solutions for the challenges outlined above are, among others, improved languages and tools for specifying secure software, improved assurance methods, privacy-aware software development, acceptable universal digital identifiers, data tainting and tracking.

Other solutions are societal and legislative, and may be achieved through the creation of technology-aware legal and law enforcement frameworks, and through the education of citizens.

18.6 Identified Priorities and Problems

Summarizing, the document identifies the following macro-priorities for research:

1. Integration of the Internet and of digital devices into users' lives and business processes, leading to increased dependence on the availability, security and privacy of these devices. Cyber-physical security and the security of mobile, cloud-connected devices will be of paramount interest.

2. Growth of the scale of the problem: the Internet will grow to include multiple billions of devices, traffic and complexity will grow, the number and prevalence of attacks will grow, perimeters will shatter and applications will become complex orchestrations of services

3. Tools, metrics and frameworks will need to evolve to cope with the unprecedented scale and integration of digital devices and processes in our lives

4. Law and education should go hand in hand in hand to protect users' privacy, increase accountability, and preserve human rights.

19 Digital Government: Building a 21st Century Platform to Better Serve the American People

The "Digital Government: Building a 21st Century Platform to Better Serve the American People" report outlines a digital strategy for the federal government in the United States to embrace new technologies, in a coordinated fashion, to better serve its citizens. Data and services should, as they say, be available *"anywhere, anytime, on any device"* in a secure fashion to encourage innovation. The roadmap outlines three major goals to be reached by the following four guiding principles: an information-centric approach for the data, a shared platform for consistency and to reduce costs, a focus on the needs of the users of the data, and, finally, an emphasis on security and privacy.

19.1 Identified Priorities

As identified in the roadmap, the exponential development of technology has changed how businesses perform in the private sector. Such advances should also influence how government agencies operate and serve data and services to citizens. By using the effectiveness of modern IT solutions, it is believed that better digital services can be built, using fewer resources. However, adopting new technologies may also pose a number of challenges.

For example, the roadmap introduces a conceptual model for digital services, with an information layer, a platform layer, and a presentation layer. Previously, systems made by the government have focused on a specific use case with a tight coupling between the information and the presentation, i.e., the presentation of data from this database should be tailored to a web page on a computer. By decoupling the information from the presentation, the presentation can be done on any device. Furthermore, by letting the data be machine readable, extensions can be built both by government agencies as well as entrepreneurs to better serve the citizens. The security of the end device becomes less important if one concentrates on securing the information itself.

The roadmap outlines the beginning of a coordinated path where data and services from the government can be better used in society. The roadmap complements a number of other directives.

19.2 Goals of the Digital Strategy

The three main goals of the strategy are the following:

1. "Enable the American people and an increasingly mobile workforce to access high-quality digital government information and services anywhere, anytime, on any device."

2. "Ensure that as the government adjusts to this new digital world, we seize the opportunity to procure and manage devices, applications, and data in smart, secure and affordable ways."

3. "Unlock the power of government data to spur innovation across our Nation and improve the quality of services for the American people."

19.3 The Guiding Principles of the Roadmap

The roadmap likewise outlines the following four principles:

1. 'An "Information-Centric" approach—moves us from managing "documents" to managing discrete pieces of open data and content that can be tagged, shared, secured, mashed up and presented in the way that is most useful for the consumer of that information.'

2. 'A "Shared Platform" approach—helps us working together, both within and across agencies, to reduce costs, streamline development, apply consistent standards, and ensure consistency in how we create and deliver information.'

3. 'A "Customer-Centric" approach—influences how we create, manage, and present data through websites, mobile applications, raw data sets, and other modes of delivery, and allows customers to shape, share and consume information, whenever and however they want it.'

4. 'A platform of "Security and Privacy"—ensures this innovation happens in a way that guarantees the safe and secure delivery and use of digital services to protect information and privacy.'

20 H2020: The Challenge of Providing Cyber Security

Under the vision of 2020 the challenges in cyber security are often the result of technological gaps and a lack of the necessary knowledge, in different quarters of our technology-filled society. To address these issues it is necessary to invest in research, and promote the resulting findings. In the rest of this section, we discuss the two different perspectives of the challenges ahead, in the context of Horizon 2020.

20.1 Industrial Perspective

20.1.1 Summary

The Trust and Security Unit at DG Connect held a workshop in July 2012 focusing on Horizon 2020 [25]. The objective of this workshop was to brainstorm on the challenges, technological gaps and necessary research directions related to cyber security and the best suited instruments to implement the tasks. The outcome of the discussion is meant to be used as input to the wider discussion on the thematic orientations of cyber security research, development and innovation in H2020.

20.1.2 Problems Identified

The group of experts identified a number of challenges during a brainstorming session:

- **"Addressing the needs and perspective of the user"** Specifically in terms of the usability of security products, perception of cyber security, education about and understanding of security issues, awareness raising, and development of appropriate security policies.

- **"Building capabilities"** That is, employing deterrence, building intelligence, sharing data and forming public-private partnerships.

- **"Making cyber security a positive business case"** Focusing on economics and explaining the cost of security, and qualifying security products.
- **"The role of technology"** Security should be built-in by design and security products should be sent to market when they are ready, instead of when they are half-baked.
- **"Defining cyber security metrics"** Security metrics must be developed, and compliance and benchmarking should be applied where possible. There is a serious lack of statistics and information about cyber security events, and this must be addressed.

20.1.3 Recommendations

The group of experts proposed the following instruments to address the challenges of cyber security:

- Fund R&D activities
- Promote demonstrators
- Support infrastructure build-up
- Provide users with education and training
- Give incentives

20.2 Societal Perspective

20.2.1 Summary

The Trust and Security Unit at DG Connect held a workshop in October 2012 focusing on Horizon 2020 [2]. The goal of the meeting was to invite participants to contribute their opinions on the societal issues of cyber security. The experts provided input on both substantive and procedural issues.

20.2.2 Procedural Issues

The experts identified the following procedural issues in terms of the societal perspective on cyber security:

- **"Widening the modalities of the public-private cooperation"** A model was suggested where a private company funds a public research lab to do research with them or for them, EU-funds could be added to that effort, in order to leverage the private investment

- **"Ensuring ethical and societal issues are duly examined"** Both ethical/human rights aspects and technical relevance should be part of the evaluation process.
- **"Ensuring a wider participation of civil society organizations or independent authorities"** Stakeholders should be included in the project proposals to provide a societal perspective.
- **"Coherence between research funding and policies orientations"** Policy-oriented research should be looked into, as well as coherence between the funding areas and policy goals.

In terms of procedural issues, the experts highlighted the following:

- **"Security: best ensured by identity or anonymity?"** Identification may not be the best way to perform online transactions. A more contextual way using synthesis from a variety of building blocks may be more appropriate.
- **"Recovering originality?"** We must reconsider the way we distinguish the original from the falsified in the online world, as it is often very hard to do in today's digital environments.
- **"Worrying about geo-strategic interdependence?"** We should consider the source of components in information systems, and look into the impact this has on cyber security.
- **"Software liability?"** We must reconsider software liability, as liability exists in other industries. The potential slow down of innovation should also be taken into account.
- **"Adopt a systemic approach"** We should look beyond security providers and the hierarchical model in our approach to cyber security.
- **"Security and the other fundamental rights"** Technology can be used to solve privacy issues and users can be given more control over technologies that affect them.
- **"Need to understand better the Internet ecosystem"** Understanding the Internet ecosystems is critical in order to shape and design the approach to cyber security.
- **"Security and consumer policy"** Smaller players should be protected, specifically those that are less able to evaluate risks in transactions.

21 Trust in the Information Society: A Report of the Advisory Board RISEPTIS

Visualizing the future of the Information Society, the RISEPTIS report is divided into 4 chapters. Chapter 1 introduces the report and gives a contextual overview of the main themes and issues addressed therein. Chapter 2 describes the use of concepts such as trust, trustworthiness, identity and accountability and explains how these relate to the EU legal framework of personal data protection and privacy. Chapter 3 highlights two important problems: a) *the dangers of our digital shadow*, and b) *the weakest links in the data storage chain*. Nine fictional scenarios demonstrate the risks associated with these problems. Finally, Chapter 4 lists 6 recommendations.

21.1 Problems Identified

The report highlighted two important problems:

The dangers of our digital shadow. People can be lured to give away private information easily. Even for insignificant prizes, people do not hesitate to reveal sensitive information about themselves. Moreover, the massive adaptation of social networks has driven users to publish information about their habits and lifestyle on the web. They essentially create a digital profile, which can be effectively used by third parties. This digital profile can be correlated and combined with information found in other places on the web. The result might be a superset of information, which the user is not aware of. This superset indicates that in certain cases a third party can learn more about a user, just by collecting public information, than the particular user intends. There are two key points that should be stressed here:

- Users are incapable of protecting themselves. Controlling and tuning privacy settings is complex rather than trivial for the average user.

- Taking advantage of the digital profile has severe consequences. A third party can take advantage of public information in arbitrary ways, giving rise to numerous fraudulent possibilities for would-be *identity-thieves*.

The weakest links in the data storage chain. Digital data can be stored on high-profile servers, where sophisticated security mechanisms are applied. However, it is still hard to guarantee that those data are never going to leak, since data are frequently transferred in data storage devices, such as CDs or USB sticks. These devices offer easy physical access. An attacker can alter the integrity of the data in transfer, break their confidentiality, or recycle the data with malicious purpose. Data encryption, if effectively used, can reduce such risks. However, data breach degrades the trust associated with victim companies or governments, even when the attacker reaps no practical benefit.

21.2 Recommendations

The report provided the following recommendations:

"Recommendation 1: The EC should stimulate interdisciplinary research, technology development and deployment that addresses the trust and security needs in the Information Society. The priority areas are:

- Security in (heterogeneous) networked, service and computing environments, including a trustworthy Future Internet,
- Trust, Privacy and Identity management frameworks, including issues of meta-level standards and of security assurances compatible with IT interoperability,
- Engineering principles and architectures for trust, privacy, transparency and accountability, including metrics and enabling technologies (e.g., cryptography),
- Data and policy governance and related socio-economic aspects, including liability, compensation and multi-polarity in governance and its management."

"Recommendation 2: The EC should support concrete initiatives that bring together technology, policy, legal and social-economic actors for the development of a trustworthy Information Society. (The Partnership for Trust in Digital Life could be a first step.)"

"Recommendation 3: The EC, together with the Member States and industrial stakeholders, must give high priority to the development of a common EU framework for identity and authentication management that ensures compliance with the legal framework on personal data protection and privacy and

allows for the full spectrum of activities from public administration or banking with strong authentication when required, through to simple web activities carried out in anonymity."

"Recommendation 4: The EC should work towards the further development of the EU data protection and privacy legal frameworks as part of an overall consistent ecosystem of law and technology that includes all other relevant frameworks, instruments and policies. It should do so in conjunction with research and technology developments."

"Recommendation 5: The EC together with industrial and public stakeholders should develop large-scale actions towards building a trustworthy Information Society which make use of Europe's strengths in communication, research, legal structures and societal values - for example, a Cloud which complies with European law."

"Recommendation 6: The EC should recognize that, in order to be effective, it should address the global dimension and foster engagement in international discussions, as a matter of urgency, to promote the development of open standards and federated frameworks for cooperation in developing the global Information Society."

22 ENISA Threat Landscape and Industrial Threat Reports

WHAT ARE THE FUTURE THREATS in cyber security? This is the main question addressed in this report produced by ENISA, the European Network and Information Security Agency: a center of network and information security expertise for the EU, its member states, the private sector and Europe's citizens. The ENISA Threat Landscape [268], a deliverable of the ENISA work programme of 2012, has been released on September 28, 2012, with the objective of describing the cyber-security threat landscape by consolidating existing threat reports. The main output of the report is a list of threats, threat agents and attack vectors.

22.1 Current Threat Trends

The report presents a list of current *threats*, along with the *threat agents* involved and, more importantly, an overall indication of the *current trend*. Besides spam, which is ranked as the only decreasing threat, the remainder threats are either increasing or stable.

The most relevant treat agents identified include, in order of novelty, *hacktivists* (a new trend that involve socially an politically-motivated individuals who target high profile websites to protest), *terrorists* (who today rely on cyber weapons to target critical infrastructures), *nation states* (which, with defense purposes, also rely on cyber weapons), *cybercriminals* (who have increased their skills toward more financial gain), *corporations* (which rely on offensive technologies or cybercriminals to gain competitive advantage over competitors), and *employess* (both hostile and non-hostile, who are still the main insider threat agent).

The above agents are involved in the following *increasing* threats, ordered by importance and frequency:

1. **Drive-by Exploits.** Vulnerable browser third-party components (primarily Java, but also Adobe Reader and Flash) are still the main target of drive-by exploits, which are mostly distributed through compromised legitimate sites.

2. **Malware.** Trojans are the most reported class of malware (also on mobile devices). Trojan Autorun and Conficker worm are still two of the top threats worldwide. Today, money making (e.g., through banking credential stealing) is the main motivation behind malware campaigns. With Koobface, the miscreant have showed that social networks are an effective distribution channel.

3. **Code Injection Attacks.** SQL injection attacks are today more popular than cross-site scripting attacks than in the past. Hacktivists rely on SQL injection attacks against their target websites.

4. **Exploit Kits.** Malware-as-a-Service (MaaS) is a new and growing criminal business. Modern criminals have a professional attitude, with support and development services. The enabling "technologies" of MaaS are (1) exploit kits (packages that automate cybercrime) and a (2) plethora of channels to deliver malware (malicious advertising, social networks, legitimate websites, malicious SEO).

5. **Botnets** Within the MaaS phenomenon, botnets have become a commodity since they switched from single-purpose (e.g., spamming, DDoS) to multi-purpose botnets. For increasing the reliability of botnets, cybercriminals rely on decentralized by peer-to-peer technologies (ZeroAccess botnets) and expand their surface to include mobile devices infected with specific malware ported from desktop-based OSs.

6. **Compromising Confidential Information.** 2011 has been addressed to as the "year of security breaches." Many sensitive database have been leaked or targeted by attacks (e.g., healthcare, law enforcement). According to the report, 9 out of 10 breaches would have been prevented through proper data protection and information security best practices. Indeed, besides targeted attacks, negligent, non-hostile insiders and web application vulnerabilities were the main cause of such breaches.

7. **Targeted Attacks** have been increasing during the first half of 2012, with spear-phishing as the topmost common infection vectors against industrial-control systems. Other tools used in targeted attacks include platform-specific malware: Stuxnet, Duqu, and Flamer.

8. **Physical Theft/Loss/Damage.** With the increased mobility of working locations and with the bring-your-own-device practice, the probability of data loss (even due to simple physical device theft) have increased in the last year. Unfortunately, the report highlights that full-device encryption is not widely adopted, although this would be a good mitigation technique.

9. Identity Theft is perpetrated through the spread of advanced trojans that perform information stealing, rather than phishing, which was more popular in the past. Mobile platforms are the main repository of sensitive information: Indeed, cybercriminals have ported information-stealing malware to such platforms (e.g., Android ZeuS in the Mobile, or ZitMo).

10. Abuse of Information Leakage is increasing, due to new technologies such as geo location and social advertising platforms (e.g., Foursquare), which create new venues for tracking users and compromising their privacy. Also aggressive advertising is abused to track users trough the information (e.g., "permanent" cookies) leaked by web browser.

11. Rogue Certificates are being leveraged to break the chain of trust. Indeed, in the last two years, the offenders have been stealing (see **6. Compromising Confidential Information**), producing and circulating rogue certificates. As a result, the criminals managed to conduct large-scale, man-in-the-middle attacks with stolen certificates.

There are threats which instead show a stable trend. Among them, *denial of service*, mostly used by hacktivists, is leaving traditional low-level protocols (e.g., UDP, ICMP, and SYN flood) in favor of application layer protocols (e.g., HTTP, service APIs), where more targeted denial of service attacks can be designed. An exception is the IPv6 layer, which was also targeted. *Phishing* has been stable (i.e., uptime of phishing sites dropped in first half of 2012), probably leaving the floor to more effective means such as information stealers (e.g., ZeuS or SpyEye), which can collect two-factor authentication credentials. However, cybercriminals are targeting VoIP systems via "vishing" (i.e., voice phishing) scams. *Rogueware/scareware* still be a problem, although the users are more aware of these scams. Indeed, the report notices little technical evolution in rogueware tools, although they are more widespread by leveraging the same distribution channels used by regular malware (e.g., SEO poisoning). Noticeably, the first fake AV product that targets Macs appeared in 2011. *Search engine poisoning* is still one of the major methods used to drive users to malware-distribution sites. SEO poisoning typically take advantage of events and trending topics to create campaigns that attract many victims.

22.2 Emerging Issues per Area

The report also presents a list of areas (i.e., assumed to grow), within which threat predictions are made from the current security issues.

Mobile Computing is affected by cross-platform malware families (e.g., ZeuS, SpyEye), which impact is exacerbated by the widespread use of mobile platforms for financial transactions. Unfortunately, the app stores are still too immature to fight back.

Social Technology is increasingly being leveraged for stealing information and, primarily, identities. This is predicted to grow into a "fake trust" effect that may one day be leveraged to build so-called social bots.

Critical Infrastructures are concerning because they integrate different systems from very diverse domains, each with peculiar security policies, practices and threats. On top of this, external factors such as political instability and financial crisis impact negatively by creating, respectively, motivation for attackers and vulnerabilities (e.g., cheap equipment). The increasing BYOD practice, along with its security issues, constitute an additional weak spot in the future critical infrastructures ecosystem.

Trust Infrastructure. Operators of trust infrastructures are likely to become targeted by offenders, whose goal is to compromise the chain of trust of the systems that rely on such infrastructures (e.g., social networks, web services). In this regard, the ENISA report highlights a need for more pervasive education and training to increase the users' awareness.

Cloud Computing and Big Data Cloud services such as remote backup and application services have become a consumer product. This, together with the massive use of social networks, yielded vast amounts of data, which are now an attractive target for attackers. Furthermore, the tight integration of cloud services in mobile devices will lead to a larger cloud attack surface, which could be exploited to compromise data privacy and to collect intelligence to prepare targeted attacks.

22.3 Recommendations

The ENISA report gives a series of recommendations that highlight the importance and usefulness of future threat landscapes in information security management. More precisely, rather than the typical list of recommendation for authorities and decision or policy makers, the report points out a list of "open issues" that need to be addressed by future threat landscapes. As this aspect is purely methodological, we present it in Section A.2

22.4 A Look at Industrial Threat Reports

22.4.1 Summary

At the beginning of each year, it is common for several security-related companies to publish reports or blog entries in which they try to summarize the trends they observed in the past and propose some threat predictions for the upcoming year. These are usually short term forecasts with a focus on technology and practical issues more than on long term research directions. However, these reports are compiled by the best experts in the area, and therefore they

represent the best information we can get to estimate what kind of problems we will have to face in the short-term future.

Therefore, we decided to complete this chapter on previous work on research roadmaps by reviewing a number of industrial reports, looking for recurrent patterns or common threats that we can reuse in our study. In particular, we covered the threat forecast published by Microsoft [377], Imperva [216], WebSense [393], McAfee [271], Symantec [367], Kaspersky [232], Bullguard [115], and by the Georgia Tech Information Security Center [193].

22.4.2 Common Threats and Recommendations

Not surprisingly, most of the predictions for 2013 have several points in common. In particular, these are the main areas on which experts from various companies seem to agree:

Mobile Malware. The emergence of mobile malware is one of the main concern we observed in the industrial reports. However, if the area itself is certainly the major threat on the landscape, the way in which it is going to materialize in the short term can vary. For instance, some experts see an increase in exploitation of vulnerabilities that target the OS and on the development of drive-by downloads; others think that malware will focus on the payment capabilities of phones to either steal information or to purchase applications developed by the attacker. Some companies even forecast the appearance of the first mass worm for Android devices. Finally, a common point in many reports is the likely increase of mobile adware, e.g., software that sends pop-up alerts to the notification bar, adds new icons, or change some of the phone settings.

Cloud-Based Malicious Activities. Attackers will leverage cloud infrastructure in general, and IAAS in particular, to perform a wide range of malicious activities. According to the experts, these can range from simple denial of service attacks (paid with stolen credit cards) to using the cloud to spread malware or to develop cloud-based botnets.

Malware fighting back. An interesting point raised by several experts is the fear that malware writers will start adopting more sophisticated techniques either to hinder the analysis and detection, or to make their command and control infrastructures more resilient. For instance, it is expected an increase in the adoption of techniques to detect virtual machines and in protection methods similar to those employed in Digital Rights Management (DRM) systems. Rootkits will also diversify, and will adopt new persistence mechanisms and bootkit techniques. McAfee also thinks that botnets will become harder to take down because malware

writers will include fall-back mechanisms in their code to reestablish control of the infected machines after a takedown.

Ransomware. The ransom business model has been tried in the past but several companies think that it will soon increase in popularity as a quick way for criminals to monetize their attacks. In fact, victims faced with the risk of losing their data are often willing to pay a ransom in the hope of regaining access. The raise of ransomware is going to affect all devices, from traditional computers to mobile phones—and it will be supported by the release of new and more sophisticated ransomware kits.

APT, Targeted Attacks, and Cyber-Espionage. Given the current trend, an easy prediction for the near future is that targeted attacks will increase both in number and in sophistication. Moreover, both nations and large organizations will become more involved in cyber-espionage, both as victims and as actors. In addition, thanks to the increased automation of these attacks, small companies will soon become the target of APT and espionage. Along the same line, experts also expect to observe more attacks for political reasons against individuals, minority groups, and non-government organizations. From a technical perspective, the security of supply chains is becoming a major problem that can lead to pre-installed malware and backdoors in many popular devices. Finally, as the distinction between cyber-espionage and cyber-war becomes more fuzzy, the use of sophisticated technologies so far observed only in "State-sponsored" attacks will become more widespread and will become part of the arsenal of common criminals.

Hacktivism. Hacktivism is probably going to change, but will not disappear any time soon. In the near future, experts think that the focus will be on "quantity over quality," with more private data stolen and published on the Web. At the same time, McAfee Threat Report discusses the fact that Anonymous' success will probably decline due to the low level of sophistication of their attacks. Unfortunately, new patriot groups may transform into cyber-armies and spread their extremist views.

As a possible consequence of the increase in hacktivist groups, in the near future we may observe a return to large and destructive attacks—not designed to gain an economic profit, but just to cause damage. Note that the possibility of such attacks was very popular in the mid-2000, when the fear of a *flash worm* that could bring down the Internet in a few minutes drove many researchers to focus on this threat. Nowadays, with critical infrastructures reachable from the network and new SCADA vulnerabilities discovered every month, the potential damage of a large scale destructive attack is higher than ever.

23 Cyber Security and Information Intelligence Research Workshop

X-RAYING, the underground economy reveals startling results: cybercrime and e-fraud are estimated to have exceeded $1 trillion globally in 2008 [244]. The 6th Annual Cyber Security and Information Intelligence Research Workshop (CSIIRW) [65] in 2010 focused on approaches that can help revolutionize the advancement of new ideas in cyber security. Since addressing narrow topics in security is no longer sufficient, it emphasized the need for comprehensive game-changing strategies, and solutions from novel multidisciplinary viewpoints.

This chapter is based on the outcomes of recent summits dedicated to defining new approaches to cyber security [348]. Apart from the CSIIRW workshop, we also consider the Federal Cybersecurity R&D Themes Kickoff Event organized by Network and Information Technology Research and Development (NITRD) [18], and the "Toward a Federal Cybersecurity Game Change Research Agenda" meeting organized by NITRD in May 2010.

23.1 Hard Problems in Security

The following list of five hard problems in cyber security was compiled by the experts.

- **Complexity.** The security of information contained in modern devices relies on a complex structure that includes hardware, operating systems, software, networks, data, and people. The components interact together in complex ways. Thus, a failure in one of them can result in a scenario where it is difficult even to determine the underlying cause. We need systems to help cope with this complexity.

- **Immense amount of data.** "Data" understood as *all electronic forms of information and knowledge* amounts to 1.8 zettabytes (1.8 trillion gigabytes) [166]. We need systems that can efficiently process this massive amount of data for attack recognition, understanding, and response.

- **Conversion of data into knowledge.** We convert raw data to information, i.e., "data in the context of other data," and then information to knowledge, i.e., "information in the context of other information." Only the latter provides us with understanding. Since not all processes can be automated with current technology, we need systems that can always create knowledge without relying on human intervention.
- **Nontechnical constraints.** These include
  - the need to protect private information,
  - usability and cost-effectiveness, e.g., law compliance, and
  - economic concerns.
- **The inadequacy of perimeter defenses in the networked world.** Instead of protecting individual components of systems, we should develop security as an integral part of the entire hardware-software combinations.

23.2 Research Directions Identified

To overcome the obstacles, and address the hard problems in cyber security, the experts devised a few game-changing themes. They aim to change the foundations of cyber security R&D. The three concepts that emerge from them are as follows:

- **Moving target (MT) defense for providing resilience through agility.** MT systems have the means to change in multiple dimensions so as to increase the degree of both uncertainty and complexity for attackers, as well as the resilience and fault tolerance within a system. As a result, attackers need to increase their costs and efforts in order to cause harm.

  Example MTs include dynamic networking, just-in-time compilation, and non-persistent virtual machines. In order to build MT systems, one might also randomize instruction sets and data, obfuscate operations by varying addresses, paths, and topologies, or decentralize cryptographic protection for credentials. Research challenges include ensuring scalability, performance, and energy consumption.

  Despite their inherent complexity, MT systems should be also easy to use, so they require management and configuration capabilities. Finally, the MT mechanisms must adapt quickly to diminish the window of opportunity for attackers, so they rely on innovative strategies to support real-time selection of MT protections.
- **Tailored trustworthy spaces (TTSs).** Since cyberspace blurs the boundaries between traditional spaces in the physical world, a TTS serves as a

flexible, distributed trust environment that can satisfy various requirements. TTSs support context-specific trust decisions—"security should be tailored to the needs of a particular transaction and not the opposite."

In order to build TTSs, "security must be end to end and top to bottom." Achieving results thus requires a shared vision, and approaches that allow control of the security in the whole system. More research is needed in the area of trust in heterogeneous environments, and the use of untrusted systems in trusted environments.

- **Cybereconomics (CE) for incentivizing good security.** Currently, there are no good ways to measure how secure a system is, so we also cannot estimate how much more secure it would become with an additional investment.

  Even though research in cybereconomics has grown, there are still numerous unanswered questions related to both cyberdefense and cyberattack. We need to analyze economic factors in attack and defense, so that we can build technologies that reduce the economic incentive for attackers. One way is to understand the structure of financial benefits of cyberattacks, and target the most vital components in this process.

23.3 Recommendations

Rather than a list of recommendations for authorities, the workshops devised and discussed a number of game-changing R&D themes that are essential for cyber security. Addressing the hard problems in security requires significant resources, and a long-term R&D vision focusing on the game-changing approaches. It is a multidisciplinary and challenging effort.

24 Cyber Security Strategy of the European Union

YOUTUBE has a press conference on EU Cyber security strategy [44] that makes it crystal clear: *the chances are that someone somewhere is attacking you—and you don't even know it*. In February 2013, the European Commission, together with the High Representative of the Union for Foreign Affairs and Security Policy, published a report on the cyber security strategy of the European Union [171]. The main goal of this report is to identify priorities for protecting and promoting citizens' rights online and for protecting cyberspace from incidents and malicious activities.

The strategy further details the roles and responsibilities of different stakeholders, both nationally and EU-wide, in working together to strengthen cyber security. While identifying national governments as the best place to deal with cyber security challenges, the report suggests actions for EU member states as well as EU institutions and the industry.

24.1 Cyber Security Principles

The cyber security strategy identified the following principles to guide the policy in the EU and internationally:

- The same core values of the EU that apply in the physical world, apply to the digital world as well.
- Cyber security should be based on protecting the fundamental rights and freedoms of individuals, and any information sharing for the purpose of cyber security should be compliant with EU data protection law.
- The Internet's integrity and security must be guaranteed in order to allow unhindered and safe access for everyone.
- Many commercial and non-governmental entities are involved in the day-to-day management of Internet resources, protocols and standards. Thus, a democratic and efficient multi-stakeholder governance of the Internet is of great importance.

- All relevant actors (public authorities, private sector and individual citizens) share responsibilities to protect themselves and strengthen cyber security.

24.2 Strategic Priorities

In order to achieve the goal of a safer EU online environment, the European Commission identified the following five strategic priorities:

1. **Achievement of Cyber Resilience.** Progress in this area has already been made based on voluntary commitments. The EU proposes to close gaps when it comes to national capabilities and coordination in the case of incidents across borders or in terms of private sector involvement and preparedness. The strategy includes a proposal for legislation to (i) establish requirements and strategies for Network and Information Security (NIS) at the national level and the need to set up a Computer Emergency Response Team (CERT) in each member state, (ii) foster the coordination of cyber security measures and information sharing amongst national NIS authorities, and (iii) improve the preparedness and engagement of the private sector by increasing incentives for private actors to embrace a cyber security culture.

   Furthermore, the strategy details the need to raise end users' awareness of cyber security by publishing reports, organizing expert workshops and developing public-private partnerships.

   Proposed Actions: *The EC should continue to identify vulnerabilities of critical infrastructure. The EC will also launch a pilot project for fighting botnets and malware via cooperation between member states, the private sector and international partners. ENISA should assist member states in building security expertise and improving the resilience of critical infrastructures. The industry should invest in cyber security and develop best practices and information-sharing mechanisms with public authorities. In order to raise awareness, the EC proposes, amongst other things, to increase national efforts towards NIS education and training. Finally, the industry should also promote cyber security awareness and reflect on the accountability for ensuring cyber security.*

2. **Drastic Reduction of Cybercrime.** Law enforcement should adopt a cross-border approach to respond to cybercrime through: (i) passing legislation such as the Council of Europe Convention of Cybercrime (Budapest Convention) and a Directive on attacks against information systems, especially through the use of botnets; (ii) enhancing operational capabilities to combat cybercrime and the use of state-of-the-art operational tools; and (iii) improving coordination at EU level.

Proposed Actions: *The EC will ensure the implementation of cybercrime-related directives and provide funding programs to support member states in strengthening their cybercrime combating capabilities. The EC will also support the cooperation between research/academia, law enforcement practitioners and the private sector. The European Cybercrime Centre (EC3) should act as the focal point in the fight against cybercrime. Furthermore, the EC proposes to increase the accountability of registrars of domain names and to ensure the accuracy of information on website ownership. Europol should support the member states' cybercrime investigations, and produce strategic and operational reports on trends, and target investigative action by cybercrime teams.*

3. **Development of Cyberdefense Policy and Capabilities.** In order to ensure EU member states' defense and national security interests, synergies between civilian and military cyber security mechanisms should be enhanced. The EC proposes to support these efforts by research and development and closer cooperation between governments, the private sector and academia as well as with international partners, such as NATO.

   Proposed Actions: *The High Representative will assess cyberdefense requirements and promote development of cyberdefense capabilities and technologies. The developed EU cyberdefense policy framework should include dynamic risk management, improved threat analysis and information sharing. The High Representative will promote the cooperation between civilian and military actors.*

4. **Development of Industrial and Technological Resources.** Hardware and software components used in critical services and infrastructures, as well as increasingly in mobile phones, need to be trustworthy and secure, and must protect personal data. Thus, the EC proposes (i) making security a priority for all actors in the value chain of cyber security products and (ii) fostering research and development investments and innovation. This should be achieved through the development of security standards, EU-wide voluntary certification schemes and the reduction of European dependence on foreign technologies.

   Proposed Actions: *The EC will develop incentives and recommendations for the adoption of secure ICT solutions and the take-up of good cyber security performance across the ICT value chain. The EC will further examine the possibilities for providers of ICT components to report detected security-critical vulnerabilities to national authorities. Public and private stakeholders should adopt security principles in their development process to ensure that new generations of software and hardware have stronger, embedded and user-friendly security features. EU member states should promote the involvement of industry and academia in developing and coordinating security solutions and should coor-*

dinate the research agendas of civilian and military organizations. Europol and ENISA should identify emerging trends and prerequisites to combat evolving cybercrime.

5. **Establishment of an International Cyberspace Policy.** The EU should participate in international cyber security efforts and promote achieving a high level of data protection. The EU should further participate in international collaborations to exchange best practices, share information and perform joint incident management exercises.

   Proposed Actions: *Together with all member states the EC will work towards an EU international cyberspace policy to increase engagement with international partners and organizations.*

25 The Dutch National Cyber Security Research Agenda

Zooming on developing a research agenda to fit the Dutch cyber security strategy, the National Cyber Security Research Agenda, or NCSRA, was embraced by the Dutch National Cyber Security Council and several ministries and has led to a substantial research program with a projected total research budget of approximately 30 million euro, of which the first 6.5M was allocated in a call in October 2012. Currently, a new research agenda is being developed. In this section, we will describe the current draft for the new NCSRA.

The process through which the research agenda is drafted is important. The initial draft is written by researchers from various research centers. Next, the writers' team invites feedback from as many stakeholders as possible in order to refine the agenda. Finally, it sends the agenda to the Security Council and various ministries to gather as much support as possible.

Another distinguishing feature is that the research agenda includes all relevant disciplines and not just ICT directly. It specifically caters to research in Law, Criminology, Economics, Sociology, Psychology, etc.

In the remainder of this section, we first provide examples of the different contexts that are relevant for the NCSRA (non-exhaustively) from both the technological and application domain point of view. Next, we describe the current list of research themes defined in the draft agenda.

25.1 Contexts

Concrete research questions typically arise in a specific context, which may involve a certain technology (e.g., cloud computing), or a particular application domain (e.g., finance), or a combination of the two. Still, similar research questions arise across different contexts, representing broader research themes. Below, we make an inventory of the most important contexts, regarding both technology and application domain. The next section then lists the underlying research themes that represent the central challenges for security across these contexts.

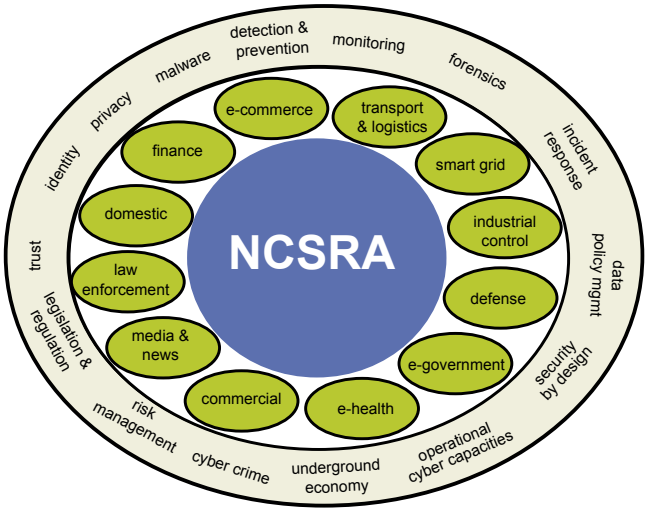

Figure 25.1: Application domains and research themes

25.1.1 Technologies

A central technology that is at the heart of most applications is of course the Internet, fixed or mobile. **Telecommunications and the Internet** are merging more and more to become an all-IP environment, where traditional telephony (voice), television (video) and data exchange are integrated into a multi-channel system. Services can be provided to large groups of users (broadcasting and information sharing), specific groups (narrowcasting and user communities) as well as single users. As many critical applications have come to rely on the Internet, the Internet itself has become an ever more critical infrastructure.

An important technology that builds on top of this is **cloud computing**. Cloud computing uses the communication infrastructure provided by the Internet to provide on-demand computation resources, in the form of raw computing power or more specialized services, by offering infrastructure, platforms or software 'as a service'. Cloud computing is increasingly used by individual citizens and companies to outsource their ICT needs. Cloud computing may offer economic benefits, by exploiting economies of scale and releasing users from maintenance tasks. However, cloud computing also introduces extra (communication) costs, and raises serious challenges for security.

Another important technological trend is **pervasive systems**: We are rapidly moving away from the desktop-model, and increasingly interact with ICT technology that is integrated into everyday objects and activities, that make up the *Internet of things*. Some of these devices are fully connected to the wider

Internet (e.g., smartphones), but many are not (e.g., wearable computing, or smart insulin pumps).

In some respects, cloud computing and pervasive systems are polar opposites: Cloud computing relies on massive centralization of data and processing power, whereas pervasive systems rely on a diverse distribution of processing power.

As we are surrounded by ever more devices with embedded electronics, the digital and physical worlds are rapidly converging to form one cyber-physical reality—in our homes and our workplaces, in semi-public places such as care homes and hospitals, in public spaces such as the (public) transport systems, and ultimately at a global level. Pervasive systems have important implications for privacy, security and trust and have a profound impact on our social lives. Also, some of the devices, for instance RFID tags, have only very limited capabilities when it comes to information storage, processing and communication, so that traditional methods for providing security are not feasible.

Besides the location of computation and hardware capabilities, the nature of the software involves a myriad of variations that have serious implications for security. Information exchange no longer has a predominantly client-server nature. Information is exchanged in a peer-to-peer fashion, more and more information is shared via social networks, and security sensitive operations (related to banking, healthcare, taxes, etc.) all occur via the Internet with a variety of technologies for such aspects as authentication and protection.

25.1.2 Application Domains

ICT technologies are used for many applications, ranging from generic use of ICT in the office or at home to more specific applications in industry, each with their own security requirements and threats. Below we highlight some—but by no means all—of these application domains in the NCSRA.

- **Domestic**. ICT and ICT networks play an increasingly important role in people's private lives, as the way to communicate and socialize (e.g., through social networks), as a source of information and entertainment (e.g., with gaming, and Internet taking over the role of television). This clearly has important security and privacy implications. Also, the huge ICT infrastructure collectively provided by Dutch citizens, with its excellent broadband connections, in itself has proved to be an interesting target for botnets.

- **Commercial**. Trust in ICT and Internet is vital for its ongoing and increasing use, and for companies to reap the economic benefits that this brings. Online commerce is increasingly important, and lack of trust in

ICT and Internet could undermine its growth: it has been estimated that increased trust in Internet by consumers could provide an additional 1.4 billion euro of online trade by 2014.

Just as private individuals are concerned with privacy, companies are concerned with their intellectual property and confidential information. Companies are faced with a rapid rise of ever more sophisticated cyber attacks aimed at corporate espionage.

- **Industrial Control Systems**. SCADA (short for "Supervisory Control and Data Acquisition") systems monitor and control large industrial establishments, such as chemical and nuclear plants, and large parts of the national critical infrastructure, such as the water, gas and electricity supply. Disruptions in SCADA systems can have disastrous consequences, but their increasing reliance on ICT—including the Internet—has made them vulnerable to remote attacks. Stuxnet is the most famous among numerous examples here. This is especially worrying as these systems are attractive targets for hacktivism, cyber terrorism, and cyber war.

  Improving the resilience of the ICT-dependent critical infrastructure requires research into these infrastructures as they exist today, to understand their interdependencies and judge their reliability in the face of attacks, and research into more secure components (hardware, software, or communication protocols) that may be needed to build a secure infrastructure.

- **Smart grid**. A new piece of technical critical infrastructure very much under development today is the smart grid, the next-generation electricity and utilities network that uses ICT technology to provide two-way digital communications between suppliers and appliances at consumers' homes, including smart meters and in the near future also batteries in electric cars. Smart grids are being promoted as a way of addressing energy independence, global warming and emergency resilience issues, but the increased reliance on ICT also introduces new threats, to both the security of the overall Grid and the privacy of individual users.

- **Finance**. Financial institutions or their customers are increasingly often victims of targeted cyberattacks, carried out by well-funded criminal organizations, which are becoming ever more sophisticated. These attacks are costing millions to consumers, retailer, and financial institutions (e.g., through skimming, stolen credit-card numbers, DoS attacks on payment infrastructure) and undermine the trust that is crucial for the financial system to operate.

  Present security solutions (firewalls, intrusion detection systems) cannot cope with this level of sophistication. There is a clear need for new

defensive approaches that can deal with targeted attacks and exploits of zero-day vulnerabilities. Identity fraud is also a major issue here. New payment schemes (e.g., using NFC mobile phones) may offer new technical and commercial possibilities, but also raise new security and privacy concerns.

- **Transport & Logistics**. Cars and transportation systems are increasingly making use of sophisticated software to carry out safety-critical processes, such as braking in cars. Drive-by-wire is already a reality, and in the near future intelligent transportation systems will make use of large-scale communication to optimize fuel consumption, reduce traffic jams, increase safety and implement smart tax charges, but this change also brings high security risks (e.g., it has been demonstrated that malware in a car could turn off the braking system). Moreover, the communication means that are needed to implement the smart mobility paradigm will turn the car into an open system that is by definition open to cyberattacks.

  Cars are only one example. Whitehat hackers have shown that new air traffic protocols are susceptible to a wide range of attacks. Several incidents in the past have shown that train services are vulnerable to software problems. There is no doubt that this is true for most modern forms of transport.

  In logistics, the main challenge in the domain is to ensure business continuity, while making the value chains as short and responsive as possible. A shorter chain has fewer participants and thus lower cost. A responsive chain delivers goods and payments faster, again lowering costs. However, in a shorter chain the risks of interruption of the logistics and transport services will increase and thus business continuity risks will increase.

- **e-Health**. Processes in the health sector are increasingly being supported by ICT. ICT is also the key enabler of new methods of providing care, as exemplified by ambient assisted living. However, patient data are often spread across many care providers such as the general practitioner, dentist, specialist, physiotherapist, hospital staff, pharmacists and, of course, the patient. Care providers must be able to access relevant information that is created and maintained by colleagues (e.g., medication records), to take action in case of emergencies while still guaranteeing the privacy of the patient's data. The security of patient data is essential to ensure that doctors obtain the correct information at the right time. The retention period for patient data is long (up to 70 years) and this poses a significant challenge for the technical infrastructure that supports the healthcare system.

- **e-Government**. The government plays different roles as far as cyber security is concerned. On the one hand, the government is a major user of ICT technology, with the increasing use of online information and services to citizens. Here the government is an important role model, and its conduct sets a standard. Also, ICT technology may provide new ways to promote democracy (e.g., through e-voting and local referenda).

  On the other hand, the government is responsible for the security and the protection of privacy for citizens, not only through legislation and law enforcement, but also through promoting awareness, by providing knowledge and expertise (e.g., via the National Cyber Security Centre, NCSC), and stimulating (inter)national collaboration. Just as governments already provide identities and means of identification for use in the physical world, they will increasingly do so in the online world, which may be crucial in combating identity theft as ever more services go online. Indeed, the introduction of an national electronic ID, the eNIK, is one of the stated objectives in the National Cyber Security Strategy. Finally, cyber espionage is a growing concern for government.

- **Military/defense**. In 2010, cyberwarfare became frontpage news, as well as a conspicuous reality with the Stuxnet attack on Iran. Cyber security is crucial to the military and the Department of Defense in terms of both defensive-reactive and proactive capabilities. Cyber defense is strongly related to resilience of the various critical infrastructures already mentioned above. Additionally, forensics and attribution are fertile grounds for research involving many disciplines. However, in most advanced countries, including the Netherlands, interest in a proactive strike force is growing, and more research and study is needed in this area.

- **Law enforcement**. Similarly, the use of ICT has become a crucial tool in many tasks related to tracking down, monitoring and apprehending criminals. Research is needed into improving these abilities without jeopardizing the safety and privacy of citizens. Some of these capabilities are extensions to existing capabilities like tapping, whereas others are entirely new. The research challenges include many different fields: technical, legal, sociological, etc. Again, attribution in particular is a difficult but hugely important research task.

- **Media and news outlets**. News outlets and mass media are important channels for disseminating information and thus make attractive targets for attackers. Both the news outlets and the threats are increasingly digital. In the past, we have witnessed compromises of government websites like that of Syria by Anonymous, but more traditional television

and radio broadcasts and printed media are possible targets too. Besides these traditional media, the domain also includes new media outlets such as blogs, social networks, tweets, etc.

25.2 Research Themes

The research agenda covers research themes that range from designing new systems in a secure way, to coping with the aftermath of attacks on existing systems. Each topic requires contributions from multiple disciplines: technical, legal, economical, etc. To structure the discussion on a potentially infinite list of research themes, the NCSRA distinguishes the following research themes:

1. **Identity, privacy and trust management**

   Managing (digital) identities, protecting user privacy and managing the trust in the online world are essential functionalities of the *future Internet* [21], which are required in each of the application domains listed above. The application domains concern important but distinct aspects of the digital life of the citizen. In each of these, different authorities, and different numbers of authorities—sometimes one (e.g., the government), sometimes many, sometimes none—will be responsible for providing and controlling identities, and different authentication mechanisms will be used. Therefore, different identity management solutions are needed to cater for the various needs. Research sub-areas include the computer science and cryptography techniques to ensure privacy and to handle identities securely, organizational rules and guidelines to delegate trust, and rules and legislation to deal with identity theft, privacy and anonymity rights, as well as private data retention and corresponding access rights.

2. **Malware and malicious infrastructures**

   Malware, short for malicious software, denotes all forms of hostile, intrusive, or annoying software or program code. The ability to run malware is essential for many types of attack, serving as a *conditio sine qua non* for generating social and economic power for the attackers.

   Thus, the threat of malware will remain critical for the foreseeable future. Currently, we experience the threat of malware most saliently in the form of botnets—millions of infected machines tied together in networks at the disposal of attackers. But malware evolves with the ICT infrastructure. We are already seeing malware on social networks, in cloud computing and on mobile devices.

   In terms of research, it poses an interdisciplinary challenge. We need advances in technology, for instance in reverse engineering, deobfuscation,

botnet tracking, analysis of criminal infrastructures, and classification and clustering of malware). Likewise, we need reliable methods for estimating the number of infected machines and the effectiveness of counter-measures. At the same time, we need arrangements to shape the socio-economic forces that fuel or mitigate the spread and impact of malware. From a historical perspective, we should study trends in malware—as doing so prepares us for new threats in time. Unless these issues are researched jointly, only partial solutions of limited value will be available.

While originating in criminal behavior, the magnitude and impact of the malware threat are also influenced by the decisions and behavior of legitimate market players, such as Internet Service Providers (ISPs), software vendors, e-commerce companies, hardware manufacturers, registrars and, last but not least, end users. Here, critical questions focus on economic incentives for the variety of market players. These can be shaped by self-regulation, state regulation and liability assignment.

3. **Attack detection, attack prevention, and monitoring**

   Malicious code and human attackers use ICT technology to launch attacks. The attacks include large-scale denial-of-service attacks, epidemic virus distribution, and stealthy attacks on high-value targets. Sometimes the attack is stealthy or even dormant, as in the case of backdoors, while in other cases, the attacks are very noisy. Monitoring systems and networks helps to detect and prevent the attacks as early as possible.

   Technological research challenges include binary hardening, network monitors, IDS and IPS systems, and attack analysis. For instance, to detect and prevent attacks, we need techniques and tools to spot and remove vulnerabilities from software, and monitoring systems to raise an alarm when a system behaves in an anomalous manner. Likewise, compliance monitoring is important for spotting vulnerabilities (in systems and organizations) as early as possible. From an organizational and management perspective, we need research into policies and protocols for monitoring, auditing, etc. From a legal perspective, we need clear rules for what is and what is not permitted in monitoring (and by whom), as well as ways to enforce these rules.

4. **Forensics and incident management**

   Forensics and incident management are related, but different activities that follow in the aftermath of an attack. Incident management consists of recovery (e.g., salvaging as much of the compromised system as possible), but also restoring systems and state, and becoming operational again at minimal cost. Part of the non-technical side of incident management

consists of setting up protocols and organizational structures for handling the incidents (who is in charge, who answers to whom, what role do various parties play, etc.)—all in the light of rules and regulations (such as the general obligation to report information leaks). Another part will be the assessment of such things like as: what was compromised, how badly was it compromised, can the data manipulation be reverted, what needs manual inspection, etc.? At the same time, it requires organizational procedures to deal with reviewing security protocols, disclosure to press and public, assembling teams to deal with the recovery, etc.

The goal of cyber forensics is to examine digital media in a sound manner to identify, preserve, recover, analyze and present facts and opinions about the information. The first decision after an incident is an economic one. How essential is the compromised system? For example, in a critical infrastructure setting such as a power station, it may be more important to get things up and running (without running the risk of a repeat) than to gather forensic evidence. In a crime scene, however, highly skilled digital forensics expertise is needed on-site as quickly as possible to collect evidence in a way that will make it admissible in a court of law. This process requires deeply technical as well as legal knowledge. Live forensics (forensics on a system that cannot be switched out, as in critical systems) and the attribution question (linking the criminal activity to the criminals behind it) are examples of issues that urgently require additional research. The same is true for the legal side: what is admitted as what sort of evidence under what circumstances? Forensic evidence has been used in a number of high profile cases and is becoming more accepted as reliable within US and European court systems. However, this is hampered by a lack of standards for digital forensic evidence, especially with multiple parties providing it. Again, research is needed into developing such standards and methods.

5. **Data, Policy and Access Management**

   In the application domains the variety of data plays a key role. However, the confidentiality, availability, authenticity and integrity requirements for different kinds of data can vary greatly, both in the technical as well as in the legal sense. For example, health records must be kept for 70 years, and therefore require strong security, whereas other data are almost ephemeral, such as the data by RFID tags. In this area, we need computer science research to develop data management techniques (possibly over very long time scales), but also organizational procedures, to ensure correct handling of sensitive data, and research to make sure that the technical policies match the user's mental models and understanding.

Security in systems that handle sensitive information requires the enforcement of information flow according to well-defined policies. We need research in novel access management techniques to help regulate who gets access to what information and under what circumstances. Given the current trend toward storing more and more data in the cloud, with the associated ambiguities regarding ownership and access, this problem is increasingly important. We need research that helps us decide how data should be managed, where it should be stored, how it will be maintained, who can do what with the data, and so on. And we need new technology to enforce these policies.

6. **Cybercrime and the underground economy**

   There is organized cyber crime, such as skimming, botnets, provision of child pornography and advance fee fraud, and unorganized (common) cyber crime, such as simple frauds, downloading child pornography, uttering threats, etc. In both cases we need to understand the (explaining) factors that lie behind the crimes, the modus operandi and the criminal careers of cyber criminals, and, in the case of organized crime, how their organizations work. We need to know more about patterns in cybercrime, who the victims are and how victimization can be explained. Since money (and consequently goods and information with a monetary value) is a key factor in many crimes, it is important to study the underground economy, its size, its characteristics and how it is intertwined with the legal economic system. In addition we need to investigate and assess the effectiveness of measures against cyber crime and the cooperation between (private and governmental; national and international) parties against cybercrime. What works and why? Do law enforcement agencies use their special powers for crime fighting in a digital world and, if so, with what result? The aim of research into the cybercrime area, is to design crime prevention strategies and measures to effectively disturb or block criminal activities.

   In addition, we often lack understanding about the socio-cultural context of the attack. Why is it doing what it is doing? The threat posed by Anonymous (the loose group of netizens and hackers that attacked companies that interfered with WikiLeaks) is very different from that of criminal organizations herding massive botnets, and that of state-sponsored cyber espionage and warfare. Studying the origin of attacks and the nature of the victims, as well as the language and socio-cultural references in malware will help linguists and sociologists to profile the attackers.

7. **Risk Management, Economics, and Regulation**

Risk management aims to assess the economic value of security, to provide a rational basis for allocating resources to improve security after identifying and assessing risks—and to determine if we are doing enough, or too much, and if we are spending resources on the right things. A scientific basis for establishing such facts is currently missing. One central problem here is that concrete data are often lacking, and more research could provide a more solid basis.

A much more fundamental problem is that risk assessment is typically done by an individual party, one of the many parties that collectively provide a complex value chain. For an individual party in such a complex value chain there may not be any economics incentives to fix a problem. Indeed, in cyber security there are many *externalities*: costs that are borne by other parties and hence not incorporated in price. For example, someone whose home PC is part of a botnet might not notice any adverse effects, and hence not be motivated to go through all the hassle of cleaning it. Perverse incentives may be a more important cause of security problems rather than the lack of suitable technical protection mechanisms. We need to carry out new studies into the incentives and externalities. We need research to get a better understanding of the economics of security—and the economic (dis)incentives that occur—and for more structural solutions of security problems.

Understanding economic drivers—and where these fail—is also crucial to determine where regulation is needed, and more generally what the government's role should be in cyber security. Different regulatory frameworks and legal constraints may apply in the various application domains, and at different levels: national, EU, and international.

8. **Secure Design, Tooling, and Engineering**

   Security engineering is a relatively new field and still lacks the methods and tools to design, build and cost-effectively test secure systems. ICT systems in use today are typically not designed and built with security in mind. As a result, security is often dealt with retrospectively, only after security problems arise. Security problems then have to be solved by an add-on in the design, when bad initial design decisions can no longer be reversed. When it comes to the software, fixing the problems requires costly bug fixes to patch implementations.

   Ideally, systems should be designed with security and privacy in mind from the start—ensuring **Security by Design** or **Privacy by Design**. They should then be implemented and tested using sound engineering principles and analysis techniques, in order to avoid security problems or detect them at an early stage. While considerable progress has been

made in some niche areas, such as security protocol analysis, sound engineering methods for security are still a long way off, especially when it comes to providing secure software.

Vastly improved tools are needed for testing and verification—to discover bugs prior to release. This includes testing for backdoors in regular and embedded systems. The general concept of fuzzing and program exploration are important research directions.

Besides software engineering, the field of economics plays an important role in this area. The cost of a secure design may be initially higher and requires a trade-off between risks and expenses. In addition, the cost over time for a secure design is likely to be quite different from that of less secure systems.

9. **Offensive Cyber Capabilities**

   In some domains, it is important to develop techniques to strike back at attackers. Besides the technical advances (often collectively referred to "hacking back"), these include ways to disrupt financial and other support infrastructures on which the adversary relies. Offensive cyber capacities are equally essential in testing the defenses of existing systems—for instance in penetration testing.

   Research challenges include the development of reliable techniques to penetrate other systems, evade defenses and escalate privileges. Non technical challenges include the development of legal guidelines to determine when offensive capacities may be used and by whom, and against which targets. Decision procedures and command structures for the use of offensive cyber force are similarly areas that require research.

Even if initially aimed at one specific application domain, research on the themes above can provide generic solutions that will apply to many application domains. For this to happen it is important that NCSRA helps to disseminate knowledge and project results across these application domains.

A Methodologies

In this appendix we outline the methodology used to create this Red Book. We list the people we mobilized, the way we organized them, and the interactions we had. We also list the procedure we followed and the meetings (physical and virtual) we had. For completeness we also include the methodologies used in the creation of the "Crisis of Prioritization" Report (section 15 in page 107) and the ENISA Threat Landscape Report (section 22 in page 131).

A.1 Cyber Security: A Crisis of Prioritization

In addition to engaging the members of the PITAC Committee, the co-chairs of the Committee organized or participated in a number of meetings as follows:

- **April 13, 2004 PITAC Meeting**: Members of US funding agencies presented the current state of funding: i.e. who funds what kind of cyber security research and to what level. PITAC members then discussed the issues addressed in the presentations. The public was then invited to make comments and ask questions.
- **June 17, 2004 PITAC Meeting**: The Subcommittee Chair gave an update on the Cyber Security Subcommittee's activities and solicited comment from PITAC members and the public.
- **July 29, 2004 Town Hall Meeting**: The purpose of the Town Hall meeting was to "solicit perspectives from the public on the current state of cyber security and the future measures needed to help ensure US leadership in this area." The participants were given a list of questions and asked to address them in their presentations.
- **November 19, 2004 PITAC Meeting**. The Subcommittee Chair presented the draft findings and recommendations and PITAC members provided feedback. Members of the public also provided comments.

A.2 ENISA Threat Landscape

The ENISA Threat Landscape [268], summarized in Chapter 22, originated from processing 120 individual reports. Threat reports from 2011 are actually

summarized documents, whereas the majority of reports were collected during 2012. The reports come from malware protection vendors, CERTS, security agencies, commercial companies in the area of security, industrial associations and committees, and networks of excellence.

The collected threat reports were prioritized into a list of current threats, which is included in the ENISA report as an annex (i.e., an actual list of phrases excerpted from the original sources). Trends were identified by identifying the emerging technologies and projecting the current threat to the respective technological areas.

A.2.1 Recommendations

ENISA recommends that future threat landscape reports and security-management actors follow some guidelines:

- Use a common terminology to refer to attacks, threats, actors, and so forth.
- Collect and develop better evidence concerning attack vectors and the impact achieved by adversaries. This is a challenging objective, but will ensure a more rigorous estimation of threat importance and trends.
- Collect information about threat agents and, more importantly, correlations among them.
- Include the user perspective, which is still absent from the majority of threat reports (i.e., users are not often the target of such reports).
- Develop use cases for threat landscapes, which will help in the analysis of the feasibility of future threats based on current and past landscapes.
- Collect security intelligence and share it across organizations as common knowledge bases.

B SysSec Threats Landscape Evolution

In this appendix a short overview of SysSec project cyberthreats and vulnerabilities landscape evolution will be briefly described, following the methodological framework and progress achievements, during the last three years since 2011.

B.1 Methodological Framework

The identification of threats and vulnerabilities for future Internet in the SysSec Network of Excellence has been organized on the basis of experts' brainstorming and q-based surveys, implementing the Delphi methodology, and jointly producing threat classification tables and a "plausible future" scenario context (see Figure B.1), initially proposed as a tool for strategic planning by RAND [158].

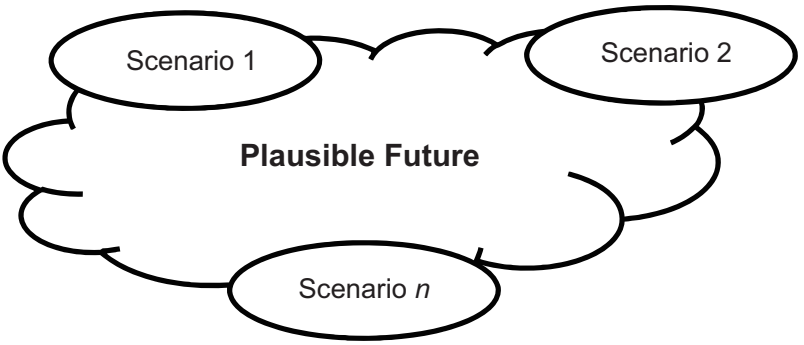

Figure B.1: Graphical interpretation of the "plausible future" idea.

The efforts of the SysSec consortium are also benefiting from ENISA [268], PITAC [316], World Economic Forum Reports [394] and NATO Comprehensive Approach [280] methodologies.

Generally, the idea is closely related to the application of the "scenario method." A method which uses scenarios as synthetic descriptions of events with driving factors, which are classified as important from the subject-matter experts.

As the implementation of the "scenario method," especially for future forecasting, is a rather complex task, the SysSec consortium has taken a stepwise application approach. SysSec is generating a research roadmap that encompasses the identified threats of future Internet from the project's yearly achievements.

B.2 SysSec 2011 Threats Landscape

In 2011, D4.1: First Report on Threats on the Future Internet and Research Roadmap [94], produced by the consortium, classified threats/enablers (represented as rows) in *personal*, *societal* and *professional* assets (columns), using a three-level colour scale: low (green), medium (yellow) and high (red). The generated results could be summarized in the table, shown in Figure B.2.

The generated context was encompassing two rather complex scenarios: "The Bank Job" and "The Peccadillo," accentuating on high priority threats, related to Internet usage of mobile devices, smart meters and the human factor, enabling obvious and hidden threats.

B.3 SysSec 2012 Threats Landscape

Evidently, the classification presented in Figure C.2 does not show the dynamics of threat trends, so a generalization was proposed in 2012. The results were published in D4.2: Second Report on Threats on the Future Internet and Research Roadmap [95].

A multi-aspect evaluation of: *System Security Aspects of Privacy*; *Collection, Detection and Prevention of Targeted Attacks*; *Security of New and Emerging Technologies*; *Security of Mobile Devices* and *Usable Security*, implementing three bilateral classification graphs (Likelihood/Impact; Technological Difficulty/Need for Research; Targets)/Time) within five years' time horizon has been performed. The results are generalized in Figures B.3 and B.4.

It should be noted here that the implemented evaluation scale was covering four degrees: blue (no votes), green (few, less than 3 votes), yellow (medium, less than 5), red (high, 5 or more votes).

Obviously, this classification gives a more dynamic and reasonable representation of the experts' beliefs for the next five years. The produced "plausible future" context is, as expected, broadened. Three rather complex scenarios are produced: "The Contact Dealer," "Portable Device in Stepping-stone Attack," "Password Reuse and Mobile Applications." These scenarios encompass the driving factors, related to mobile malware, mobile networks, social engineering, and password problems of human factors. The study is, in fact, covering the technological threats from Web 2.0 and the upcoming Web 3.0 technologies.

Threat-Enabler \ Assets	**Personal** Assets				**Societal** Assets		**Professional** Assets
	Privacy (Human Rights)	Digital Identity	Financial Assets	Health Safety	Critical Infrastructures	GRIDS Clouds	Data Sales etc.
Anonymous Internet Access	Medium	Medium	Low	Low	Medium	Low	Medium
Ubiquitous networks	High	High	High	High	Low	Low	Low
Human Factors	High	High	High	High	High	High	High
Insider attacks	High	High	High	High	High	High	High
Botnets	High	High	High	High	High	High	High
Program Bugs	High	High	High	High	High	High	High
Scale and Complexity	High	High	High	High	High	High	High
Mobile Devices	High	High	High	High	Medium	Low	High
24/7 connectivity	High	High	High	High	Low	Low	High
more private info available	High	High	Medium	High	Low	Low	Low
smart meters	High	High	Medium	High	High	Low	Low
Tracking	High	High	Medium	High	Low	Low	High
Smart Environments	High	High	Medium	High	Medium	Low	High
Unsecured Devices	High	High	High	High	Low	Low	High
Social networks	High	High	Medium	Medium	Low	Low	Low
Cyber-physical connectivity for Infrastructures, cars etc.	High	Low	Medium	High	High	Low	High
Organized Cyber Crime	High	High	High	High	High	Low	High
Mobile Malware	High	High	High	High	Medium	Low	High
SCADA Malware	Low	Low	Low	Low	High	Low	Medium
	Privacy (Human Rights)	Digital Identity	Financial Assets	Health Safety	Critical Infrastructures	GRIDS Clouds	Data Sales etc.

Figure B.2: SysSec 2011 threats/enablers experts' based classification in *personal*, *societal* and *professional* assets.

B.4 SysSec 2013 Threats Landscape

In 2013, the SysSec consortium is going even further, trying to forecast the future Internet threats in a more global context with the present Red Book: A Roadmap for System Security Research. This time a broader threats and vulnerabilities observation is being produced, trying to understand the nature of threats and the users' necessities for prevention in the "plausible future." The study is accentuating on cybersecurity landscape, structuring with: horizontal research areas (see Chapter 1)) and taking into account the human factor understanding about threats and vulnerabilities evolution, "assets we value," "domain of the game," and "what-if" scenarios (Chapter 2) for the possible dynamics overview. The classification is again not surprisingly related to Web

2.0/Web 3.0 technologies progress touching the future artificial interactiveness and noting: malware, targeted attacks and social engineering in the domains of: mobile devices, critical infrastructure protection and social networks.

Finally, four grand challenges (Chapter 14) are outlined: (i) development of non-compromisable devices, (ii) provision of users' data self-control mechanisms, (iii) enabling privacy on public places, and (iv) design of compromise-tolerant systems by means of adequate security.

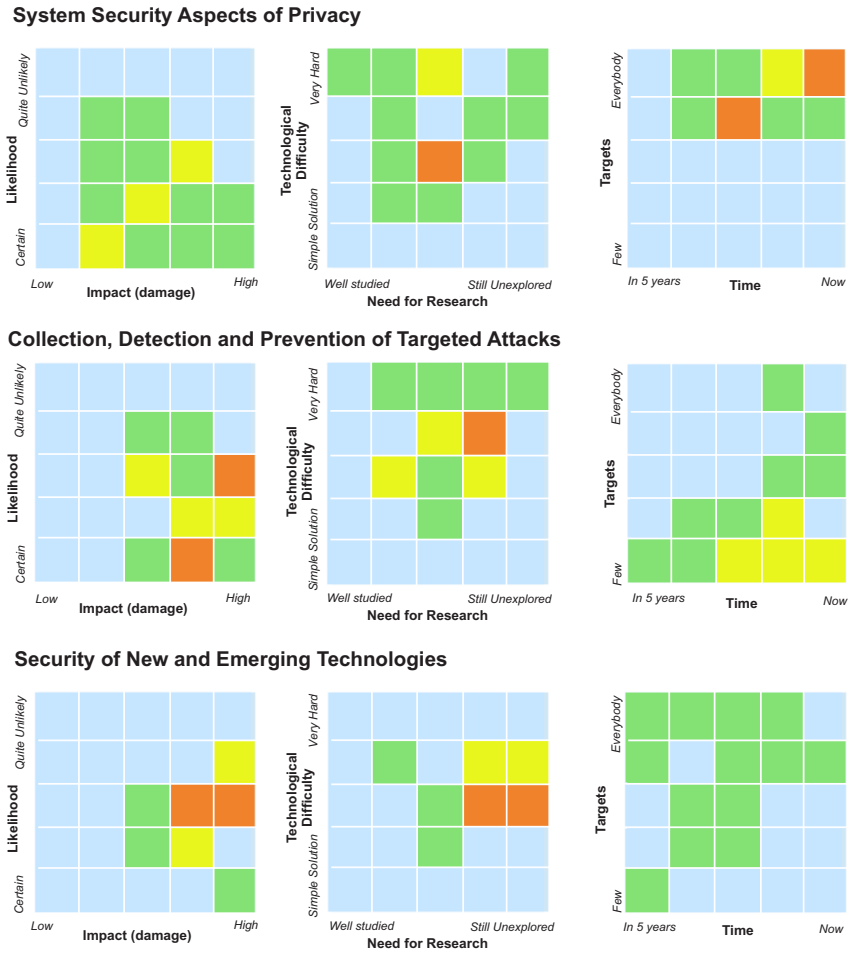

Figure B.3: Generalized results from D4.2: Second Report on Threats on the Future Internet and Research Roadmap about *System Security Aspects of Privacy*, *Collection, Detection and Prevention of Targeted Attacks* and *Security of New and Emerging Technologies* trends.

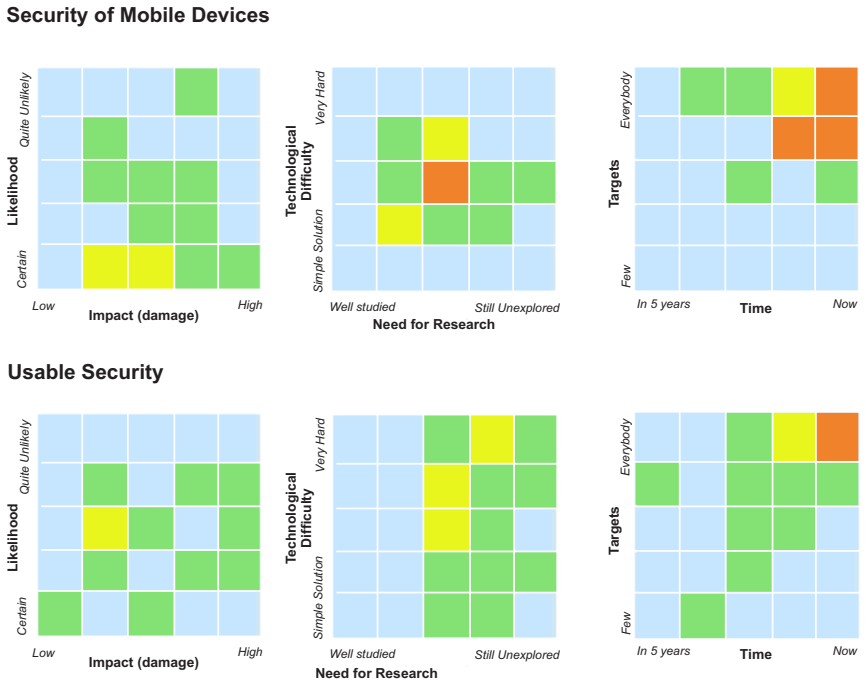

Figure B.4: Generalized results from D4.2: Second Report on Threats on the Future Internet and Research Roadmap about *Security of Mobile Devices* and *Usable Security* trends.

Bibliography

[1] 10 Questions for Kevin Mitnick. http://www.time.com/time/magazine/article/0,9171,2089344,00.html.

[2] A societal perspective on Cybersecurity. http://cordis.europa.eu/fp7/ict/security/docs/societal.pdf.

[3] A Spike in Phone Phishing Attacks? http://krebsonsecurity.com/2010/06/a-spike-in-phone-phishing-attacks/.

[4] Adblock Plus. http://adblockplus.org/.

[5] Adobe CoolType SING Table "uniqueName" Stack Buffer Overflow. http://www.metasploit.com/modules/exploit/windows/fileformat/adobe_cooltype_sing.

[6] Adobe Flash Player 11.3 Kern Table Parsing Integer Overflow. http://www.metasploit.com/modules/exploit/windows/browser/adobe_flash_otf_font.

[7] Andrubis: A tool for analyzing unknown android applications. http://anubis.iseclab.org/.

[8] China mobile users warned about large botnet threat. http://www.bbc.co.uk/news/technology-21026667.

[9] CloudCracker. https://www.cloudcracker.com. January 2013.

[10] Cnn security clearance: Hagel ties china to cyber attacks. http://security.blogs.cnn.com/2013/06/01/hagel-ties-china-to-cyber-attacks/.

[11] Copperdroid. http://copperdroid.isg.rhul.ac.uk.

[12] Cracking Android Passwords: The Need for Speed. http://linuxsleuthing.blogspot.com.es/2013/01/cracking-android-passwords-need-for_19.html. January 2013.

[13] Cyber crime now bigger than the drugs trade. http://www.theregister.co.uk/2011/09/07/cost_is_more_than_some_drug_trafficking/.

[14] DBLP Computer Science Bibliography. http://dblp.uni-trier.de/.

[15] Do Not Track - Universal Web Tracking Opt Out. http://donottrack.us/.

[16] Do Not Track: A Universal Third-Party Web Tracking Opt Out. http://tools.ietf.org/html/draft-mayer-do-not-track-00.

[17] Facebook Blocker. http://webgraph.com/resources/facebookblocker/.

[18] Federal Cybersecurity R&D Themes Kickoff Event 2010. http://www.nitrd.gov/CSThemes.aspx.

[19] Financial loss from identity theft increasing. http://californiawatch.org/dailyreport/financial-loss-identity-theft-increasing-report-says-16845.

[20] FTC Cracks Down on Senders of Spam Text Messages Promoting "Free" Gift Cards. http://www.ftc.gov/opa/2013/03/textmessages.shtm.

[21] Future Internet Assembly. http://www.future-internet.eu/.

[22] Google Accounts Authentication and Authorization. https://developers.google.com/accounts/docs/GettingStarted.

[23] Google Declares War on the Password. http://www.wired.com/wiredenterprise/2013/01/google-password/all/. January 2013.

[24] Google Glass. http://www.google.com/glass/start/.

[25] H2020: The challenge of providing cybersecurity. http://cordis.europa.eu/fp7/ict/security/ws-report-cybersec-v2_en.pdf.

[26] Hacker Posts 6.4 Million LinkedIn Passwords. http://www.technewsdaily.com/7839-linked-passwords-hack.html. December 2012.

[27] How Companies Learn Your Secrets. http://www.nytimes.com/2012/02/19/magazine/shopping-habits.html.

[28] How Target Figured Out A Teen Girl Was Pregnant Before Her Father Did. http://www.forbes.com/sites/kashmirhill/2012/02/16/how-target-figured-out-a-teen-girl-was-pregnant-before-her-father-did/.

[29] IEEE data breach: 100K passwords leak in plain text. http://www.neowin.net/news/ieee-data-breach-100k-passwords-leak-in-plain-text. December 2012.

[30] Ieee spectrum: The real story of stuxnet. http://spectrum.ieee.org/telecom/security/the-real-story-of-stuxnet.

[31] Industrial Instrumentation Process Lab. http://www.bcit.ca/appliedresearch/tc/facilities/industrial.shtml.

[32] MODBUS over serial line specification and implementation guide V1.0. http://www.modbus.org/docs/Modbus_over_serial_line_V1.pdf.

[33] MS12-063 Microsoft Internet Explorer execCommand Use-After-Free Vulnerability. http://www.metasploit.com/modules/exploit/windows/browser/ie_execcommand_uaf.

[34] Nagoya University and Fujitsu Develop World's First Technology to Detect Overtrust Situations Based on Voice Pitch and Level. http://www.fujitsu.com/global/news/pr/archives/month/2012/20120319-01.html.

[35] National SCADA Test Bed Program. http://www.inl.gov/scada/index.shtml.

[36] National SCADA Testbed. http://energy.sandia.gov/?page_id=7107.

[37] Nbc news: Chinese hacked obama, mccain campaigns. http://openchannel.nbcnews.com/_news/2013/06/06/18807056-chinese-hacked-obama-mccain-campaigns-took-internal-documents-officials-say.

[38] New york times: In cyberspace, new cold war. http://www.nytimes.com/2013/02/25/world/asia/us-confronts-cyber-cold-war-with-china.html.

[39] New York Times: Iran Suggests Attacks on Computer Systems Came From the U.S. and Israel. http://www.nytimes.com/2012/12/26/world/middleeast/iran-says-hackers-targeted-power-plant-and-culture-ministry.html.

[40] No Likie. https://chrome.google.com/webstore/detail/pockodjapmojcccdpgfhkjldcnbhenjm.

[41] OpenID Authentication 2.0. http://openid.net/specs/openid-authentication-2_0.html.

[42] PayPal Leads Industry Effort to Move Beyond Passwords. https://www.thepaypalblog.com/2013/02/paypal-leads-industry-effort-to-move-beyond-passwords/. February 2013.

[43] Plain Text Offenders. http://plaintextoffenders.com/. February 2013.

[44] Pres conference on EU Cyber Security Strategy. http://www.youtube.com/watch?v=qYOIlT9hqPA.

[45] SCADA & Security of Critical Infrastructures. http://resources.infosecinstitute.com/scada-security-of-critical-infrastructures/.

[46] SCADA HoneyNet Project. http://scadahoneynet.sourceforge.net/.

[47] SCADA/ICS Vulnerability Reference. http://scadahacker.com/vulndb/ics-vuln-ref-list.html.

[48] ShareMeNot. http://sharemenot.cs.washington.edu/.

[49] SHODAN - Computer Search Engine. http://www.shodanhq.com/.

[50] Sony Hacked Again, 1 Million Passwords Exposed. http://www.informationweek.com/security/attacks/sony-hacked-again-1-million-passwords-ex/229900111.

[51] State of SCADA Security 'Laughable,' Researchers Say. http://threatpost.com/state-scada-security-laughable-researchers-say-020312/76171.

[52] Stolen passwords re-used to attack Best Buy accounts. http://www.zdnet.com/stolen-passwords-re-used-to-attack-best-buy-accounts-7000000741/. February 2013.

[53] Study: 88 percent of Europeans victims of cybercrime. http://www.proofpoint.com/about-us/security-compliance-and-cloud-news/articles/study-88-percent-of-europeans-victims-of-cybercrime-800788986.

[54] Swiss soccer player banned from Olympics for racist tweet. http://usatoday30.usatoday.com/sports/olympics/london/soccer/story/2012-07-30/swiss-athlete-banned-michel-morganella-olympics/56591966/1.

[55] Tech Support Phone Scams Surge. http://krebsonsecurity.com/2012/08/tech-support-phone-scams-surge/.

[56] The Domino Effect of the Password Leak at Gawker. http://voices.yahoo.com/the-domino-effect-password-leak-gawker-10566853.html. February 2013.

[57] The OAuth 2.0 Authorization Framework. http://www.ietf.org/rfc/rfc6749.txt.

[58] Twitter Authentication & Authorization. https://dev.twitter.com/docs/auth.

[59] Twitter detects and shuts down password data hack in progress. http://arstechnica.com/security/2013/02/twitter-detects-and-shuts-down-password-data-hack-in-progress/. February 2013.

[60] Update: New 25 GPU Monster Devours Passwords In Seconds. http://securityledger.com/new-25-gpu-monster-devours-passwords-in-seconds/. December 2012.

[61] wiki.debian.org security breach. https://lwn.net/Articles/531727/. January 2013.

[62] Worldwide Observatory of Malicious Behaviors and Attack Threats. http://www.wombat-project.eu/.

[63] Council Directive 2008/114/EC, Dec. 2008. http://eur-lex.europa.eu/LexUriServ/LexUriServ.do?uri=OJ:L:2008:345:0075:0082:EN:PDF.

[64] Malware detected at the International Space Station. http://www.zdnet.com/blog/security/malware-detected-at-the-international-space-station/1806, 2008.

[65] 6th Cyber Security and Information Intelligence Research Workshop. http://csiir.ornl.gov/csiirw/10/, Apr. 2010.

[66] 2011 CWE/SANS Top 25 Most Dangerous Software Errors, 2011. http://cwe.mitre.org/top25/archive/2011/2011_cwe_sans_top25.pdf.

[67] PandaLabs Annual Report. http://press.pandasecurity.com/wp-content/uploads/2012/01/Annual-Report-PandaLabs-2011.pdf, 2011.

[68] N-gram Against the Machine: On the Feasibility of the N-gram Network Analysis for Binary Protocols. In *Research on Attacks, Intrusions and Defences Symposium*. Springer Verlag, Springer Verlag, 2012.

[69] Obama Order Sped Up Wave of Cyberattacks Against Iran, June 2012. http://www.nytimes.com/2012/06/01/world/middleeast/obama-ordered-wave-of-cyberattacks-against-iran.html?pagewanted=all&_r=1&.

[70] U.S. Team and Israel Developed Iran Worm, June 2012. http://online.wsj.com/article/SB10001424052702304821304577440703810436564.html.

[71] Cyberattack leaves natural gas pipelines vulnerable to sabotage, Feb. 2013. http://www.csmonitor.com/Environment/2013/0227/Exclusive-Cyberattack-leaves-natural-gas-pipelines-vulnerable-to-sabotage.

[72] Hello, Unit 61398, Feb. 2013. http://www.economist.com/blogs/analects/2013/02/chinese-cyber-attacks.

[73] McAfee Threats Report: First Quarter 2013. http://www.mcafee.com/us/resources/reports/rp-quarterly-threat-q1-2013.pdf, 2013.

[74] SANS SCADA and Process Control Security Survey, Feb. 2013. https://www.sans.org/reading_room/analysts_program/sans_survey_scada_2013.pdf.

[75] Symantec Internet Security Threat Report 2013. http://www.symantec.com/content/en/us/enterprise/other_resources/b-istr_main_report_v18_2012_21291018.en-us.pdf, 2013.

[76] M. Abadi, M. Budiu, U. Erlingsson, and J. Ligatti. Control-flow integrity. In *Proceedings of the 12th ACM conference on Computer and Communications Security (CCS)*, 2005.

[77] B. Adida. Beamauth: two-factor web authentication with a bookmark. In *Proceedings of the 14th ACM conference on Computer and communications security*, CCS '07, pages 48–57, New York, NY, USA, 2007. ACM.

[78] G. Aggrawal, E. Bursztein, C. Jackson, and D. Boneh. An analysis of private browsing modes in modern browsers. In *Proceedings of 19th Usenix Security Symposium*, 2010.

[79] D. Akhawe and A. P. Felt. Alice in warningland: A large-scale field study of browser security warning effectiveness. In *Proceedings of the 22th USENIX Security Symposium*, 2013.

[80] P. Akritidis, C. Cadar, C. Raiciu, M. Costa, and M. Castro. Preventing memory error exploits with WIT. In *Proceedings of the 2008 IEEE Symposium on Security and Privacy*, S&P'08, 2008.

[81] C. Albanesius. Google: Wi-Fi Sniffing Collected Whole E-Mails, URLs, Passwords. PCMAG.COM, October 2010. http://www.pcmag.com/author-bio/chloe-albanesius.

[82] S. Alexander. Defeating compiler-level buffer overflow protection. *USENIX ;login:*, 30(3):59–71, June 2005.

[83] S. Anand, M. Naik, H. Yang, and M. Harrold. Automated concolic testing of smartphone apps. In *Proc. of FSE*, 2012.

[84] Anonymous. Why we protest. http://whyweprotest.net/community/.

[85] M. Antonakakis, R. Perdisci, Y. Nadji, N. Vasiloglou, S. Abu-Nimeh, W. Lee, and D. Dagon. From Throw-Away Traffic to Bots: Detecting the Rise of DGA-Based Malware. In *Proceedings of the 21st USENIX Security Symposium*, 2012.

[86] C. Arthur. Conficker is a lesson for MPs - especially over ID cards. The Guardian, http://www.guardian.co.uk/technology/2009/apr/02/conficker-parliament-security-charles-arthur, 2009.

[87] E. Athanasopoulos, A. Makridakis, S. Antonatos, D. Antoniades, S. Ioannidis, K. G. Anagnostakis, and E. P. Markatos. Information Security. In T.-C. Wu, C.-L. Lei, V. Rijmen, and D.-T. Lee, editors, *ISC '08 Proceedings of the 11th international conference on Information Security*, volume 5222 of *Lecture Notes in Computer Science*, pages 146–160. Springer Berlin Heidelberg, 2008.

[88] A. Avizienis, J.-C. Laprie, B. Randell, and C. Landwehr. Basic concepts and taxonomy of dependable and secure computing. *IEEE Trans. Dependable Secur. Comput.*, 1(1):11–33, Jan. 2004.

[89] R. Baden, A. Bender, N. Spring, B. Bhattacharjee, and D. Starin. Persona: an online social network with user-defined privacy. In *Proceedings of the ACM SIGCOMM 2009 conference on Data communication - SIGCOMM '09*, page 135. ACM Press, 2009.

[90] G. Balakrishnan and T. Reps. Analyzing memory accesses in x86 binary executables. In *Procedings of the Conference on Compiler Construction*, CC'04, 2004.

[91] M. Balduzzi, C. Gimenez, D. Balzarotti, and E. Kirda. Automated discovery of parameter pollution vulnerabilities in web applications. In *Proceedings of the 18th Network and Distributed System Security Symposium*, 2011.

[92] M. Balduzzi, C. Platzer, T. Holz, E. Kirda, D. Balzarotti, and C. Kruegel. Abusing social networks for automated user profiling. In *International Symposium on Recent Advances in Intrusion Detection (RAID 2010)*, 9 2010.

[93] J. Baltazar, J. Costoya, and R. Flores. The Real Face of KOOBFACE : The Largest Web 2 . 0 Botnet Explained, 2009.

[94] D. Balzarotti(Ed.). D4.1: First Report on Threats on the Future Internet and Research Roadmap. Technical report, SySSeC Consortia, Sept. 2011.

[95] D. Balzarotti(Ed.). D4.2: Second Report on Threats on the Future Internet and Research Roadmap. Technical report, SySSeC Consortia, Sept. 2012.

[96] A. Baratloo, N. Singh, and T. Tsai. Transparent run-time defense against stack smashing attacks. In *Proceedings of the USENIX Annual Technical Conference*, June 2000.

[97] A. Barth, J. Caballero, and D. Song. Secure Content Sniffing for Web Browsers or How to Stop Papers from Reviewing Themselves. In *Proceedings of the 30th IEEE Symposium on Security & Privacy*, Oakland, CA, May 2009.

[98] A. Barth, C. Jackson, and J. C. Mitchell. Robust Defenses for Cross-Site Request Forgery. In *Proceedings of the 15th ACM Conference on Computer and Communications Security (CCS)*, 2008.

[99] U. Bayer, C. Kruegel, and E. Kirda. Ttanalyze: A tool for analyzing malware. In *Proc. of EICAR*, 2006.

[100] M. Benioff and E. Lazowska, editors. *Cyber Security: A Crisis of Prioritization*. National Coordination Office for Information Technology Research and Development, Feb. 2005.

[101] J. Bennett, Y. Lin, and T. Haq. The Number of the Beast, 2013. `http://blog.fireeye.com/research/2013/02/the-number-of-the-beast.html`.

[102] E. Bhatkar, D. C. Duvarney, and R. Sekar. Address obfuscation: an efficient approach to combat a broad range of memory error exploits. In *In Proceedings of the 12th USENIX Security Symposium*, 2003.

[103] R. Biddle, S. Chiasson, and P. Van Oorschot. Graphical passwords: Learning from the first twelve years. *ACM Comput. Surv.*, 44(4):19:1–19:41, Sep 2012.

[104] L. Bilge, T. Strufe, D. Balzarotti, and E. Kirda. All your contacts are belong to us. In *Proceedings of the 18th international conference on World wide web - WWW '09*, page 551. ACM Press, 2009.

[105] H. Bojinov, D. Boneh, R. Cannings, and I. Malchev. Address space randomization for mobile devices. In *Proceedings of the fourth ACM conference on Wireless network security*, WiSec '11, pages 127–138, New York, NY, USA, 2011. ACM.

[106] H. Bojinov, E. Bursztein, and D. Boneh. XCS: Cross Channel Scripting and Its Impact on Web Applications. In *CCS '09: Proceedings of the 16th ACM conference on Computer and communications security*, pages 420–431, New York, NY, USA, 2009. ACM.

[107] J. Bonneau. Statistical metrics for individual password strength. In *Proceedings of the 20th international conference on Security Protocols*, pages 76–86, 2012.

[108] J. Bonneau, C. Herley, P. C. v. Oorschot, and F. Stajano. The quest to replace passwords: A framework for comparative evaluation of web authentication schemes. In *Proceedings of the 2012 IEEE Symposium on Security and Privacy*, pages 553–567, 2012.

[109] Y. Boshmaf, I. Muslukhov, K. Beznosov, and M. Ripeanu. The socialbot network: when bots socialize for fame and money. In *Proceedings of the Annual Computer Security Applications Conference*. ACM, 2011.

[110] E. Bosman, A. Slowinska, and H. Bos. Minemu: The World's Fastest Taint Tracker. In *Proceedings of 14th International Symposium on Recent Advances in Intrusion Detection*, RAID'11, 2011.

[111] P. Boykin and V. Roychowdhury. Leveraging social networks to fight spam. *Computer*, 38(4):61–68, Apr. 2005.

[112] C. Braz, A. Seffah, and D. MŠRaihi. Designing a trade-off between usability and security: A metrics based-model. *Human-Computer Interaction–INTERACT 2007*, pages 114–126, 2007.

[113] T. Bukowski. ZeuS v3 P2P Network Monitoring, 2012. Technical Report by CERT.pl.

[114] Bulba and Kil3r. Bypassing StackGuard and StackShield. *Phrack*, 10(56), 2001.

[115] BullGuard. Security predictions for 2013. Internet. `http://blog.bullguard.com/2013/01/bullguards-security-predictions-for-2013.html`, 2013.

[116] P.-M. Bureau. Same Botnet, Same Guys, New Code: Win32/Kelihos. In *VirusBulletin*, 2011.

[117] J. J. C. Cowan, S. Beattie and P. Wagle. Pointguard: Protecting pointers from buffer overflow vulnerabilities. In *Proceedings of the* 12^{th} *USENIX Security Symposium*, August 2003.

[118] C. Cadar, D. Dunbar, and D. Engler. KLEE: Unassisted and automatic generation of high-coverage tests for complex systems programs. In *Proceedings of the 8th USENIX Symposium on Operating Systems Design and Implementation*, OSDI'08, 2008.

[119] J. Carr. *Inside Cyber Warfare*. Mapping the Cyber Underworld. O'Reilly Media, Dec. 2011.

[120] C. Castillo. Spitmo vs Zitmo: Banking Trojans Target Android, Sept. 2011. `http://blogs.mcafee.com/mcafee-labs/spitmo-vs-zitmo-banking-trojans-target-android`.

[121] L. Cavallaro, P. Saxena, and R. Sekar. On the limits of information flow techniques for malware analysis and containment. In *DIMVA*, pages 143–163, 2008.

[122] D. Chappell. Introducing Windows CardSpace. msdn, April 2006. `http://msdn.microsoft.com/en-us/library/aa480189.aspx`.

[123] S. Checkoway, D. McCoy, B. Kantor, D. Anderson, H. Shacham, S. Savage, K. Koscher, A. Czeskis, F. Roesner, and T. Kohno. Comprehensive experimental analyses of automotive attack surfaces. In *Proceedings of the 20th USENIX Security Symposium*, 2011.

[124] S. Chen, J. Xu, E. C. Sezer, P. Gauriar, and R. K. Iyer. Non-control-data attacks are realistic threats. In *Proceedings of the 14th USENIX Security Symposium*, 2005.

[125] W. Cheng, Q. Zhao, B. Yu, and S. Hiroshige. TaintTrace: Efficient Flow Tracing with Dynamic Binary Rewriting. In *Proc. of ISCC*, pages 749–754, 2006.

[126] W. Cheswick. Rethinking passwords. *Communications of the ACM*, 56(2):40–44, 2013.

[127] E. Chien, L. OMurchu, and N. Falliere. W32.Duqu: the precursor to the next stuxnet. In *USENIX conference on Large-Scale Exploits and Emergent Threats*. USENIX Association, Apr. 2012.

[128] V. Chipounov, V. Kuznetsov, and G. Candea. S2E: A platform for in vivo multi-path analysis of software systems. In *Proceedings of the 16th Intl. Conference on Architectural Support for Programming Languages and Operating Systems*, ASPLOS'11, 2011.

[129] L. S. Clair, L. Johansen, W. Enck, M. Pirretti, P. Traynor, P. McDaniel, and T. Jaeger. Password exhaustion: Predicting the end of password usefulness. In A. Bagchi and V. Atluri, editors, *ICISS*, volume 4332 of *Lecture Notes in Computer Science*, pages 37–55. Springer, 2006.

[130] G. Cluley. 600,000+ compromised account logins every day on facebook, official figures reveal. nakedsecurity news from SOPHOS, October 2011. `http://nakedsecurity.sophos.com/2011/10/28/compromised-facebook-account-logins/`.

[131] CNBC. False rumor of explosion at white house causes stocks to briefly plunge. `http://www.cnbc.com/id/100646197`.

[132] F. B. Cohen. Operating system protection through program evolution. *Computers and Security*, 12:565–584, Oct. 1993.

[133] T. F. Consortium. White book: Emerging ICT Threats, Jan. 2010. http://www.ict-forward.eu/media/publications/forward-whitebook.pdf.

[134] M. Costa, J. Crowcroft, M. Castro, A. Rowstron, L. Zhou, L. Zhang, and P. Barham. Vigilante: end-to-end containment of internet worms. In *Proceedings of the 20th ACM Symposium on Operating Systems Principles*, SOSP'05, 2005.

[135] A. Costin, J. Isacenkova, M. Balduzzi, A. Francillon, and D. Balzarotti. The role of phone numbers in understanding cyber-crime. In *PST 2013, 11th International Conference on Privacy, Security and Trust, July 10-12, 2013, Tarragona, Catalonia, Spain*, 2013.

[136] C. Cowan, C. Pu, D. Maier, M. Hinton, J. Walpole, P. Bakke, S. Beattie, A. Grier, P. Wagle, and Q. Zhang. Stackguard: Automatic adaptive detection and prevention of buffer-overflow attacks. In *Proceedings of the 7th USENIX Security Symposium*, January 1998.

[137] C. Cowan, P. Wagle, C. Pu, S. Beattie, and J. Walpole. Buffer overflows: attacks and defenses for the vulnerability of the decade. In *Proceedings of the DARPA Information Survivability Conference and Exposition (DISCEX)*, 2000.

[138] J. R. Crandall and F. T. Chong. Minos: Control Data Attack Prevention Orthogonal to Memory Model. In *Proceedings of the 37th annual IEEE/ACM International Symposium on Microarchitecture*, MICRO 37, 2004.

[139] A. Cui and M. Costello. Hacking Cisco Phones, 2012.

[140] A. Cui, M. Costello, and S. J. Stolfo. When firmware modifications attack: A case study of embedded exploitation. In *Proceedings of the ISOC Symposium on Network and Distributed Systems Security (NDSS)*, 2013.

[141] A. Cui and S. J. Stolfo. A quantitative analysis of the insecurity of embedded network devices: results of a wide-area scan. In *Proceedings of the 26th Annual Computer Security Applications Conference (ACSAC)*, 2010.

[142] L. Cutillo, R. Molva, and T. Strufe. Safebook: A privacy-preserving online social network leveraging on real-life trust. *IEEE Communications Magazine, Consumer Communications and Networking Series*, 47(12):94–101, 2009.

[143] CWE/SANS. CWE/SANS TOP 25 Most Dangerous Software Errors. `www.sans.org/top25-software-errors`, 2011.

[144] D. Dagon, G. Gu, C. P. Lee, and W. Lee. A Taxonomy of Botnet Structures. In *Proceedings of the 23rd Annual Computer Security Applications Conference*, 2007.

[145] G. Danezis and P. Mittal. Sybilinfer: Detecting sybil nodes using social networks. In *Network and Distributed System Security Symposium - NDSS*, 2009.

[146] S. K. Das, K. Kant, and N. Zhang. Handbook on Securing Cyber-Physical Critical Infrastructure. 2012.

[147] C. R. Davis, S. Neville, J. M. FernÃąndez, J.-M. Robert, and J. McHugh. Structured Peer-to-Peer Overlay Networks: Ideal Botnet Command and Control Infrastructures? In *Proceedings of the 13th European Symposium on Research in Computer Security*, 2008.

[148] S. J. Delany, M. Buckley, and D. Greene. Review: Sms spam filtering: Methods and data. *Expert Syst. Appl.*, 39(10):9899–9908, Aug. 2012.

[149] d. e. denning and p. j. denning. certification of programs for secure information flow. *commun. acm*, 20(7):504–513, 1977.

[150] D. Desai. Malware Analysis Report: Trojan: AndroidOS/Zitmo, Semptember 2011. http://www.kindsight.net/sites/default/files/android_trojan_zitmo_final_pdf_17585.pdf.

[151] S. Designer. Non-executable user stack. http://www.openwall.com/linux/.

[152] A. Dey and S. Weis. Pseudoid: Enhancing privacy in federated login. In *Hot Topics in Privacy Enhancing Technologies*, 2010.

[153] R. Dhamija, J. Tygar, and M. Hearst. Why Phishing Works. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, pages 581–590. ACM New York, NY, USA, 2006.

[154] R. Dhamija, J. D. Tygar, and M. Hearst. Why phishing works. In *Proceedings of the SIGCHI conference on Human Factors in computing systems*. ACM, 2006.

[155] T. Dierks and E. Rescorla. The Transport Layer Security (TLS) Protocol Version 1.2. IETF, August 2008. https://tools.ietf.org/html/rfc5246.

[156] Q. Ding, N. Katenka, P. Barford, E. Kolaczyk, and M. Crovella. Intrusion as (anti)social communication. In *Proceedings of the 18th ACM SIGKDD conference on Knowledge discovery and data mining - KDD '12*, page 886, 2012.

[157] D. Dittrich. So You Want to Take Over a Botnet. In *Proceedings of the 5th USENIX conference on Large-Scale Exploits and Emergent Threats*, 2012.

[158] P. Durance and M. Godet. Scenario building: Uses and abuse. *Technological Forecasting and Social Change*, 77:1488–1492, 2010.

[159] T. Durden. Bypassing PaX ASLR protection. *Phrack*, 11(59), 2002.

[160] P. Eckersley. How unique is your web browser? In *Proceedings of the Privacy Enhancing Technologies Symposium (PETS)*, 2010.

[161] M. Egele. Invited talk: The state of mobile security. In *DIMVA*, 2012.

[162] M. Egele, C. Kruegel, E. Kirda, and G. Vigna. PiOS: Detecting privacy leaks in ios applications. In *NDSS*, 2011.

[163] M. Egele, A. Moser, C. Kruegel, and E. Kirda. Pox: Protecting users from malicious facebook applications. In *Proceedings of 3rd Workshop on Security and Social Networking*, 2011.

[164] M. Egele, G. Stringhini, C. Kruegel, and G. Vigna. Compa: Detecting compromised accounts on social networks. In *ISOC Network and Distributed System Security Symposium (NDSS)*, 2013.

[165] M. W. Eichin and J. A. Rochlis. With Microscope and Tweezers: An Analysis of the Internet Virus of November 1988. In *Proceedings of the IEEE Symposium on Security & Privacy*, pages 326–344, 1989.

[166] EMC. The Digital Universe is Still Growing. http://www.emc.com/leadership/programs/digital-universe.htm, 2009.

[167] W. Enck, P. Gilbert, B. Chun, L. Cox, J. Jung, P. McDaniel, and A. Sheth. Taintdroid: an information-flow tracking system for realtime privacy monitoring on smartphones. In *Proc. of USENIX OSDI*, 2010.

[168] W. Enck, D. Octeau, P. McDaniel, and S. Chaudhuri. A study of android application security. In *USENIX Security*, Berkeley, CA, USA, 2011.

[169] W. Enck, M. Ongtang, and P. McDaniel. On lightweight mobile phone application certification. In *CCS*, 2009.

[170] Ú. Erlingsson. Low-level software security: Attack and defenses. Technical Report MSR-TR-07-153, Microsoft Research, 2007. http://research.microsoft.com/pubs/64363/tr-2007-153.pdf.

[171] European Commission. Cybersecurity Strategy of the European Union: An Open, Safe and Secure Cyberspace, Feb. 2013.

[172] D. Evans. Top 25 technology predictions. Technical report, CISCO, 2009.

[173] Facebook. Social Authentication. 2011. http://www.facebook.com/blog.php?post=486790652130.

[174] Facebook. Social Plugins. 2013. http://developers.facebook.com/docs/plugins/.

[175] N. Falliere. Sality: Story of a Peer-to-Peer Viral Network, 2011. Technical Report by Symantec Labs.

[176] N. Falliere, L. O. Murchu, and E. Chien. W32. stuxnet dossier. Technical report, Symantec Corporation, 2011.

[177] A. Felt and D. Evans. Privacy protection for social networking platforms. In *Proceedings of the 2008 Workshop on Web 2.0 Security and Privacy*.

[178] A. P. Felt, E. Chin, S. Hanna, D. Song, and D. Wagner. Android permissions demystified. In *CCS*, 2011.

[179] J. Fildes. Stuxnet worm 'targeted high-value iranian assets'. http://www.bbc.co.uk/news/technology-11388018, September 2010.

[180] J. Fildes. Stuxnet virus targets and spread revealed. http://www.bbc.co.uk/news/technology-12465688, February 2011.

[181] D. Florencio and C. Herley. A large-scale study of web password habits. In *Proceedings of the 16th international conference on World Wide Web*, WWW '07, pages 657–666, New York, NY, USA, 2007. ACM.

[182] D. Florencio and C. Herley. Is everything we know about password stealing wrong? *IEEE Security & Privacy*, 10(6):63–69, 2012.

[183] S. Forrest, A. Somayaji, and D. Ackley. Building diverse computer systems. In *Proceedings of the 6th Workshop on Hot Topics in Operating Systems (HotOS-VI)*, 1997.

[184] M. Frantzen and M. Shuey. Stackghost: Hardware facilitated stack protection. In *Proceedings of the* 10^{th} *USENIX Security Symposium*, August 2001.

[185] M. Fredrikson and B. Livshits. RePriv: Re-envisioning in-browser privacy. In *IEEE Symposium on Security and Privacy*, May 2011.

[186] G. Fresi Roglia, L. Martignoni, R. Paleari, and D. Bruschi. Surgically returning to randomized lib(c). In *Proceedings of the 25th Annual Computer Security Applications Conference (ACSAC)*, 2009.

[187] T. Garfinkel and M. Rosenblum. A Virtual Machine Introspection Based Architecture for Intrusion Detection. In *Proc. of NDSS*, 2003.

[188] S. Garriss, M. Kaminsky, M. J. Freedman, B. Karp, D. Mazières, and H. Yu. RE: Reliable Email. *3rd Symposium on Networked Systems Design and Implementation*, 2006.

[189] S. Gaw and E. W. Felten. Password management strategies for online accounts. In *Proceedings of the second symposium on Usable privacy and security*, SOUPS '06, pages 44–55, New York, NY, USA, 2006. ACM.

[190] M. Gegick, L. Williams, J. Osborne, and M. Vouk. Prioritizing software security fortification through code-level metrics. In *Proc. of the 4th ACM workshop on Quality of protection*, QoP'08. ACM Press, Oct. 2008.

[191] B. Genge, C. Siaterlis, and M. Hohenadel. AMICI: An Assessment Platform for Multi-Domain Security Experimentation on Critical Infrastructures. *ibs.ro*, 2012.

[192] B. Genge, C. Siaterlis, I. Nai Fovino, and M. Masera. A cyber-physical experimentation environment for the security analysis of networked industrial control systems. *Computers & Electrical Engineering*, 38(5):1146–1161, Sept. 2012.

[193] GeogiaTech Research Institute and GeorgiaTech Information Security Center. Emerging cyber threats report 2013. Internet. http://www.gatech.edu/newsroom/release.html?nid=170981, 2012.

[194] V. George, T. Piazza, and H. Jiang. Technology Insight: Intel©Next Generation Microarchitecture Codename Ivy Bridge. `www.intel.com/idf/library/pdf/sf_2011/SF11_SPCS005_101F.pdf`, September 2011.

[195] P. Godefroid, M. Y. Levin, and D. A. Molnar. Automated Whitebox Fuzz Testing. In *Proceedings of the 15th Annual Network and Distributed System Security Symposium*, NDSS'08, 2008.

[196] I. Goldberg, D. Wagner, R. Thomas, and E. A. Brewer. A secure environment for untrusted helper applications (confining the wily hacker). In *Proceedings of the 5th USENIX Security Symposium*, 1996.

[197] L. H. Gomes, R. B. Almeida, and L. M. A. Bettencourt. Comparative Graph Theoretical Characterization of Networks of Spam and Legitimate Email. In *Conference on Email and Anti-Spam (CEAS)*, 2005.

[198] C. Grier, S. Tang, and S. King. Secure Web Browsing with the OP Web Browser. In *Security and Privacy, 2008.*, pages 402–416. IEEE, 2008.

[199] E. Grosse. Gmail account security in Iran. Google Blog, September 2011. `http://googleonlinesecurity.blogspot.com/2011/09/gmail-account-security-in-iran.html`.

[200] T. Guardian. China suspected of facebook attack on nato's supreme allied commander. `http://www.guardian.co.uk/world/2012/mar/11/china-spies-facebook-attack-nato`.

[201] M. V. Gundy and H. Chen. Noncespaces: Using Randomization to Enforce Information Flow Tracking and Thwart Cross-Site Scripting Attacks. In *Proceedings of the 16th Annual Network and Distributed System Security Symposium (NDSS)*, San Diego, CA, Feb. 8-11, 2009.

[202] P. Gutmann and I. Grigg. Security usability. *Security & Privacy, IEEE*, 3(4):56–58, 2005.

[203] D. Hadziosmanovic, D. Bolzoni, P. Hartel, and S. Etalle. MELISSA: Towards Automated Detection of Undesirable User Actions in Critical Infrastructures. 2011.

[204] D. Hadziosmanovic, D. Bolzoni, and P. H. Hartel. A log mining approach for process monitoring in SCADA. *International Journal of Information Security*, 11(4):231–251, Apr. 2012.

[205] J. Haldeman. *The Forever War*. S. F. Masterworks Series. Orion, 2011.

[206] D. Halperin, T. Kohno, T. Heydt-Benjamin, K. Fu, and W. Maisel. Security and privacy for implantable medical devices. *Pervasive Computing, IEEE*, 7(1):30 –39, jan.-march 2008.

[207] S. Hanna, L. Huang, E. X. Wu, S. Li, C. Chen, and D. Song. Juxtapp: A scalable system for detecting code reuse among android applications. In *DIMVA*, 2012.

[208] M. Hayes, A. Walenstein, and A. Lakhotia. Evaluation of Malware Phylogeny Modelling Systems Using Automated Variant Generation. *Journal in Computer Virology*, 5(4):335–343, 2009.

[209] A. Ho, M. Fetterman, C. Clark, A. Warfield, and S. Hand. Practical taint-based protection using demand emulation. In *Proceedings of the 1st ACM SIGOPS/EuroSys European Conference on Computer Systems*, EuroSys'06, 2006.

[210] P. Holme, B. J. Kim, C. N. Yoon, and S. K. Han. Attack Vulnerability of Complex Networks. *Physical Review E, vol. 65*, 2002.

[211] T. Holz, M. Steiner, F. Dahl, E. Biersack, and F. Freiling. Measurements and Mitigation of Peer-to-Peer-based Botnets: A Case Study on Storm Worm. In *Proceedings of the 1st USENIX Workshop on Large-Scale Exploits and Emergent Threats*, 2008.

[212] M. Honan. How apple and amazon security flaws led to my epic hacking. Wired Magazine, August 2012. `http://www.wired.com/gadgetlab/2012/08/apple-amazon-mat-honan-hacking/`.

[213] R. Hund, M. Hamann, and T. Holz. Towards Next-Generation Botnets. In *Proceedings of the 2008 European Conference on Computer Network Defense*, 2008.

[214] ICS-CERT. Monthly monitor, march 2012. http://ics-cert.us-cert.gov/sites/default/files/ICS-CERT_Monthly_Monitor_March_2012_0.pdf.

[215] V. Igure and R. Williams. Taxonomies of attacks and vulnerabilities in computer systems. *Communications Surveys Tutorials, IEEE*, 10(1):6–19, 2008.

[216] Imperva ADC Team. Security trends 2013. Internet. http://blog.imperva.com/2012/12/security-trends-2013-trend-1.html, 2012.

[217] Interagency Working Group on Cyber Security and Information Assurance. Federal plan for cyber security and information assurance research and development, April 2006.

[218] D. Irani, M. Balduzzi, D. Balzarotti, E. Kirda, and C. Pu. Reverse Social Engineering Attacks in Online Social Networks. In *Proceedings of the 8th international conference on Detection of intrusions and malware, and vulnerability assessment (DIMVA)*, 2011.

[219] J. Isacenkova, O. Thonnard, A. Costin, D. Balzarotti, and A. Francillon. Inside the SCAM jungle: A closer look at 419 scam email operations. In *IWCC 2013, International Workshop on Cyber Crime (co-located with the 34th IEEE Symposium on Security and Privacy (IEEE S&P 2013))*, 2013.

[220] Iseclab. Anubis. http://anubis.iseclab.org.

[221] C. Jackson, A. Bortz, D. Boneh, and J. C. Mitchell. Protecting browser state from web privacy attacks. In *Proceedings of the 15th International World Wide Web Conference (WWW)*, 2006.

[222] D. Jacoby. Facebook security phishing attack in the wild. http://www.securelist.com/en/blog/208193325/Facebook_Security_Phishing_Attack_In_The_Wild.

[223] T. Jim, N. Swamy, and M. Hicks. Defeating Script Injection Attacks with Browser-Enforced Embedded Policies. In *WWW '07: Proceedings of the 16th international conference on World Wide Web*, pages 601–610, New York, NY, USA, 2007. ACM.

[224] R. Johnson. A castle made of sand: Adobe Reader X sandbox. CanSecWest, 2011.

[225] M. Jurek. Google Explores +1 Button To Influence Search Results. 2011. http://www.tekgoblin.com/2011/08/29/google-explores-1-button-to-influence-search-results/.

[226] E. Kalige. A Case Study of Eurograbber: How 36 Million Euros were Stolen via Malware. https://www.checkpoint.com/products/downloads/whitepapers/Eurograbber_White_Paper.pdf.

[227] S. Kamkar. Evercookie. http://samy.pl/evercookie/.

[228] B. B. H. Kang, E. Chan-Tin, C. P. Lee, J. Tyra, H. J. Kang, C. Nunnery, Z. Wadler, G. Sinclair, N. Hopper, D. Dagon, and Y. Kim. Towards Complete Node Enumeration in a Peer-to-Peer Botnet. In *Proceedings of the 4th International Symposium on Information, Computer, and Communications Security*, 2009.

[229] C. Kanich, K. Levchenko, B. Enright, G. M. Voelker, and S. Savage. The Heisenbot Uncertainty Problem: Challenges in Separating Bots from Chaff. In *Proceedings of the 1st USENIX Workshop on Large-Scale Exploits and Emergent Threats*, 2008.

[230] M. Karim, A. Walenstein, A. Lakhotia, and L. Parida. Malware Phylogeny Generation Using Permutations of Code. *Journal in Computer Virology*, 1(1):13–23, 2005.

[231] Kaspersky. Teamwork: How the ZitMo Trojan Bypasses Online Banking Security, October 2011. http://www.kaspersky.com/about/news/virus/2011/Teamwork_How_the_ZitMo_Trojan_Bypasses_Online_Banking_Security.

[232] Kaspersky Labs. Kaspersky security bulletin 2012. malware evolution. Internet. http://usa.kaspersky.com/about-us/press-center/press-releases/kaspersky-lab-outlines-key-security-trends-2012-predicts-core-t, 2012.

[233] G. S. Kc, A. D. Keromytis, and V. Prevelakis. Countering code-injection attacks with instruction-set randomization. In *Proceedings of the 10th ACM conference on Computer and communications security (CCS)*, pages 272–280, 2003.

[234] P. G. Kelley, S. Komanduri, M. L. Mazurek, R. Shay, T. Vidas, L. Bauer, N. Christin, L. F. Cranor, and J. Lopez. Guess again (and again and again): Measuring password strength by simulating password-cracking algorithms. In *Proceedings of the 2012 IEEE Symposium on Security and Privacy*, Oakland '12, pages 523–537, Washington, DC, USA, 2012. IEEE Computer Society.

[235] V. P. Kemerlis, G. Portokalidis, K. Jee, and A. D. Keromytis. libdft: Practical Dynamic Data Flow Tracking for Commodity Systems. In *Proc. of VEE*, 2012.

[236] A. Ki-Chan and H. Dong-Joo. Malware migrating to gaming consoles, 2010. `http://defcon.org/images/defcon-18/dc-18-presentations/Chan-Joo/DEFCON-18-Chan-Joo-Malware-to-Gaming-Consoles.pdf`.

[237] J. Kinder, F. Zuleger, and H. Veith. An abstract interpretation-based framework for control flow reconstruction from binaries. In *Proceedings of the 10th International Conference on Verification, Model Checking, and Abstract Interpretation*, VMCAI'09, 2009.

[238] V. Kiriansky, D. Bruening, and S. Amarasinghe. Secure execution via program shepherding. In *Proceedings of the 11th USENIX Security Symposium*, 2002.

[239] G. Kontaxis, M. Polychronakis, A. D. Keromytis, and E. P. Markatos. Privacy-Preserving Social Plugins. In *Security'12 Proceedings of the 21st USENIX conference on Security symposium*, 2012.

[240] J. Koo, X. Lin, and S. Bagchi. PRIVATUS: Wallet-Friendly Privacy Protection for Smart Meters. pages 343–360. Springer Berlin Heidelberg, Berlin, Heidelberg, 2012.

[241] M. Kosinski, D. Stillwell, and T. Graepel. Private traits and attributes are predictable from digital records of human behavior. *Proceedings of the National Academy of Sciences*, pages 2–5, Mar. 2013.

[242] B. Krishnamurthy and C. E. Wills. Characterizing privacy in online social networks. In *Proceeedings of the Workshop on Online Social Networks*, 2008.

[243] N. Kroes. Internet security: everyone's responsibility, Feb. 2012. `http://europa.eu/rapid/press-release_SPEECH-12-68_en.htm`.

[244] N. Kshetri. The Global Cybercrime Industry. Springer, 2010.

[245] S. Kuo. Execute disable bit functionality blocks malware code execution, 2005. Intel Corp. `http://cache-www.intel.com/cd/00/00/14/93/149307_149307.pdf`.

[246] M. La Polla, F. Martinelli, and D. Sgandurra. A survey on security for mobile devices. *Communications Surveys Tutorials, IEEE*, 15(1):446 –471, quarter 2013.

[247] H.-y. Lam and D.-y. Yeung. A Learning Approach to Spam Detection based on Social Networks. In *Conference on Email and Anti-Spam (CEAS)*, 2007.

[248] L. Lamport. Password authentication with insecure communication. *Commun. ACM*, 24(11):770–772, Nov 1981.

[249] S. Landau, M. R. Stytz, C. E. Landwehr, and F. B. Schneider. Overview of Cyber Security: A Crisis of Prioritization. *IEEE Security and Privacy*, 03(3):9–11, 2005.

[250] R. Langner. Enumerating Stuxnet's exploits, 2011. `http://www.langner.com/en/2011/06/07/enumerating-stuxnet%E2%80%99s-exploits/`.

[251] A. Lelli. Zeusbot/Spyeye P2P Updated, Fortifying the Botnet, 2012. Technical Report by Symantec Labs: `http://www.symantec.com/connect/node/2135671`.

[252] J. Leskovec, J. Kleinberg, and C. Faloutsos. Graphs Over Time: Densification Laws, Shrinking Diameters and Possible Explanations. In *Proceeding of the eleventh ACM SIGKDD international conference on Knowledge discovery in data mining - KDD '05*, page 177. ACM Press, 2005.

[253] C. Lever, M. Antonakakis, B. Reaves, P. Traynor, and W. Lee. The Core of the Matter: Analyzing Malicious Traffic in Cellular Carriers. In *NDSS*, 2013.

[254] J. Leyden. Conficker left Manchester unable to issue traffic tickets, 2009. `http://www.theregister.co.uk/2009/07/01/conficker_council_infection/`.

[255] J. Leyden. London hospital recovers from Conficker outbreak. The Register, `http://www.theregister.co.uk/2009/08/24/nhs_hospital_conficker/`, 2009.

[256] H. Li. Understanding and exploiting Flash ActionScript vulnerabilities. CanSecWest, 2011.

[257] H. Lin-Shung, W. Zack, E. Chris, and J. Collin. Protecting Browsers from Cross-Origin CSS Attacks. In *CCS 10: Proceedings of the 17th ACM Conference on Computer and Communications Security*, New York, NY, USA, 2010. ACM.

[258] B. Liu, L. Shi, Z. Cai, and M. Li. Software vulnerability discovery techniques: A survey. In *Proceedings of the 4th International Conference on Multimedia Information Networking and Security (MINES)*, pages 152–156, 2012.

[259] Y. Liu, P. Ning, and M. K. Reiter. False data injection attacks against state estimation in electric power grids. *Transactions on Information and System Security (TISSEC*, 14(1), May 2011.

[260] H. Lockheimer. Bouncer. `http://googlemobile.blogspot.it/2012/02/android-and-security.html`.

[261] Lookout. 2013 mobile threat predictions. https://blog.lookout.com/blog/2012/12/13/2013-mobile-threat-predictions/.

[262] M. M. Lucas and N. Borisov. Flybynight: mitigating the privacy risks of social networking. In *Proceedings of the 7th ACM workshop on Privacy in the electronic society*, 2008.

[263] W. Luo, Q. Xie, and U. Hengartner. Facecloak: An architecture for user privacy on social networking sites. In *Proceedings of the International Conference on Computational Science and Engineering*, 2009.

[264] F. Maggi. Are the con artists back? a preliminary analysis of modern phone frauds. In *Proceedings of the International Conference on Computer and Information Technology (CIT)*, pages 824–831. IEEE Computer Society, 2010.

[265] F. Maggi, A. Frossi, S. Zanero, G. Stringhini, B. Stone-Gross, C. Kruegel, and G. Vigna. Two years of short urls internet measurement: security threats and countermeasures. In *Proceedings of the 22nd international conference on World Wide Web*, WWW '13, pages 861–872, Republic and Canton of Geneva, Switzerland, 2013. International World Wide Web Conferences Steering Committee.

[266] F. Maggi, A. Sisto, and S. Zanero. A social-engineering-centric data collection initiative to study phishing. In *Proceedings of the First Workshop on Building Analysis Datasets and Gathering Experience Returns for Security (BADGERS)*, pages 107–108, New York, NY, USA, 2011. ACM.

[267] T. Mai. Android Reaches 500 Million Activations Worldwide. http://www.tomshardware.com/news/Google-Android-Activation-half-billion-Sales,17556.html, 2012.

[268] L. Marinos and A. Sfakianakis. ENISA Threat Landscape. Technical report, ENISA, Sept. 2012.

[269] Mashable. What is the syrian electronic army? `http://mashable.com/2012/08/10/syrian-electronic-army/`.

[270] D. Mashima and A. A. Cárdenas. Evaluating Electricity Theft Detectors in Smart Grid Networks. *Research on Attacks, Intrusions and Defences Symposium*, 2012.

[271] McAfee Labs. 2013 threats predictions, 2012.

[272] G. McDonald, L. O. Murchu, S. Doherty, and E. Chien. Stuxnet 0.5: The Missing Link, Feb. 2013. http://www.symantec.com/content/en/us/enterprise/media/security_response/whitepapers/stuxnet_0_5_the_missing_link.pdf.

[273] M. McGlohon, S. Bay, M. G. Anderle, D. M. Steier, and C. Faloutsos. SNARE: A Link Analytic System for Graph Labeling and Risk Detection. In *Proceedings of the 15th ACM SIGKDD international conference on Knowledge discovery and data mining - KDD '09*, page 1265. ACM Press, 2009.

[274] S. McLaughlin and P. McDaniel. SABOT: specification-based payload generation for programmable logic controllers. In *ACM conference on Computer and Communications Security*. ACM Request Permissions, Oct. 2012.

[275] R. McMillan. Stolen Twitter accounts can fetch $1000. 2010. http://www.computerworld.com/s/article/9150001/Stolen_Twitter_accounts_can_fetch_1_000.

[276] K. McNamee. Malware Analysis Report: ZeroAccess/Sirefef, 2012. Technical Report by Kindsight Security Labs.

[277] M. Miculan and C. Urban. Formal analysis of facebook connect single sign-on authentication protocol. In *SOFSEM*, volume 11, pages 22–28, 2011.

[278] B. Miller and D. Rowe. A survey of SCADA and critical infrastructure incidents. In *Annual conference on Research In Information Technology*. ACM Request Permissions, Oct. 2012.

[279] M. Miller, T. Burrell, and M. Howard. Mitigating software vulnerabilities, July 2011. http://www.microsoft.com/download/en/details.aspx?displaylang=en&id=26788.

[280] Z. Minchev and V. Shalamanov. Scenario Generation and Assessment Framework Solution in Support of the Comprehensive Approach. In *RTO-MP-SAS-081, Symposium on "Analytical Support to Defence Transformation", Boyana, Bulgaria, April 26-28*, pages 22–1–22–16, 2010.

[281] D. Misener. Rise of the socialbots: They could be influencing you online. 2011. http://www.cbc.ca/news/technology/story/2011/03/29/f-vp-misener-socialbot-armies-election.html.

[282] A. Mislove, M. Marcon, K. P. Gummadi, P. Druschel, and B. Bhattacharjee. Measurement and analysis of online social networks. In *Proceedings of the 7th ACM SIGCOMM conference on Internet measurement - IMC '07*, page 29. ACM Press, 2007.

[283] A. Mislove, A. Post, and P. Druschel. Ostra: Leveraging trust to thwart unwanted communication. In *Proceedings of the 5th USENIX Symposium on Networked Systems Design and Implementation*, number i, pages 15–30, 2008.

[284] Mitre. Common Vulnerabilities and Exposures (CVE). http://cve.mitre.org/, 2011.

[285] A. Mohaisen, A. Yun, and Y. Kim. Measuring the mixing time of social graphs. In *Proceedings of the 10th annual conference on Internet measurement - IMC '10*, page 383. ACM Press, 2010.

[286] M. Moore. Houses of Parliament computers infected with Conficker virus, 2009. http://www.telegraph.co.uk/technology/microsoft/5057605/Houses-of-Parliament-computers-infected-with-Conficker-virus.html.

[287] F. Moradi, T. Olovsson, and P. Tsigas. An Evaluation of Community Detection Algorithms on Large-Scale Email Traffic. In *11th International Symposium on Experimental Algorithms*. Springer-Verlag, 2012.

[288] F. Moradi, T. Olovsson, and P. Tsigas. Towards modeling legitimate and unsolicited email traffic using social network properties. In *Proceedings of the Fifth Workshop on Social Network Systems - SNS '12*, 2012.

[289] E. Morozov. Swine flu: Twitter's power to misinform. 2009. http://neteffect.foreignpolicy.com/posts/2009/04/25/swine_flu_twitters_power_to_misinform.

[290] T. Morris, A. Srivastava, B. Reaves, and W. Gao. A control system testbed to validate critical infrastructure protection concepts. *... Infrastructure Protection ...*, 2011.

[291] Mozilla. Browserid specification. https://github.com/mozilla/id-specs/blob/prod/browserid/index.md.

[292] Mozilla. Verified e-mail protocol. https://wiki.mozilla.org/Labs/Identity/VerifiedEmailProtocol.

[293] Y. Nadji, P. Saxena, and D. Song. Document Structure Integrity: A Robust Basis for Cross-site Scripting Defense. In *Proceedings of the 16th Annual Network and Distributed System Security Symposium (NDSS)*, San Diego, CA, Feb. 8-11, 2009.

[294] NakedSecurity. Facebook glitch lets spear phishers impersonate users' friends and family. http://nakedsecurity.sophos.com/2012/08/31/facebook-glitch-spear-phishing/.

[295] A. Nappa, A. Fattori, M. Balduzzi, M. Dell'Amico, and L. Cavallaro. Take a Deep Breath: a Stealthy, Resilient and Cost-Effective Botnet Using Skype. In *GI SIG SIDAR Conference on Detection of Intrusions and Malware & Vulnerability Assessment (DIMVA)*, July 2010.

[296] J. Nazario and T. Holz. As the Net Churns: Fast-Flux Botnet Observations Tracking Fast-Flux Domains. In *Proceedings of the 3rd International Conference on Malicious and Unwanted Software*, 2008.

[297] D. Nebenzahl and M. Sagiv. Install-time vaccination of windows executables to defend against stack smashing attacks. *IEEE Transactions on Dependable and Secure Computing*, 3(1):78–90, 2006.

[298] M. Newman and J. Park. Why social networks are different from other types of networks. *Physical Review E*, 68(3), Sept. 2003.

[299] J. Newsome and D. Song. Dynamic taint analysis for automatic detection, analysis, and signature generation of exploit attacks on commodity software. In *Proceedings of the Network and Distributed Systems Security Symposium (NDSS)*, 2005.

[300] V. H. Nguyen and L. M. S. Tran. Predicting vulnerable software components with dependency graphs. In *Proc. of the 6th International Workshop on Security Measurements and Metrics*, MetriSec'10. ACM Press, Sept. 2010.

[301] NIST. National Vulnerability Database. http://web.nvd.nist.gov/view/vuln/search, 2011.

[302] J. Oberheide, M. Bailey, and F. Jahanian. Polypack: an automated online packing service for optimal antivirus evasion. In *Proceedings of the 3rd USENIX Workshop on Offensive Technologies (WOOT)*, 2009.

[303] J. Oberheide and C. Miller. Dissecting the Android's Bouncer. *SummerCon*, 2012. http://jon.oberheide.org/files/summercon12-bouncer.pdf.

[304] J. L. Obes and J. Schuh. A Tale of Two Pwnies (Part 1), 2012. http://blog.chromium.org/2012/05/tale-of-two-pwnies-part-1.html.

[305] G. Ollmann. The vishing guide. Technical report, IBM Global Technology Services, 2007. http://www.infosecwriters.com/text_resources/pdf/IBM_ISS_vishing_guide_GOllmann.pdf.

[306] K. Onarlioglu, L. Bilge, A. Lanzi, D. Balzarotti, and E. Kirda. G-Free: defeating return-oriented programming through gadget-less binaries. In *Proceedings of the 26th Annual Computer Security Applications Conference (ACSAC)*, 2010.

[307] K. Onarlioglu, U. O. Yilmaz, E. Kirda, and D. Balzarotti. Insights into user behavior in dealing with internet attacks. In *Network and Distributed Systems Security Symposium (NDSS)*, 2012.

[308] V. Pappas, M. Polychronakis, and A. D. Keromytis. Smashing the gadgets: Hindering return-oriented programming using in-place code randomization. In *Proceedings of the 33rd IEEE Symposium on Security & Privacy (S&P)*, 2012.

[309] E. Passerini, R. Paleari, L. Martignoni, and D. Bruschi. Fluxor: Detecting and monitoring fast-flux service networks. In D. Zamboni, editor, *DIMVA*, volume 5137 of *Lecture Notes in Computer Science*, pages 186–206. Springer, 2008.

[310] PaX Team. Address space layout randomization. `http://pax.grsecurity.net/docs/aslr.txt`.

[311] PaX Team. PaX non-executable pages design & implementation. `http://pax.grsecurity.net/docs/noexec.txt`.

[312] H. Peng, C. Gates, B. Sarma, N. Li, Y. Qi, R. Potharaju, C. Nita-Rotaru, and I. Molloy. Using probabilistic generative models for ranking risks of android apps. In *CCS*, 2012.

[313] C. Percival. Stronger key derivation via sequential memory-hard functions. *BSDCan 2009*, 2009.

[314] L. Piètre-Cambacédès, M. Tritschler, and G. N. Ericsson. Cybersecurity Myths on Power Control Systems: 21 Misconceptions and False Beliefs. *Power Delivery, IEEE Transactions on*, 26(1), 2011.

[315] M. Pistoia, S. Chandra, S. J. Fink, and E. Yahav. A survey of static analysis methods for identifying security vulnerabilities in software systems. *IBM Syst. J.*, 46(2):265–288, Apr. 2007.

[316] PITAC. Cyber Security: A Crisis of Prioritization. Technical report, President's Information Technology Advisory Committee - PITAC, Feb. 2005.

[317] I. Polakis, G. Kontaxis, S. Antonatos, E. Gessiou, T. Petsas, and E. P. Markatos. Using social networks to harvest email addresses. In *Proceedings of the 9th Annual ACM Workshop on Privacy in the Electronic Society*, pages 11–20. ACM, 2010.

[318] I. Polakis, M. Lancini, G. Kontaxis, F. Maggi, S. Ioannidis, A. D. Keromytis, and S. Zanero. All your face are belong to us. In *Proceedings of the 28th Annual Computer Security Applications Conference on - ACSAC '12*, page 399. ACM Press, 2012.

[319] P. Porras, H. Saidi, and V. Yegneswaran. An analysis of conficker's logic and rendezvous points. Technical Report SRI International Technical Report, 2009. `http://mtc.sri.com/Conficker`.

[320] G. Portokalidis, A. Slowinska, and H. Bos. Argos: an Emulator for Fingerprinting Zero-Day Attacks. In *Proceedings of the 1st ACM SIGOPS/EuroSys European Conference on Computer Systems 2006*, EuroSys'06, 2006.

[321] M. Prasad and T. cker Chiueh. A binary rewriting defense against stack based overflow attacks. In *Proceedings of the USENIX Annual Technical Conference*, June 2003.

[322] R. W. Proctor, M.-C. Lien, K.-P. L. Vu, E. E. Schultz, and G. Salvendy. Improving computer security for authentication of users: influence of proactive password restrictions. *Behav Res Methods Instrum Comput*, 34(2):163–9, 2002.

[323] N. Provos and D. Mazières. A future-adaptive password scheme. ATEC 1999.

[324] S. J. Prowell, M. Pleszkoch, K. D. Sayre, and R. C. I. S. G. T. I. . I. P. Linger. Automated vulnerability detection for compiled smart grid software. *Innovative Smart Grid Technologies (ISGT), 2012 IEEE PES*, 2012.

[325] F. Qin, C. Wang, Z. Li, H.-s. Kim, Y. Zhou, and Y. Wu. LIFT: A Low-Overhead Practical Information Flow Tracking System for Detecting Security Attacks. In *Proc. of MICRO*, pages 135–148, 2006.

[326] C. Queiroz, A. Mahmood, and Z. S. G. I. T. o. Tari. SCADASim—A Framework for Building SCADA Simulations. *Smart Grid, IEEE Transactions on*, 2(4), 2011.

[327] M. Raciti and S. Nadjm-Tehrani. Embedded Cyber-Physical Anomaly Detection in Smart Meters. 2012.

[328] R. Radvanovsky and J. Brodsky. *Handbook of Scada/Control Systems Security*. CRC PressI Llc, Feb. 2013.

[329] F. Raja, K. Hawkey, S. Hsu, K.-L. C. Wang, and K. Beznosov. A brick wall, a locked door, and a bandit: a physical security metaphor for firewall warnings. In *Proceedings of the Seventh Symposium on Usable Privacy and Security*, SOUPS '11, pages 1:1–1:20, New York, NY, USA, 2011. ACM.

[330] V. Reding. the EU Data Protection Reform 2012: Making Europe the Standard Setter for Modern Data Protection Rules in the Digital Age. `http://europa.eu/rapid/press-release_SPEECH-12-26_en.htm`.

[331] A. Reina, A. Fattori, and L. Cavallaro. A system call-centric analysis and stimulation technique to automatically reconstruct android malware behaviors. In *EuroSec*, April 2013.

[332] C. Reis and S. Gribble. Isolating web programs in modern browser architectures. In *Proceedings of the 4th ACM European Conference on Computer Systems (EuroSys)*, pages 219–232. ACM, 2009.

[333] E. Rescorla. Security holes... Who cares? In *Proceedings of the 12th USENIX Security Symposium*, pages 75–90, Aug. 2003.

[334] R. Richmond. Stolen Facebook Accounts for Sale. 2010. `http://www.nytimes.com/2010/05/03/technology/internet/03facebook.html?_r=0`.

[335] R. Roberts. Malware Development Life Cycle. *Virus Bulletin Conf.*, (October), 2008.

[336] W. Robertson and G. Vigna. Static Enforcement of Web Application Integrity Through Strong Typing. In *Proceedings of the 18th USENIX Security Symposium*, Montreal, Quebec, August 2009.

[337] C. Rossow, D. Andriesse, T. Werner, B. Stone-Gross, C. J. Dietrich, and H. Bos. P2pwned — modeling and evaluating the resilience of peer-to-peer botnets. In *Security & Privacy (Oakland)*, San Francisco, CA, USA, May 2013.

[338] I. Rouf, H. Mustafa, M. Xu, W. Xu, R. Miller, and M. Gruteser. Neighborhood watch: security and privacy analysis of automatic meter reading systems. In *ACM conference on Computer and Communications Security*. ACM Request Permissions, Oct. 2012.

[339] RSA. Apt summit findings. `http://www.rsa.com/innovation/docs/APT_findings.pdf`.

[340] A. Rubin and D. Geer. A survey of web security. *Computer*, 31(9):34–41, 1998.

[341] G. Sarwar, O. Mehani, R. Boreli, and D. Kaafar. On the Effectiveness of Dynamic Taint Analysis for Protecting Against Private Information Leaks on Android-based Devices. In *10th International Conference on Security and Cryptography (SECRYPT)*, 2013.

[342] P. Saxena, S. Hanna, P. Poosankam, and D. Song. FLAX: Systematic Discovery of Client-side Validation Vulnerabilities in Rich Web Applications. In *Proceedings of the 17th Annual Network and Distributed System Security Symposium (NDSS)*.

[343] Security Focus. Vulnerabilities. `http://www.securityfocus.com/bid`, 2011.

[344] D. Seeley. Password cracking: a game of wits. *Commun. ACM*, 32(6):700–703, June 1989.

[345] R. Sekar. An Efficient Black-box Technique for Defeating Web Application Attacks. In *Proceedings of the 16th Annual Network and Distributed System Security Symposium (NDSS)*, San Diego, CA, Feb. 8-11, 2009.

[346] F. J. Serna. CVE-2012-0769, the case of the perfect info leak, Feb. 2012. `http://zhodiac.hispahack.com/my-stuff/security/Flash_ASLR_bypass.pdf`.

[347] H. Shacham, M. Page, B. Pfaff, E.-J. Goh, N. Modadugu, and D. Boneh. On the effectiveness of address-space randomization. In *Proceedings of the 11th ACM conference on Computer and Communications Security (CCS)*, 2004.

[348] F. T. Sheldon and C. Vishik. Moving toward trustworthy systems: R&d essentials. *IEEE Computer*, 2010.

[349] Y. Shin and L. Williams. An initial study on the use of execution complexity metrics as indicators of software vulnerabilities. In *Proc. of the 7th international workshop on Software engineering for secure systems*, SESS'11. ACM Press, May 2011.

[350] S. Sidiroglou and A. D. Keromytis. Countering network worms through automatic patch generation. *IEEE Security and Privacy*, 3(6):41–49, 2005.

[351] G. Sinclair, C. Nunnery, and B. Kang. The Waledac Protocol: The How and Why, 2009. Technical Report by Infrastructure Systems Research Lab/University of North Carolina.

[352] K. Singh, S. Bhola, and W. Lee. xbook: Redesigning privacy control in social networking platforms. In *Proceedings of the 18th USENIX Security Symposium*, 2009.

[353] M. Sirivianos, K. Kim, and X. Yang. SocialFilter: Introducing social trust to collaborative spam mitigation. *2011 Proceedings IEEE INFOCOM*, pages 2300–2308, Apr. 2011.

[354] A. Slowinska and H. Bos. The Age of Data: Pinpointing Guilty Bytes in Polymorphic Buffer Overflows on Heap or Stack. In *Proceedings of the 23rd Annual Computer Security Applications Conference*, ACSAC'07, 2007.

[355] A. Slowinska and H. Bos. Pointless tainting?: evaluating the practicality of pointer tainting. In *EuroSys*, pages 61–74, 2009.

[356] A. Slowinska, T. Stancescu, and H. Bos. Howard: a dynamic excavator for reverse engineering data structures. In *Proceedings of NDSS 2011*, San Diego, CA, 2011.

[357] A. Slowinska, T. Stancescu, and H. Bos. Body armor for binaries: preventing buffer overflows without recompilation. In *Proceedings of USENIX Annual Technical Conference*, Boston, MA, June 2012.

[358] M. Srivatsa and M. Hicks. Deanonymizing mobility traces. In *Proceedings of the 2012 ACM conference on Computer and communications security - CCS '12*, page 628. ACM Press, 2012.

[359] J. I. S. G. T. I. . I. P. Stamp. The SPIDERS project - Smart Power Infrastructure Demonstration for Energy Reliability and Security at US military facilities. *Innovative Smart Grid Technologies (ISGT), 2012 IEEE PES*, 2012.

[360] G. Starnberger, C. Kruegel, and E. Kirda. Overbot: A Botnet Protocol Based on Kademlia. In *Proceedings of the 4th International Conference on Security and Privacy in Communication Networks*, 2008.

[361] A. Stefanov and . I. P. Chen-Ching Liu Innovative Smart Grid Technologies ISGT. Cyber-power system security in a smart grid environment. *Innovative Smart Grid Technologies (ISGT), 2012 IEEE PES*, 2012.

[362] B. Stock, M. Engelberth, F. C. Freiling, and T. Holz. Walowdac – Analysis of a Peer-to-Peer Botnet. In *Proceedings of the European Conference on Computer Network Defense*, 2009.

[363] B. Stone-Gross, M. Cova, L. Cavallaro, B. Gilbert, M. Szydlowski, R. Kemmerer, C. Kruegel, and G. Vigna. Your Botnet is My Botnet: Analysis of a Botnet Takeover. In *Proceedings of the 16th ACM Conference on Computer and Communications Security*, 2009.

[364] G. Stringhini, C. Kruegel, and G. Vigna. Detecting Spammers on Social Networks. In *Proceedings of the 26th Annual Computer Security Applications Conference*, 2010.

[365] Symantec. Spam report: Hacked personal email accounts used to scam contacts. `http://www.symantec.com/articles/article.jsp?aid=20080729_spam_report`.

[366] Symantec. Stuxnet Using Three Additional Zero-Day Vulnerabilities. `http://www.symantec.com/connect/blogs/stuxnet-using-three-additional-zero-day-vulnerabilities`, September 2010.

[367] Symantec Official Blog. Top 5 security predictions for 2013 from symantec. Internet. `http://www.symantec.com/connect/blogs/top-5-security-predictions-2013-symantec-0`, 2012.

[368] P. Ször. *The Art of Computer Virus Research and Defense*. Addison-Wesley Professional, February 2005.

[369] S. Tang, H. Mai, and S. King. Trust and Protection in the Illinois Browser Operating System. In *Proceedings of the 10th USENIX conference on Operating Systems Design and Implementation (OSDI)*. USENIX, 2010.

[370] T. Telegraph. Bogus' ap tweet about explosion at the white house wipes billions off us markets. http://www.telegraph.co.uk/finance/markets/10013768/Bogus-AP-tweet-about-explosion-at-the-White-House-wipes-billions-off-US-markets.html.

[371] M. Ter Louw and V. Venkatakrishnan. Blueprint: Precise Browser-neutral Prevention of Cross-site Scripting Attacks. In *Proceedings of the 30th IEEE Symposium on Security & Privacy*, Oakland, CA, May 2009.

[372] The Honeynet Project. Droidbox. https://code.google.com/p/droidbox/.

[373] The SysSec Consortium. Deliverable D7.1: Review of the state-of-the-art in cyberattacks, June 2011.

[374] K. Theocharoulis, I. Papaefstathiou, and C. Manifavas. Implementing rainbow tables in high-end fpgas for super-fast password cracking. In *Proceedings of the 2010 International Conference on Field Programmable Logic and Applications*, pages 145–150, 2010.

[375] K. Thomas, C. Grier, and V. Paxson. Adapting social spam infrastructure for political censorship. In *Proceedings of the USENIX Workshop on Large-Scale Exploits and Emergent Threats (LEET)*, 2012.

[376] O. Thonnard, L. Bilge, G. O'Gorman, S. Kiernan, and M. Lee. Industrial espionage and targeted attacks: understanding the characteristics of an escalating threat. In *RAID'12: Proceedings of the 15th international conference on Research in Attacks, Intrusions, and Defenses*. Springer-Verlag, Sept. 2012.

[377] Tim Rains - Microsoft. Using the past to predict the future: Top 5 threat predictions for 2013. Internet. http://blogs.technet.com/b/security/archive/2012/12/13/using-the-past-to-predict-the-future-top-5-threat-predictions-for-2013.aspx?Redirected=true, 2012.

[378] V. Toubiana, A. Narayanan, D. Boneh, H. Nissenbaum, and S. Barocas. Adnostic: Privacy preserving targeted advertising. In *Proceedings of the 17th Network and Distributed System Security Symposium*, 2010.

[379] N. Tran, B. Min, J. Li, and L. Subramanian. Sybil-resilient online content voting. In *Proceedings of the 6th USENIX symposium on Networked systems design and implementation*, pages 15–28, 2009.

[380] C.-Y. Tseng and M.-S. Chen. Incremental SVM Model for Spam Detection on Dynamic Email Social Networks. *2009 International Conference on Computational Science and Engineering*, pages 128–135, 2009.

[381] K.-Y. Tseng, D. Chen, Z. Kalbarczyk, and R. K. Iyer. Characterization of the error resiliency of power grid substation devices. In *International Conference on Dependable Systems and Networks*. IEEE Computer Society, June 2012.

[382] V. van der Veen, N. dutt Sharma, L. Cavallaro, and H. Bos. Memory Errors: The Past, the Present, and the Future. In *In Proceedings of the 15th International Symposium on Research in Attacks Intrusions and Defenses (RAID)*, September 2012.

[383] R. Vigo. The Cyber-Physical Attacker. In *dl.acm.org*, pages 347–356. Springer Berlin Heidelberg, Berlin, Heidelberg, 2012.

[384] B. Viswanath, M. Mondal, K. P. Gummadi, A. Mislove, and A. Post. Canal: scaling social network-based Sybil tolerance schemes. In *Proceedings of the 7th ACM european conference on Computer Systems - EuroSys '12*, page 309. ACM Press, 2012.

[385] B. Viswanath, A. Post, K. P. Gummadi, and A. Mislove. An Analysis of Social Network-Based Sybil Defenses. In *Proceedings of the ACM SIGCOMM 2010 conference*, page 363, New York, New York, USA, 2010. ACM Press.

[386] P. Vreugdenhil. Pwn2Own 2010 Windows 7 Internet Explorer 8 exploit. http://vreugdenhilresearch.nl/Pwn2Own-2010-Windows7-InternetExplorer8.pdf.

[387] D. Wagner, J. S. Foster, E. A. Brewer, and A. Aiken. A first step towards automated detection of buffer overrun vulnerabilities. In *Proceedings of the Network and Distributed Systems Security Symposium (NDSS)*, 2000.

[388] D. Wang, L. Kaplan, H. Le, and T. Abdelzaher. On truth discovery in social sensing. In *Proceedings of the 11th international conference on Information Processing in Sensor Networks - IPSN '12*, page 233. ACM Press, 2012.

[389] H. J. Wang, X. Fan, J. Howell, and C. Jackson. Protection and Communication Abstractions for Web Browsers in MashupOS. In T. C. Bressoud and M. F. Kaashoek, editors, *SOSP*, pages 1–16. ACM, 2007.

[390] H. J. Wang, C. Grier, A. Moshchuk, S. T. King, P. Choudhury, and H. Venter. The Multi-Principal OS Construction of the Gazelle Web Browser. In *Proceedings of the 18th USENIX Security Symposium*, Montreal, Canada, August 2009.

[391] R. Wang, S. Chen, and X. Wang. Signing me onto your accounts through facebook and google: A traffic-guided security study of commercially deployed single-sign-on web services. In *Proceedings of the 2012 IEEE Symposium on Security and Privacy*, SP '12, pages 365–379, Washington, DC, USA, 2012. IEEE Computer Society.

[392] W. Wang and Z. Lu. Cyber security in the Smart Grid: Survey and challenges. *Computer Networks*, Jan. 2013.

[393] Websense. 2013 threat report. Internet. `http://www.websense.com/content/websense-2013-threat-report.aspx?intcmp=hp-promo-pod-en-2013-threat-report-preorder`, 2012.

[394] WEF. Global Risk 2012 Report. Seventh Edition. Technical report, World Economic Forum - WEF, 2012.

[395] D. Wei, Y. Lu, M. Jafari, P. M. Skare, and K. S. G. I. T. o. Rohde. Protecting Smart Grid Automation Systems Against Cyberattacks. *Smart Grid, IEEE Transactions on*, 2(4), 2011.

[396] M. Weir, S. Aggarwal, M. Collins, and H. Stern. Testing metrics for password creation policies by attacking large sets of revealed passwords. In *Proceedings of the 17th ACM conference on Computer and communications security*, CCS '10, pages 162–175, New York, NY, USA, 2010. ACM.

[397] T. Werner. Botnet Shutdown Success Story: How Kaspersky Lab Disabled the Hlux/Kelihos Botnet, 2011. Technical Report: `http://www.securelist.com/en/blog/208193137/`.

[398] K. Wilhoit. Who's Really Attacking Your ICS Equipment? pages 1–18, Mar. 2013.

[399] H. Wimberly and L. M. Liebrock. Using fingerprint authentication to reduce system security: An empirical study. In *Proceedings of the 2011 IEEE Symposium on Security and Privacy*, pages 32–46, 2011.

[400] J. Wyke. ZeroAccess, 2012. Technical Report by SophosLabs.

[401] L. Xu, F. Sun, and Z. Su. Constructing precise control flow graphs from binaries. Technical report, Department of Computer Science, UC Davis, 2009.

[402] R. Xu, H. Saıdi, and R. Anderson. Aurasium: Practical policy enforcement for android applications. In *Proc. of USENIX Security*, 2012.

[403] G. Yan, S. Chen, and S. Eidenbenz. RatBot: Anti-enumeration Peer-to-Peer Botnets. In *Lecture Notes in Computer Science, vol. 7001*, 2011.

[404] G. Yan, D. T. Ha, and S. Eidenbenz. AntBot: Anti-Pollution Peer-to-Peer Botnets. In *Journal of Computer Networks, vol. 55*, 2011.

[405] L.-K. Yan and H. Yin. DroidScope: Seamlessly Reconstructing OS and Dalvik Semantic Views for Dynamic Android Malware Analysis. In *Proc. of USENIX Security*, 2012.

[406] C. Yang, R. Harkreader, J. Zhang, S. Shin, and G. Gu. Analyzing spammers' social networks for fun and profit. In *Proceedings of the 21st international conference on World Wide Web - WWW '12*, page 71. ACM Press, 2012.

[407] W. Yang, N. Li, Y. Qi, W. Qardaji, S. McLaughlin, and P. McDaniel. Minimizing private data disclosures in the smart grid. In *ACM conference on Computer and Communications Security*. ACM Request Permissions, Oct. 2012.

[408] Z. Yang, C. Wilson, X. Wang, T. Gao, B. Y. Zhao, and Y. Dai. Uncovering social network sybils in the wild. *Proceedings of the 2011 ACM SIGCOMM conference on Internet measurement conference - IMC '11*, page 259, 2011.

[409] T.-F. Yen and M. K. Reiter. Revisiting Botnet Models and Their Implications for Takedown Strategies. In *Proceedings of the 1st Conference on Principles of Security and Trust*, 2012.

[410] H. Yu, P. B. Gibbons, M. Kaminsky, and F. Xiao. SybilLimit: A Near-Optimal Social Network Defense against Sybil Attacks. In *2008 IEEE Symposium on Security and Privacy (sp 2008)*, number Figure 1, pages 3–17. IEEE, May 2008.

[411] H. Yu, M. Kaminsky, P. B. Gibbons, and A. Flaxman. SybilGuard: Defending Against Sybil Attacks via Social Networks. In *Proceedings of the 2006 conference on Applications, technologies, architectures, and protocols for computer communications - SIGCOMM '06*, number September, pages 267–278, New York, New York, USA, 2006. ACM Press.

[412] S. Zanero and Z. Simic. Emergent phenomena testbed simulator for improving SCADA performance in power system security management. Technical report, 2013.

[413] C. Zheng, S. Zhu, S. Dai, G. Gu, X. Gong, X. Han, and W. Zou. SmartDroid: an automatic system for revealing UI-based trigger conditions in Android applications. In *Proc. of SPSM*, 2012.

[414] W. Zhou, Y. Zhou, X. Jiang, and P. Ning. Detecting repackaged smartphone applications in third-party android marketplaces. In *Proceedings of the second ACM conference on Data and Application Security and Privacy (CODASPY)*, 2012.

[415] Y. Zhou and X. Jiang. Android Malware Genome Project. `http://www.malgenomeproject.org/`.

[416] Y. Zhou and X. Jiang. Dissecting android malware: Characterization and evolution. In *Proc. of the IEEE Symposium on Security and Privacy*, 2012.

[417] Y. Zhou, Z. Wang, W. Zhou, and X. Jiang. Hey, you, get off of my market: Detecting malicious apps in official and alternative Android markets. In *Proceedings of the 19th Annual Network & Distributed System Security Symposium*, Feb. 2012.

[418] T. Zimmermann, N. Nagappan, and L. Williams. Searching for a Needle in a Haystack: Predicting Security Vulnerabilities for Windows Vista. In *Proc. of the 3rd International Conference on Software Testing, Verification and Validation*, ICST'10, Apr. 2010.

[419] M. Zuckerberg. Facebook across the web. `http://www.facebook.com/blog/blog.php?post=41735647130`.

SEVENTH FRAMEWORK PROGRAMME

Information & Communication Technologies
Trustworthy ICT

Grant Agreement No. 257007

A European Network of Excellence in Managing Threats and Vulnerabilities in the Future Internet: Europe for the World

After the completion of its second year of operation, the SysSec Network of Excellence produced this "Red Book of Cybersecurity" to serve as a roadmap in the area of systems security. To realize this book, SysSec put together a Task Force of top-level young researchers in the area, steered by the advice of SysSec WorkPackage Leaders. The Task Force had vibrant consultations with the Working Groups of SysSec, the Associated members of SysSec, and the broader Systems Security Community. Capturing their feedback in an on-line questionnaire and in forward-looking "what if" questions, the Task Force distilled their knowledge, their concerns, and their vision for the future.

The result of this consultation has been captured in this book, which we hope will serve as a roadmap of systems security research, and as an advisory document for policy makers and researchers who would like to have an impact on the security of the future Internet.