Dagstuhl Seminar 06371 "From Security to Dependability"

Christian Cachin[1], Felix C. Freiling[2], and Jaap-Henk Hoepman[3]

[1] IBM Zurich Research Laboratory, Switzerland
[2] University of Mannheim, Germany
[3] Radboud University Nijmegen, The Netherlands

Abstract. This seminar brought together researchers and practitioners from the different areas of dependability and security, in particular, from fault-tolerance, safety, distributed computing, langelanguage-based security, and cryptography. The aim was to discuss common problems faced by research in these areas, the differences in their respective approaches, and to identify research challenges in this context.

Keywords. fault-tolerance, safety, distributed computing, language-based security, cryptography

1 Motivation

Security remains an elusive property for many systems today. Despite the research efforts of the last decades, the tremendous progress made, for example in the area of cryptography, and the impressive security technology being deployed with modern operating systems, security problems have not gone away. One reason why security technology may not have been able to fulfill its promise may be a lack of integration with the existing systems, and in particular with the technologies for fault tolerance.

Although fault tolerance and security are both necessary attributes of dependable systems, these properties have traditionally been treated separately and lead to distinct and orthogonal research areas. Both research areas are based on formal models, but their separation has lead to different approaches on achieving and validating the respective properties, and the approaches have become the subject of different communities.

As one particular example, consider the area of fault-tolerant systems on the one hand and secure systems (in particular those using cryptography) on the other: Researchers in fault-tolerance often make statements about systems by treating cryptographic primitives as black boxes. This is done to keep the model tractable, i.e., to simplify analysis and (sometimes) avoid number and probability theory. In the area of safety-critical systems, such models have been successfully applied in practice, with support from automated analysis and verification tools. However, abstracting away the basic properties of cryptographic primitives severely constrains the ability to conduct rigorous security proofs. Various

Dagstuhl Seminar Proceedings 06371
From Security to Dependability
http://drops.dagstuhl.de/opus/volltexte/2007/851

examples of the past have shown how important attributes were neglected due to over-abstraction, hence contributing to weaknesses in the resultant protocols.

The separate areas are only recently being viewed as complementary, with work underway to unify the two approaches. We mention the current work on tool-supported formal verification of cryptographic protocols and the concept of intrusion-tolerant systems, i.e., systems that continue to provide their service despite the corruption or failure of some of their parts.

As indicated by the above and confirmed by many researchers, there are strong similarities between the ways of modeling and handling uncertainty in the different areas of dependable systems. But there also seem to be fundamental tradeoffs that lead different communities into different directions.

2 Topics of the Seminar

The Dagstuhl seminar brought together researchers and practitioners from the different areas of dependability (in particular, from fault-tolerance, safety, security, and cryptography) in order to discuss the foundations of these areas, their similarities and differences.

Some of the research questions discussed during the seminar included:

- What are the relations between safety, fault-tolerance, security, and cryptography with respect to methodologies and models?
- What classifications and metrics for dependability and security properties exist and how can they be compared?
- What are the differences between methods to specify, model and analyse fault-tolerant and secure systems?
- Under which circumstances can fault-tolerance techniques be used to achieve security and security methods be used to achieve fault-tolerance?
- What is the role of cryptography in the development of protocols that are both secure and fault-tolerant?

3 Participation

The seminar was attended by about 50 researchers from industry and academia, with backgrounds ranging from safety and dependability to cryptography and formal verification of security protocols. The participants were a truly international group. Most of them were working in Europe at the time and some in North America, although their countries of origin were distributed over a much wider area on the globe.

4 Seminar Organization

We organized the seminar as a sequence of talks in which we mixed the contributions of the different communities as much as possible. Talks were restricted

to 30 minutes to allow at least 15 minutes of discussion after every presentation. The morning featured four talks, the afternoon usually 5, giving a total of 35 presentation during the week. Following Dagstuhl traditions, we fixed the programme from day to day. This gave us the flexibility to react to participants' requests and also to organize an "open air" discussion session on Wednesday afternoon. This session took place in the garden behind the castle and featured several 5-minute presentations (with no slides, only a flip chart was available) for which participants could sign up during the morning. Results of this session are reported below in the summary of findings. The week was completed by an excursion to the village of Riol on the Mosel river featuring two equally extensive events: an 11.1 km hike along the river and through the vineyards and tasting of excellent Riesling wines. The wheather throughout the week was extremely pleasant, contributing to the success of the seminar.

5 Summary of Findings

The following points summarize the main findings of the seminar from the point of view of the organizers.

Terminology

Despite continued efforts to unify and harmonize terminology, termonology is probably the most frequent source of misunderstanding between the communities involved in this seminar. This in particular refers to very basic terms like "security" or "safety." Interestingly, there seems to be the tendency that researchers estimate the other communities to be more unified and disciplined in their research methods and usage of speach.

Byzantine Failures and Failure Independence

A large part of this seminar was spent discussing distributed protocols that tolerated a certain number of arbitrary failures (commonly called Byzantine failures in the literature). Stemming from early work in the fault-tolerance field, arbitrary failures can also be regarded as malicious and targeted attacks. Therefore today, tolerance to Byzantine failures is now an accepted area of study in both fields of security and dependability. It was however argued that Byzantine failures should be applied in a security setting with care. The main point for this is the fact that it is hard to quantify the coverage of the classical threshold assumption that at most t out of n processes can behave arbitrarily. In the dependability area, where Byzantine faults were intended to model random hardware failures, failure independence can be justified and therefore assumption coverage can often be calculated. In the security area, recent work on concepts like adversary structures or core/survivor sets has lead to protocols which arguably can map their failure assumptions better to practical scenarios. Quantification of coverage, however, is still out of reach.

Requirements vs. Techniques

It seems that the existence of an intelligent and strategic adversary creates a significant difference between the areas of security and dependability. The coverage problem of Byzantine failured as discussed above is only one manifestation of this fact. Another example is that the requirements of secure and dependable applications sound very similar, the techniques used to implement the requirements are however very different. Integrity requirements for example are common to systems in both areas. But while systems tolerant to random hardware faults can use invertible functions like CRCs or Hamming codes for integrity checking, secure systems must revert to non-invertible hash functions and other cryptographic techniques.