Part I

Introduction

1 Introduction

Chapter overview

1. Network security
   - importance and relevance of network security
   - challenges of hardware, software, and networking complexity
   - multifaceted nature of network security
2. Approaches
   - why and how are decision and game theories relevant
   - interplay between theory and practice
   - detection, decision, and response
3. Motivating examples
   - security games
   - security risk-management
   - optimal malware epidemic response

Chapter summary

Network security is an important, challenging, and multi-dimensional research field. In its study, theory and practice should function together as parts of a feedback loop. Game, optimization, and control theories, among others, provide a mathematical foundation to formalize the multitude of decision-making processes in network security. The analytical approaches and quantitative frameworks presented lead to better allocation of limited resources and result in more informed responses to security problems in complex networked systems and organizations.

1.1 Network security

Networked computing and communication systems are of **vital importance** to the modern society simply because the civilization we know today would cease to exist without them. A good illustration of this fact is provided by the Internet, the epitome of networks that has evolved to a global virtual environment and become an indispensable part of our lives. Nowadays our communication, commerce, and entertainment are all based on networked systems in one way or another. Once they are disrupted its cost to society is hard to measure, but enormous, for sure. As an example, the Code Red worm, which infected about 360,000 servers in 2001, has cost – according to estimates – hundreds of millions of dollars globally in lost productivity and clean-up of systems afterwards [115].

The security of computers and networks has become an increasingly important concern as they have grown out of research environments where they fulfilled only specific duties at the hands of well-trained specialists. Security problems emerged once such systems entered general public life and started to be used for a multitude of different purposes in business, entertainment, and communication. Today, an overwhelming majority of users are no longer trained professionals who would know the nature and limitations of these systems. To complicate matters further, there is no shortage of malicious individuals and groups to exploit weaknesses of networked systems and their users for financial gain or other purposes.

Despite its vital importance and the ongoing research efforts, network security remains an **open problem**. This is partly because networked systems are difficult to *observe* and *control* even by their legitimate users and owners due to their *complexity* and *interconnected nature*. A regular user has only limited observational capabilities, for the user interface is only the tip of the iceberg. A significant number of automated system processes run hidden in the background, since it is simply infeasible to expose users to all of them. Furthermore, many of these processes involve communication with multiple other computing systems across networks, creating tightly coupled supersystems. As a simple example, the system clock of a computer is usually controlled by a specific program that runs in the background and corrects it by connecting to a time server possibly on the other side of the world, all of which is unknown to most regular users.

The first factor that is responsible for loss of observability and control of networked systems is **complexity**. Complexity is due to both hardware and software. The complexity of system and application software has increased significantly as a result of enormous advances in hardware in the last few decades. The microprocessor revolution is probably the best example of these advances (Figure 1.1). Thanks to the progress in microprocessors, the personal computers today are as powerful as the supercomputers of two decades ago. As a natural consequence, the software running on computers has become more layered and complex. The widely used Microsoft Windows and Linux operating systems on personal computers consist of tens of millions of lines of code (Figure 1.2).

A related issue contributing to the complexity is the unintentional flaws (bugs) in software. In current software architectures, there is always a mismatch between the

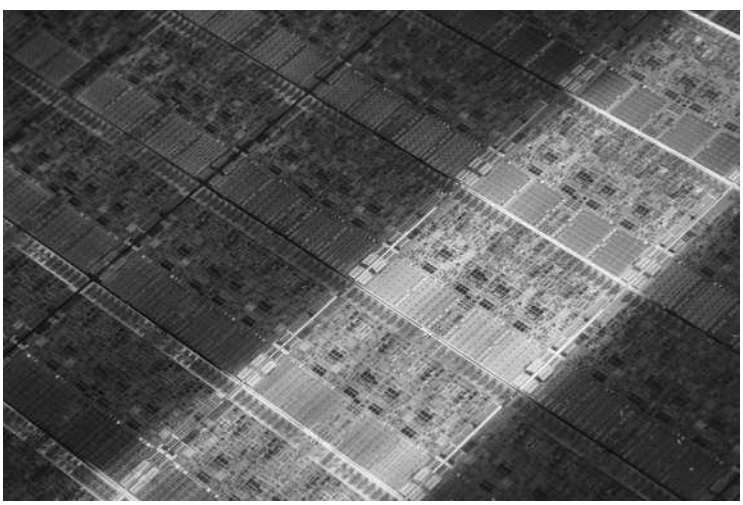

Figure 1.1 Intel processors with Nehalem architecture have multicores and between 500 and 1000 million transistors. Many of them are sold for home and small office usage. (Image courtesy of Intel Corporation.)

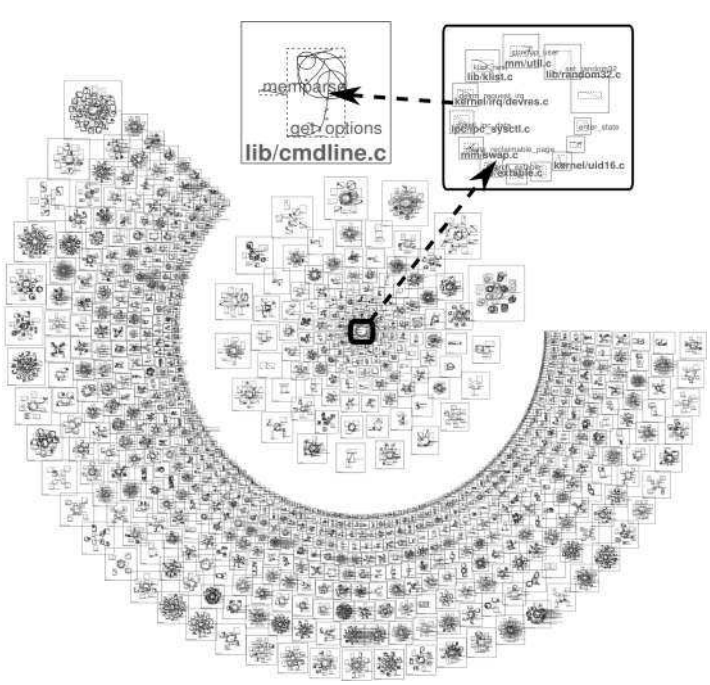

Figure 1.2 Linux 2.6 operating system kernel, visually depicted here, is widely used on a variety of devices and has more than 10 million source lines of code.

intentions of the developer and the actual behavior of the program, which exhibits itself as **software bugs**. Unfortunately, this is another permanently open problem due to the fundamental miscommunication between humans, who are by nature imprecise, and computers, which are based on rigid mathematical principles.

The second factor that contributes toward making networked systems difficult to observe and control is their **interconnected** nature. The distributed architecture of contemporary networks along with the complexity of the underlying computing and communication environments prevent systems administrators and organizations from having absolute control over their networks. Network boundaries are often vague and administrators cannot exercise control outside their local domain, which leaves networked systems vulnerable to distant security attacks due to global connectivity. This issue is often half-jokingly captured by the phrase "the most secure computer is the one unplugged from the network" (Figure 1.3).

The difficulty of observing and exercising control on networked systems can be illustrated by the following everyday **example**. Consider a laptop computer connected to the Internet. Although the connection makes the laptop a part of the most complex systems ever built, it still shares a simple property of the simplest man-made tools: it can be used for "good" purposes as well as exploited for "bad" ones. A malicious *attacker* can potentially run a malicious program (*malware*) on this laptop without permission of its owner by exploiting various *vulnerabilities*. While the laptop is physically next to its rightful owner, it can thus be partially controlled by the attacker who may be on the other side of the world. Furthermore, the owner may not even know or *observe* the security problem, allowing the attacker to maintain partial control over an extended time period.

In addition to limited observability and control of the underlying complex systems, another defining aspect of network security is its social dimension or so-called human factors. Network security is not a problem that can be solved once and for all by engineering a solution. It should be seen as a **problem** that needs **to be managed**, similarly to the security of a city. In other words, there is no such thing as a fully secured network, just as there is no such thing as a city without crime. Networks are man-made systems.

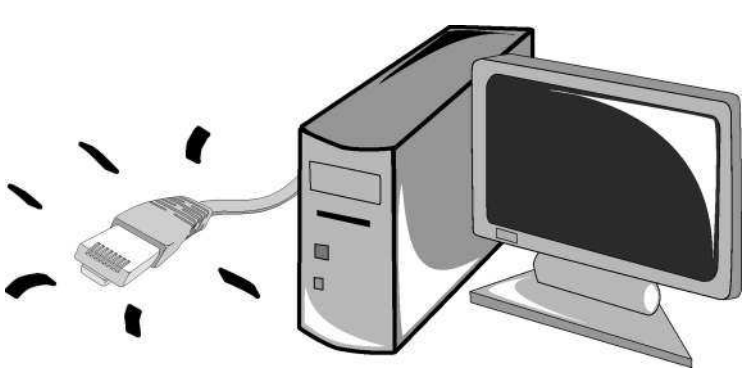

Figure 1.3 "*The most secure computer is the one unplugged from the network.*" The US Department of Defense C2 rating of Windows NT 3.5 only applied to a computer unplugged from the network!

Since the system engineers are human beings who make mistakes, future networks will have vulnerabilities in some form, no matter how carefully they are designed. As long as there are people who would benefit from exploiting vulnerabilities for selfish reasons, there will always be security threats and attacks.

Given the **dynamic nature** of network security, attackers cannot be stopped by purely static measures such as classical firewalls. When targeting the vulnerabilities of networks, attackers update their strategies from day to day. Hence, it is crucial for the defense side to also take dynamic measures and address security both in the design phase and afterwards. Many existing defense mechanisms already adopt such a dynamic approach and offer various *security services*. Automatic (patching) updates of antivirus programs, browsers, and operating systems are well-known examples. In a sense, network security is like a "game" played between attackers and defenders on the board of complex networks using attacks and defensive measures as pieces (Figure 1.4).

The complexity, multi-dimensionality, and importance of network security makes it hard to describe with a single definition. The **multifaceted nature** of network security puts one in mind of an ancient story about *blind men and an elephant*. According to the tale, a group of blind men, who have never heard of elephants, want to learn about them by observing an elephant that comes to their village for the first time. Once they approach the elephant, each blind man touches a different part of the animal and tells others his opinion: the first man touches the body and says "elephant is like a wall," the second one touching the tusk says "elephant is like a spear," the third one holds the trunk and says "elephant is like a snake," and so on ...

The field of network security can be compared to the elephant in the story (Figure 1.5) and security researchers to the blind men trying to understand it. For researchers in the field of cryptography, security is all about cryptographic algorithms and hash functions.

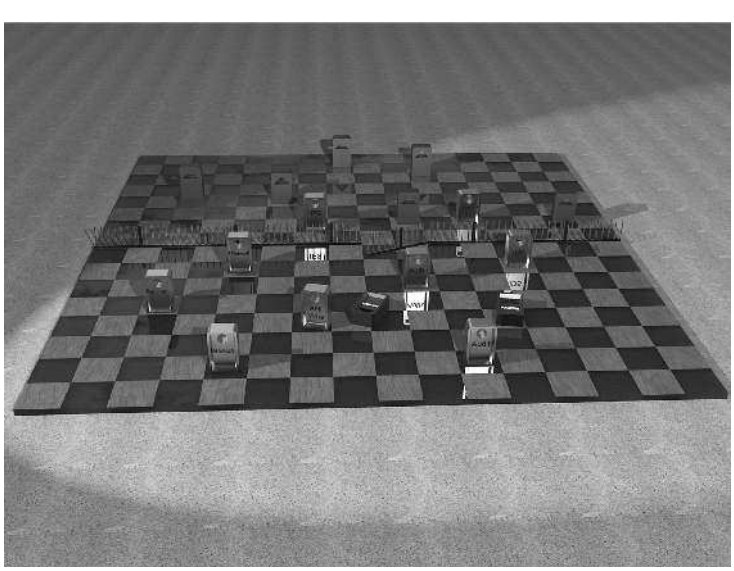

Figure 1.4 Network security is like a "game" played between attackers and defenders.

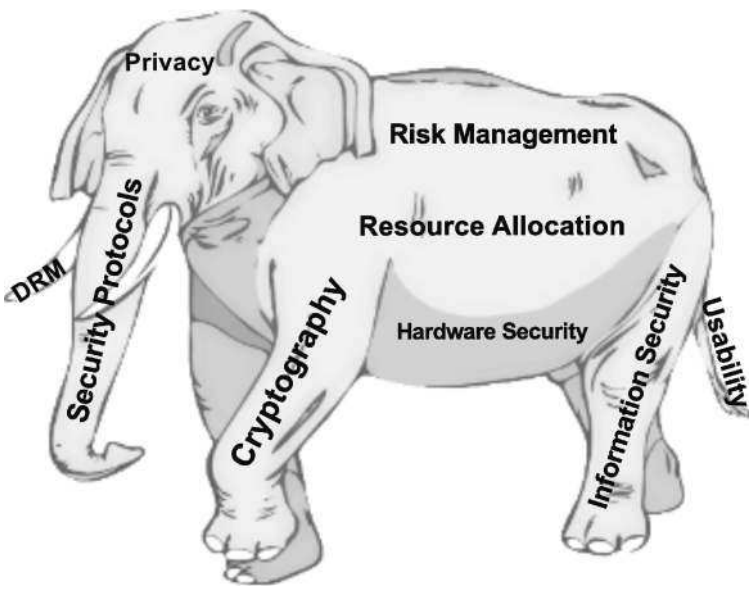

Figure 1.5 Elephant as metaphor for network security, from the ancient tale of *blind men and an elephant*.

Those who are in information security focus mainly on privacy, watermarking, and digital rights management systems. For researchers with an interest in hardware, security is about tamper-resistant architectures and trusted computing. Network security encompasses all these aspects and more. Researchers, unlike the blind men, are of course aware of this fact regardless of their specific focus. However, a wide field such as this calls for specialization and different perspectives.

This book also adopts a specific view of network security and introduces a **decision and game-theoretic approach**. The upcoming chapters study incentive mechanisms, analytical models, and resource allocation aspects of network security in terms of detection, decision making, and response. But first, the next section presents the adopted decision and game-theoretic approach, which is subsequently illustrated with examples in Section 1.3.

1.2 The approach

There is a fundamental relationship between **security and decision making**. Whether it is about buying a simple lock versus installing an expensive alarm system in a house, deploying a security suite on a personal computer, or applying a patch to a production server, decisions on allocating limited resources while balancing risks are at the center of network security. Making such decisions in a principled way instead of relying on heuristics provides numerous advantages and simply corresponds to following the celebrated *scientific method*.

It is, therefore, not surprising to observe **theoretical models** at the system level play an increasing role in network security as the field matures from its earlier qualitative and empirical nature. The increasing number of books, journal articles, and conference publications that study the problem analytically is clear evidence of the emerging interest in this approach to network security. The mathematical abstraction provided by quantitative models is useful for generalization of problems, combining the existing ad-hoc schemes under a single umbrella, and opening doors to novel solutions. Hence, the analytical approach provides a unique advantage over heuristic schemes that are problem specific. One of the main objectives of this book is to develop a deeper understanding of existing and future network security problems from a decision and game-theoretic perspective.

Securing information, controlling access to systems, developing protocols, discovering vulnerabilities, and detecting attacks are among the well-known topics of network security. In practice, all these involve decision making at multiple levels. **Security decisions** allocate limited resources, balance perceived risks, and are influenced by the underlying incentive mechanisms. Although they play an important role in everyday security, they are often overlooked in security research and are usually made in a very heuristic manner.

The human brain is undoubtedly a wonderful decision-making engine which current technology has not been able even remotely to replicate. A security expert can make very balanced decisions taking into account multiple factors and anticipating future developments. Admittedly, none of the techniques discussed in this book can come close to the performance of such an expert who relies on "intuition" which combines enormous pattern recognition capabilities with years of experience.

However, relying on **human expertise** in this manner has its own set of **shortcomings**. The first one is scale. Given the complexity and the number of networked systems around us, a security expert cannot oversee all systems all the time. A second issue is the availability of good experts. The number of experts is very limited due to the long training period required. In many cases, an organization has to work with available people of limited knowledge in less than ideal circumstances. A third problem is the timescale. Computers operate on a much faster timescale than humans and some security problems (e.g. malware epidemics) require an immediate response of the order of seconds. The human brain, despite its wonderful properties, operates on a much slower timescale.

It is hence unavoidable to have to rely on **computer assistance** in some form to address network security problems. Consider, for example, an organization that employs multiple experts to secure its networked systems. Given the scale and availability limitations, the organization has to naturally direct the attention of its experts primarily to the most important systems. This strategic resource allocation decision is often made by the management again relying on human expertise, but manifests itself electronically through scheduling systems or spreadsheets. In addition, the organization has to equip its security experts with a variety of computational tools to effectively observe and respond to security problems in a timely manner. Thus, computer assistance is essential in network security regardless of the degree of reliance on human expertise.

Both computerized support systems and security managers *implicitly* make numerous strategic decisions in terms of resource allocation, detection, or response. Decisions such as a log file viewer not showing some fields due to limited screen estate or a manager ordering a security administrator to patch a certain server have nontrivial consequences for overall security. In most of the existing security structures, such decisions are made in a heuristic manner. Therefore, the issues with human experts discussed above directly apply.

An alternative to implicit and heuristic decision making is the **analytical approach** based on mathematical models. For example, a manager can pose the problem of allocation of limited security experts within the organization as one of optimization where the available resources (e.g. in terms of man-hours) and degree of importance of specific subsystems (e.g. in terms of monetary loss when attacked or down) are quantified explicitly. Then, the problem can be solved automatically with computer assistance on a large scale. Another example is a packet-filtering system to decide on whether or not to drop a packet, based on a preset threshold. Unlike the previous case, this decision is made within milliseconds and repeated millions of times. How to (dynamically) determine the threshold value used as decision criterion can be investigated analytically and solved within an optimal and robust control framework based on given preferences.

The **quantitative approach** described has multiple **advantages** over the ad-hoc ones. First, the knowledge of the decision maker is expressed through mathematical models in a transparent and durable manner. Second, the decision making can now be made on a large scale. Third, it can be made as fast as numerical solution methods allow, in many cases of the order of milliseconds. While some security decisions, such as ones on investments or policies, are made over days or months, there are many security decisions made on much smaller timescales, as in the packet-filtering example mentioned. Finally, the decision-making process captured by the model can now be checked experimentally and improved upon, providing a way of aggregating the knowledge of multiple experts.

Developing a sound decision and game-theoretic framework for security requires building a **feedback loop** between high-quality theoretical research and real-life problems experienced by practitioners on a daily basis, as depicted in Figure 1.6. Actual problems lead to deep questions for fundamental research whereas theoretical advances provide novel algorithms and solutions. It has been repeatedly observed in many fields of science that weakening this feedback loop is detrimental to both theoretical and practical efforts. Therefore, the intention here is to bridge the gap between theory and practice as much as possible through simulations and proof-of-concept demonstrations addressing existing and emerging network security topics.

A variety of well-established **mathematical theories** and tools can be utilized to model, analyze, and address network security problems. Adopting a defensive approach, many resource allocation aspects of protecting a networked system against malicious attacks can be formulated as optimization problems. Especially convex **optimization** problems are well understood and a plethora of tools exists to solve them efficiently. When the underlying system dynamics play a significant role in security, **control theory** provides a large field of expertise for extending static optimization formulations to control of dynamic systems. Fundamental concepts such as observability and controllability

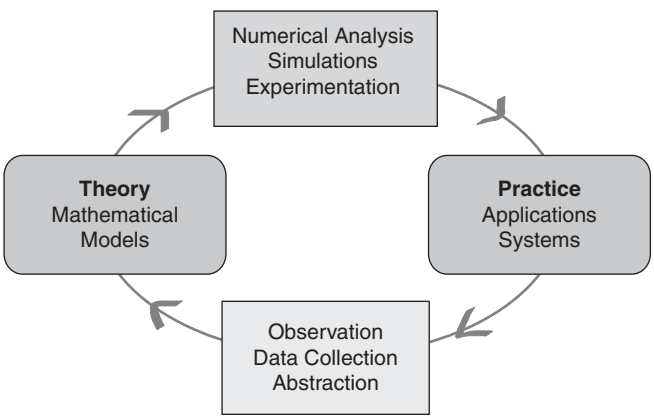

Figure 1.6 Feedback loop between theory and practice is of fundamental importance for the success of both.

relevant to security have been mathematically formalized and utilized in control theory for decades.

Beyond single-person decision making, **game theory** provides a rich set of mathematical tools and models for investigating multi-person strategic decision-making where the players (decision makers) compete for limited and shared resources. As a special case, security games study the interaction between malicious attackers and defenders. Security games and their solutions are used as a basis for formal decision-making and algorithm development as well as to predict attacker behavior. Security games vary from simple deterministic ones to more complex stochastic and limited information formulations and are applicable to security problems in a variety of areas ranging from intrusion detection to social, wireless, and vehicular networks.

Despite the broad scope and extent of available mathematical models in optimization, control, and game theories, one should not consider these fields as static. There are many open problems in each of them and they are themselves progressing. For example, how to incorporate information structures and formalize decision making under information limitations in single- and multiple-person dynamic problems are active research areas of great relevance to security applications. The models in this book should be interpreted merely as first steps toward developing realistic frameworks for decision making in security. Hence, establishing mature and relevant models is one of the important research challenges in the decision-theoretic approach to security.

Organization of the book

Based on the presented approach, the remainder of the book is organized into three parts encompassing detection, analysis, and optimized response as illustrated in Figure 1.7. Decision and game-theoretic schemes utilizing optimization, control, and machine learning are applied to a variety of network security topics. Principles for scalable, robust, effective security architectures for autonomous operation as well as computer