06371 Abstracts Collection
From Security to Dependability
— Dagstuhl Seminar —

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

[1] IBM Research - Zürich, CH
cachin@acm.org
[2] Univ. Mannheim, DE
freiling@informatik.uni-mannheim.de
[3] Radboud Univ. of Nijmegen, NL
jhh@cs.ru.nl

Abstract. From 10.09.06 to 15.09.06, the Dagstuhl Seminar 06371 "From Security to Dependability" was held in the International Conference and Research Center (IBFI), Schloss Dagstuhl. During the seminar, several participants presented their current research, and ongoing work and open problems were discussed. Abstracts of the presentations given during the seminar as well as abstracts of seminar results and ideas are put together in this paper. The first section describes the seminar topics and goals in general. Links to extended abstracts or full papers are provided, if available.

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

06371 Executive Summary – From Security to Dependability

This seminar brought together researchers and practitioners from the different areas of dependability and security, in particular, from fault-tolerance, safety, distributed computing, language-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

Joint work of: Cachin, Christian; Freiling, Felix C.; Hoepman, Jaap-Henk

Full Paper: http://drops.dagstuhl.de/opus/volltexte/2007/851

Dagstuhl Seminar Proceedings 06371
From Security to Dependability
http://drops.dagstuhl.de/opus/volltexte/2007/853

Reasoning with MAD Distributed Systems

Lorenzo Alvisi (Univ. of Texas at Austin, USA)

Distributed services spanning multiple administrative domains (MADs) have become increasingly popular. In a MAD distributed service, nodes collaborate to provide some service that benefits each node, but there is no central authority that controls the nodes' actions. Examples of such services include Internet routing, wireless mesh routing, file distribution, archival storage, and cooperative backup.

Unfortunately, there currently exists no satisfactory way to model MAD services. In these systems, the classical dichotomy between correct and faulty nodes becomes inadequate. Nodes in MAD systems may depart from protocols for two distinct reasons. First, as in traditional systems, nodes may be *broken* and arbitrarily deviate from a protocol because of component failure, misconfiguration, security compromise, or malicious intent. Second, nodes may be *selfish* and alter the protocol in order to increase their utility. Byzantine Fault Tolerance (BFT) handles the first class of deviations well. However, the Byzantine model classifies all deviations as faults and requires a bound on the number of faults in the system; this bound is not tenable in MAD systems where *all* nodes may benefit from selfish behavior and be motivated to deviate from the protocol. Models that only account for rational behavior handle the second class of selfish deviations, but may be vulnerable to arbitrary disruptions if even a single node is broken and deviates from expected rational behavior.

As MAD distributed systems become more commonplace, it becomes increasingly important to develop a solid foundation for constructing this class of services. The challenge is (at least) threefold:

1. to develop a model for MAD services in which it is possible to prove that a given MAD service meets its goals, no matter what strategies nodes may concoct within the scope of the adversary model.
2. to demonstrate that MAD services developed under this model can be practical.
3. to understand how to simplify the development of MAD services under the new model.

This talk will describe some initial answers to these challenges.

Keywords: Byzantine faul-tolerance, game theory, peer-to-peer, Nash equilibrium, gossip, live content distribution

Designing Modular Services in the Scattered Byzantine Failure Model

Emmaunuelle Anceaume (CNRS Rennes, F)

In this talk, I present the scattered byzantine failure model. In this model processes alternate correct and faulty periods.

Specifically, during its faulty periods, a process behaves arbitrarily (one cannot expect anything from it during these periods) whereas during its correct periods, it behaves according to its specification. In that sense, the scattered Byzantine failure model generalizes the classical Byzantine failure model. We have characterized two reliable services guaranteeing timeliness properties in the presence of Byzantine failures, namely the Clock Synchronization and the Δ-Atomic Broadcast. We identify necessary and sufficient conditions to ensure the correctness of both services in the scattered byzantine failure model.

Keywords: Synchronous system, moving byzantine failures

A Holistic Approach to Event-Based Modeling and Testing of System Vulnerabilities

Fevzi Belli (Universität Paderborn, D)

Man-machine systems have several *desirable* global system properties such as user friendliness, reliability, safety, and security. System vulnerability is the lack, or the exposure to breaches, of any such property, potentially leading to an *undesirable* situation from the user's point of view. Such undesirable situations could arise from internal faults, unintended environmental failures or malicious attacks from the system environment. The undesirable system features, viewed here as the sum of situations, which are complementary to the desirable ones, must be taken into account in the system development process from the very beginning in assuring a stable system behavior and a robust operation. In this respect, this presentation proposes an event-based approach to modeling, analysis and testing of systems that exhibit various forms of vulnerabilities, in particular, those encountered in user interface design and safety critical systems. The emphasis of the work is on the *holistic* treatment of both desirable and undesirable system features in a similar manner at an identical level of abstraction.

The presentation introduces an elementary test terminology, based on finite-state automata theory and Petri nets, and demonstrates the applicability as well as the effectiveness of the approach using realistic examples drawn from different domains.

Keywords: Modeling, testing, fault tolerance

Verifiable Agreement: Limits of Non-Repudiation in Peer-to-Peer Ad Hoc Groups

Zinaida Benenson (Universität Mannheim, D)

We introduce verifiable agreement as a fundamental service for securing peer-to-peer ad hoc groups, and investigate its solvability.

Verifiability of a protocol result means that the participants can prove that the protocol reached a particular result to any third party (the verifier) which was not present in the network at the time of the protocol execution.

Verifiable agreement is a joint generalization of agreement and contract signing problems. Consider a network of n parties, some of which might be faulty (dishonest) and deviate arbitrarily from their programs. Agreement problems, such as Byzantine Generals or Interactive Consistency, require all honest parties to agree on a common output which depends on their initial inputs. Contract signing can be considered as verifiable agreement on a contract text.

We prove the necessary and sufficient conditions for solvability of verifiable agreement in both synchronous and asynchronous networks, and give protocols for all possible cases.

Keywords: Consensus, Byzantine Agreement, contract signing, impossibility

Joint work of: Benenson, Zinaida; Freiling, Felix C.; Pfitzmann, Birgit; Rohner, Christian; Waidner, Michael

See also: Zinaida Benenson, Felix C. Freiling, Birgit Pfitzmann, Christian Rohner, and Michael Waidner. Verifiable Agreement: Limits of Non-Repudiation in Mobile Peer-to-Peer Ad Hoc Networks. In Third European Workshop on Security and Privacy in Ad hoc and Sensor Networks, 2006. L. Buttyan, V. Gligor, and D. Westhoff (Eds.): ESAS 2006, LNCS 4357, pp. 165-178, 2006. Springer-Verlag Berlin Heidelberg 2006

Authenticated Query Flooding in Sensor Networks

Zinaida Benenson (Universität Mannheim, D)

We propose a novel mechanism for authentication of queries in a sensor network in case these queries are flooded. In our protocol, the base station appends an authenticator to every query, such that each sensor can verify with certain probability that the query is sent by the base station. Implicit cooperation between sensor nodes during the flooding process ensures that legitimate queries propagate quickly in the network, whereas the propagation of illegitimate queries is limited to only a small part of the network.

Keywords: Sensor network, authenticated broadcast, flooding, probabilistic security, symmetric cryptography

Joint work of: Benenson, Zinaida; Freiling, Felix; Hammerschmidt, Ernest; Lucks, Stefan; Pimenidis, Lexi

See also: Zinaida Benenson, Felix C. Freiling, Ernest Hammerschmidt, Stefan Lucks, Lexi Pimenidis. Authenticated Query Flooding in Sensor Networks. 21st IFIP International Information Security Conference SEC 2006, May 2006, Karlstad University, Karlstad, Sweden.

On the Space Requirements of Robust Storage Implementations

Gregory Chockler (IBM - Haifa, IL)

We study the space complexity of distributed storage algorithms that emulate a read/write shared memory abstraction over distributed base storage objects. We focus on the robust storage implementations, i.e., the implementations that tolerate contention, asynchrony, the crash failure of the writer, as well as arbitrary (Byzantine) failures of any number of readers and a threshold of base objects. Motivated by practical space complexity considerations, we introduce the notion of a constrained storage implementation which only bounds the amount of data being stored while allowing for potentially unbounded metadata (such as unbounded timestamps). The only restriction is that the metadata cannot be used to encode the data.

We prove the impossibility of devising a constrained robust storage algorithm with even one arbitrary failure of a base object. Our result sheds light on the inherent limitations on the robustness of practical distributed storage algorithms, and helps to explain how these limitations are circumvented in the existing systems. In addition, as a corollary, we get the first sharp separation between safe and regular storage.

Keywords: Robust storage, arbitrary (Byzantine) failures, fault-tolerant memory emulations, read/write registers

Joint work of: Chockler, Gregory; Guerraoui, Rachid; Keidar, Idit

Combining reliability and privacy in networks in general: a survey

Yvo Desmedt (Univ. College London, GB)

Current TCP/IP protocols (even including IPSEC) do not have the resilience to deal with routers taken over by the adversary. Evidently, if the adversary can control all nodes (routers) in the network, no solution exists. We distinguish between a point-to-point network and one that allows partial broadcast (such as ethernet). We distinguish between an adversary that can attack a threshold number of nodes, or the general case that is described by an adversary structure. The adversary can be passive (i.e. interested in violating the privacy only) or Byzantine (i.e. undermine the reliability only), or both. For the case of broadcast, the issue of protecting networks against moreover jamming is also discussed. The talk surveys the results achieved and the remaining open problems.

Keywords: Combinatorics, graph theory, reliability, privacy, secret sharing, multi-party computation

Application of Dependability Techniques to Security Problems

Yves Deswarte (LAAS - Toulouse, F)

Some techniques that are commonly used in developing dependable systems have been successfully applied to cope with security problems:

- intrusion tolerance consists in applying fault tolerance techniques to develop secure systems using insecure components (e.g., COTS), in insecure environments (e.g., the Internet), with untrustworthy users (including possible malicious privileged users such as operators, administrators and security officers).
- quantitative evaluation is commonly used to assess certain dependability characteristics (reliability, availability, safety, maintainability, etc.). This is much less common in security, but some attempts have been made to assess quantitatively the capability of a computing system to resist to malicious attacks.

This talk will present some approaches that can be contemplated to reach these goals, illustrated by examples of real implementations.

Keywords: Dependability, security, intrusion-tolerance, quantitative security evaluation

Latency-Efficient Byzantine Atomic Broadcast for WANs

Dan Dobre (TU Darmstadt, D)

With the emerging dependency on Web-Services and their growing exposure to security threats, service availability has become a major issue in the community. Replicating at geographically dispersed sites helps avoiding location-dependent catastrophic failures and attacks such as outages and DoS. To this end, Byzantine FT replication suited for the internet has been subject to recent research.

Ramasamy and Cachin recently proposed an asynchronous atomic broadcast protocol PABC, that is the first to have an amortized message complexity of $O(n)$. Although protocol PABC needs only $O(n)$ messages per payload, compared to previous best $O(n^2)$ solutions, it has a higher latency. That is inherently due to its centralized (many-to-one and one-to-many) communication pattern and its specific recovery mechanism. We show how to reduce the latency of PABC from 8 to 6 message rounds per payload.

PABC is structured in a parsimonious and a recovery mode. Under normal conditions, the protocol operates in the parsimonious mode, and under faulty or unstable conditions, it temporarily switches to the recovery mode.To improve the latency, we have modified the recovery part of PABC. Our solution avoids byzantine agreement on payloads, however it comes at the expense of larger messages during recovery. Since the algorithm operates mostly in the parsimonious mode, the penalty incurred has little or no impact on the overall efficiency.

Keywords: Atomic Broadcast, Consensus, Byzantine

Joint work of: Dobre, Dan; Ramasamy, HariGovind; Suri, Neeraj; Cachin, Christian

Efficient Secure Multiparty Computation with Smartcards

Lucia Draque Penso (Universität Mannheim, D)

We discuss how the use of smartcards enables efficient solutions for secure multiparty computation and related problems such as voting and fair exchange. We show how such problems can be reduced to consensus and then present a novel implementation, for any number of participants, which outperforms existing ones.

Keywords: Smartcards, security units, secure multiparty computation, voting, fair exchange, consensus, reduction

Joint work of: Draque Penso, Lucia; Fort, Milan; Freiling, Felix; Benenson, Zinaida; Kesdogan, Dogan

See also: M. Fort, F. Freiling, L. D. Penso, Z. Benenson, D. Kesdogan, TrustedPals: Secure Multiparty Computation Implemented with Smartcards, ESORICS 2006 (LNCS 4189, pages 34-48). | F. Freiling, M. Herlihy, L. D. Penso, Optimal Randomized Fair Exchange with Secret Shared Coins, OPODIS 2005 (LNCS volume to be published). | G. Avoine, F. Freiling, R. Guerraoui, M. Vukolic, Gracefully Degrading Fair Exchange with Security Modules, EDCC 2005 (LNCS 3463, pages 55-71).

Signatures in Fault-Tolerant Protocols

Klaus Echtle (Universität Duisburg-Essen, D)

Besides their wide use for security purposes digital signatures are also applied in various fault-tolerant protocols. They allow for an efficient detection of special malfunctions, in particular wrong replication and wrong forwarding of messages.

This talk presents an example of a typical "fault tolerance signature" as well as two special signature types: the relative signature and the UniSig method, which solve specific problems of some protocols. Potential counterparts of cryptographically strong signatures for the security field are discussed.

Keywords: Fault tolerance, protocol, digital signature, relative signature, UniSig

Full Paper:
http://dc.uni-due.de

From Dependability to Security

Christof Fetzer (TU Dresden, D)

I will talk about how software encoded processing and automatic patching can be used to increase the dependability and security of systems.

Keywords: Dependability, security, software encoded processing, automatic patching

From Security to Dependability - Déjà vu

Dieter Gollmann (TU Hamburg-Harburg, D)

In this talk, I will try to summarize an earlier attempt to provide a quantitative and probabilistic basis for (operational) security. In the ESPRIT research project PDCS, a formal framework was developed to model security as a function of attacker effort, and some experiments were conducted to provide empirical data on the viability of this research direction. There has been little follow-up work since, so one might discuss whether there are fundamental limitations to such an approach, or whether the security landscape has changed sufficiently to warrant a new attempt.

Keywords: Security, dependability, quantitative security analysis

Toward Integrating Safety and Security

Maritta Heisel (Universität Duisburg-Essen, D)

We discuss safety and security, as far as requirements, system and software architectures, and mechanisms are concerned. We identify similarities and differences between them. Finally, we point out how safety and security can jointly be taken into account in a system and software development process, where we take a pattern-based approach.

Keywords: Safety, security, software engineering

Joint work of: Heisel, Maritta; Hatebur, Denis

New Abstractions for a New World

Matti Hiltunen (AT&T Research - Florham Park, USA)

Security attacks have become a prominent dependability risk in our everyday lives.

The abstract fault models (fail-stop, ..., Byzantine) have provided the fault-tolerance community with a uniform foundation on which to develop solutions and algorithms for accidental faults. This talk considers security attacks from the point of view of these fault models. While the Byzantine model is often used to model a system subject to security attacks, this talk presents and justifies a stronger fault model, intrusion-stop, as a new foundation for building attack-tolerant Internet services. We will also outline new abstractions and new roles for the network (Internet) in providing support for services in this new world.

Malicious Mobile Code - A Look at Network Attacks

Thorsten Holz (Universität Mannheim, D)

In this talk, we present the concept of honeypots. A honeypot is an electronic decoy, i.e. a network resource deployed to be probed, attacked, and compromised. This methodology is used in the area of IT security to learn more about attack patterns and attacker behavior. A honeypot is usually a computer system with no conventional task in the network. This assumption aids in attack detection: every interaction with the system is suspicious and could point to a possibly malicious action.

We focus on malicious mobile code, especially in the form of bots. We introduce a possibility to automatically collect autonomous spreading malware and present preliminary results. In addition, we present several attack patterns and attack statistics collected during the last year.

Keywords: Malicious Mobile Code, Honeypots, Intrusion Detection

On the utility of informed replication

Flavio Junqueira (Univ. California - San Diego, USA)

Failures of processes in distributed systems are often not independent in practice. To capture dependencies of failures, we use a technique called informed replication, which consists in using attributes of processes to determine failure patterns. In this talk, we discuss two cases in which the utilization of informed replication leads to more efficient replication techniques in systems with a large number of participants. In the first case, we discuss quorum selection in systems formed by multiple sites spread across a wide-area network. By using location of processes to select quorums, we are able to obtain higher availability. In the second example, we leverage the diversity of software configurations in a large population of hosts. Our goal for such a system is to protect data against large-scale Internet attacks that exploit a software vulnerability to compromise hosts. By using informed replication, we are able to reduce significantly the storage overhead of replicated data. We conclude with a discussion on the utilization of informed replication to secure systems.

Dependability and Security Benchmarks for Computer Systems

Karama Kanoun (LAAS - Toulouse, F)

Benchmarking the dependability of a system consists of evaluating dependability or performance-related measures, experimentally or based on experimentation and modelling, in order to characterize objectively the system behaviour in the presence of faults. Such an evaluation should allow non-ambiguous comparison of alternative solutions. Non-ambiguity, confidence in results and meaningfulness are ensured by a set of properties a benchmark should satisfy. A benchmark can be performed with respect to accidental non-malicious faults (design faults, physical faults or operator non-malicious faults, causing system unavailability for example) or with respect to malicious attacks (causing denial of services, illegal access to confidential information or improper modification of the system).

The talk will present the basic concepts and components of dependability and security benchmarking, and examples of concrete benchmark implementations and execution results.

Keywords: Dependability / security evaluation, benchmarking

See also: K. Kanoun, Y. Crouzet, A. Kalakech, A.-E. Rugina and Ph. Rumeau. Benchmarking the Dependability of Windows and Linux using PostMark Workloads. in Proc. 16th IEEE International Symposium on Software Reliability Engineering (ISSRE 2005), pp. 11-20, Chicago, IL, USA, November 8, 2005

Denial of Service Protection with Beaver

Idit Keidar (Technion - Haifa, IL)

We present Beaver, a method and architecture to "build dams" to protect servers from Denial of Service (DoS) attacks. Beaver allows efficient filtering of DoS traffic using low-cost, high-performance, readily-available packet filtering mechanisms. Beaver improves on previous solutions by not requiring cryptographic processing of messages, allowing the use of efficient routing (avoiding overlays), and establishing keys and state as needed. We present two prototype implementations of Beaver, one as part of IPSec in a Linux kernel, and a second as an NDIS hook driver on a Windows machine.

Preliminary measurements illustrate that Beaver withstands severe DoS attacks without hampering the client-server communication. Moreover, Beaver is simple and easy to deploy.

Keywords: Denial of Service

Joint work of: Badishi, Gal; Keidar, Idit; Herzberg, Amir; Romanov, Oleg; Yachin, Avital

Full Paper: http://drops.dagstuhl.de/opus/volltexte/2007/849

Abstracting out Byzantine Behavior

Petr Kouznetsov (MPI für Software Systeme - Saarbrücken, D)

Many distributed systems are designed to tolerate the presence of *Byzantine* failures: an individual process may arbitrarily deviate from the algorithm assigned to it.

Depending on the application requirements, systems enjoy various levels of fault-tolerance. Systems based on state machine replication are able to *mask* failures so that their effect is not visible by the application. In contrast, cooperative peer-to-peer systems can tolerate bounded deviant behavior to some extent and therefore do not require masking, as long as each faulty node is *exposed* eventually. Finding an abstract way to reason about the levels of fault-tolerance is thus of immanent importance.

We discuss how the information of deviant behavior can be abstracted out in the form of a *Byzantine failure detector* (BFD). We formally define a BFD abstraction, and we discuss two ways of using the abstraction: (1) monitoring systems in order to retroactively detect Byzantine failures and (2) enforcing systems in order to boost their level of fault-tolerance. Interestingly, the BFD formalism allowed us to determine the relative hardness of implementing two popular abstractions in distributed computing: state machine replication and weak interactive consistency.

Keywords: Fault-tolerance, Byzantine failures, masking, detection, total order broadcast, weak interactive consistency

Joint work of: Druschel, Peter; Haeberlen, Andreas; Kouznetsov, Petr

Extended Abstract: http://drops.dagstuhl.de/opus/volltexte/2007/850

Dependability and security: positioning

Jean-Claude Laprie (LAAS - Toulouse, F)

The presentation is aimed at positioning dependability and security with respect to each other. This will be performed via their definitions, their attributes, the threats that can affect them. Regarding the latter, statistics will be given helping to appreciate the attention which they should be devoted.

The Failure of Cryptographic Hash Functions - Some Lessons to Learn

Stefan Lucks (Universität Mannheim, D)

Cryptographic hash functions are an important tool for many crypto-applications, such as digital signatures. Recently, we have seen a flood of attacks against such hash functions.

As it turned out, the hash function's security collapses badly, if an underlying building-block just slightly fails to meet its security requirement. We discuss new approaches for the "failure-friendly" design of such hash functions, providing a decent amount of security, even if an underlying building-block fails.

Keywords: Cryptography, hash function, provable security

Towards more practical security type systems

Heiko Mantel (RWTH Aachen, D)

After a brief introduction to the research area, I will discuss three directions for making type systems for informational flow security more practical for analyzing concurrent programs:

- Characterizations of information flow security are too conservative in the sense that they are violated by many programs although they are intuitively secure. As a solution, we proposed the combining calculus, which allows one to use analysis techniques for multiple characterizations in combination during the analysis of a program.
  (joint work with Henning Sudbrock and Tina Kraußer)
- Many common security mechanisms necessarily release some information about secrets, e.g., password-based authentication or encryption.
  This leads to exceptions in information flow policies and gives rise to the question how such exceptions should be controlled. A recent taxonomy demonstrates that the existing techniques for controlling declassification each focus on only one aspect of declassification (what is declassified, who controls the declassification, and where or when does the declassification occur). We demonstrate that, though the what? and where? dimensions are orthogonal, two prominent analysis techniques for these dimensions are not compatible. We developed a solution that addresses both dimensions.
  (joint work with Alexander Reinhard)
- When a program is rejected as possibly insecure by the security type system then it is usually left to the programmer to make the program secure. We present a unification-based transformation for making some programs automatically secure during the type checking. Among other advantage over previous proposals, the transformation is idempotent, which provides a basis for multi-pass transformations to enforce security policies with more than two security levels by a fix-point computation. A similar technique might be applicable for making programs secure and fault tolerant, given an idempotent transformation for enforcing fault tolerance.
  (joint work with Boris Köpf)

Keywords: Security, static analysis, security type systems, information flow security

Long-term Security and Graceful Degradation

Jörn Müller-Quade (Universität Karlsruhe, D)

Computational assumptions are either true or false. However, infeasibility assumptions for concrete security parameters may be true now, but can become false in future. In this talk I will highlight several aspects of long-term security, i.e. security with respect to adversaries whose computational power increases over time.

Long-term security is not only important for many applications (e.g. those involving medical records), but opens up a new area of academic research between computational security and unconditional security. The talk focuses on three questions:

1. Secure Function Evaluation: Which functions can securely be computed in presence of an adversary which becomes unlimited after termination of the protocol?
2. Universal Composability: under which conditions do long-term secure protocols compose?
3. Digital Signatures: What level of security can, in the long run, be maintained by digital signatures?

Especially the last point hints in the direction of graceful degradation, where not all aspects of security can be guaranteed after the infeasibility assumptions are broken, but some authenticity properties can still be verified.

Point 2. is joint work with Dominique Unruh and will appear ath the Theory of Cryptography Conference TCC 2007. Point 3. is joint work with Stefan Röhrich.

Keywords: Long-term Security, Models of Security, Graceful Degradation, Signatures

Parsimonious Service Replication for Tolerating Malicious Attacks in Asynchronous Environments

HariGovind V. Ramasamy (IBM Research - Zürich, CH)

We present the parsimonious approach for constructing fault-tolerant protocols and as an example, we describe a parsimonious asynchronous atomic broadcast protocol that tolerates Byzantine faults of up to $t < n/3$ parties. It is the first protocol in this model with an amortized expected message complexity of $O(n)$ per delivered payload. The most efficient previous solutions are the BFT protocol by Castro and Liskov and the KS protocol by Kursawe and Shoup, both of which have message complexity $O(n^2)$. Like the BFT and KS protocols, our protocol is optimistic and uses inexpensive mechanisms during periods when no faults occur; when network instability or faults are detected, it switches to a more expensive recovery mode. In contrast to the BFT and KS and protocols, however, the parsimonious protocol uses more expensive public-key cryptography (in particular, digital signatures). We discuss the trade-off between message complexity and computational complexity. We also describe the experimental evaluation of our protocol in the context of a representative application implemented in the CoBFIT framework over LAN and WAN environments.

Keywords: Byzantine faults, atomic broadcast, state machine replication

Joint work of: Ramasamy, HariGovind V.; Cachin, Christian; Agbaria, Adnan; Seri, Mouna; Sanders, Bill

Full Paper:
http://www.crhc.uiuc.edu/PERFORM/Papers/USAN_papers/05RAM05.pdf

Failures: a basic issue in dependability and security - Their definition, modelling & analysis

Brian Randell (University of Newcastle, GB)

There are severe terminological and conceptual confusions in the dependability & security field(s). These come into particular prominence when one takes an adequately general view of dependability & security problems by avoiding the (naive) assumptions that systems always have well-established boundaries and fully-adequate specifications. This talk seeks to address these confusions by clarifying the concept of failure, and discussing the modelling of failure processes in complex systems. It takes as one of its starting points the concept of "Fault-Error-Failure" chains, and introduces the notion of "Structured Occurrence Nets" as a potentially promising means of formalising the definition and analysis of such chains.

Keywords: Dependability, security, failure, modelling, analysis, structured occurrence nets

Dependability: Relation between its Quality Attributes from a Software Engineering

Ralf Reussner (Universität Karlsruhe, D)

The short talk will briefly introduce the term "dependability" as a set of quality attributes, as seen in software engineering. After a clearification to the term "trustworthiness", two kinds of dependencies between quality attributes are elaborated and the quality attributes are grouped according to the scientific process used in their communites. It concludes with an attempt to relate security to the other quality attributes.

Keywords: Quality attributes, dependencies, research method of dependability

Anonymity and Reputation

Luis Rodrigues (University of Lisboa, P)

In ubiquitous networks, the multiple devices carried by an user may unintentionally expose information about her habits or preferences. A common approach to increase privacy is to hide the user real identity under a pseudonym that is changed frequently. Unfortunately, pseudonyms may interfere with the reputation systems that are often used to assert the reliability of the information provided by the participants in the network.

This talks addresses the problem of transferring reputation information from one pseudonym to another pseudonym in an untraceable manner. It proposes a technique to perform such transfer. Unfortunately, the link between two pseudonyms may still be established based on the values of reputation transfered. The purpose of the talk is to promote discussion on possible ways to circumvent this problem.

Keywords: Anonymity, Reputation

Joint work of: Rodrigues, Luis; Miranda, Hugo

See also: A Framework to Provide Anonymity in Reputation Systems. H. Miranda, L. Rodrigues. Proceedings of the 3rd Annual International Conference on Mobile and Ubiquitous Systems: Networks and Services (MOBIQUITOUS 2006). San Jose, California, July 2006. (Research-in-Progress Paper)

Enhanced Grid Security and Dependability Using Trusted Virtualization

Ahmad-Reza Sadeghi (Ruhr-Universität Bochum, D)

Grid applications have increasingly sophisticated functional, security, and dependability requirements. Most current techniques aim to protect only the Grid resource provider from attacks by the Grid user, while leaving the user comparatively dependent on the well-behavior of the resource provider. We describe an ongoing effort to address the trust asymmetry by using a combination of trusted computing and virtualization technologies. We present the key components for a trustworthy Grid architecture and outline a possible implementation. We propose a scalable offline attestation protocol, which allow selection of hosts with known and trusted configurations for job execution and delegation. We also propose a user-transparent and resource-efficient way of providing dependability guarantees for Grid jobs, based on secure delegation and migration of the jobs. By providing multilateral security, i.e., security for both the Grid user and the Grid provider, we believe that our architecture increases the confidence that can be placed on the correctness of a Grid computation and on the protection of user-provided assets.

This talk is based on joint work with Stefan Schulz, HariGovind Ramasamy, Christian Stueble, Hans Loehr, and Matthias Schunter.

Why not Use Several Models? - Dependability Properties under Refinement

Thomas Santen (TU Berlin, D)

Model-based approaches to develop dependable / secure systems often implicitly suggest that a single model suffices to capture and analyze all relevant properties of a system (with respect to a kind of property such as confidentiality or reliability). However, it is well-known that many dependability properties are not preserved under refinement. As a practical implication of this theoretical result, a correct implementation of the given model does not necessarily satisfy the desired dependability properties, as the model specifies them.

We describe a framework to investigate the conditions under which confidentiality properties are preserved under behavior refinement. We argue that several models at different levels of abstraction are needed to satisfactorily describe adversary capabilities and analyze their consequences on the security of a system. Confidentiality properties of an abstract model capture the requirements. The development process than needs to preserve those properties under successively extended adversary models.

Concluding, we suggest implications of the results for confidentiality on dependability in general.

Secure and Fault-Tolerant Clock Synchronization in Sensor Networks

Elad Schiller (Chalmers UT - Göteborg, S)

We present our recent results on secure clock synchronization mechanisms for sensor networks. Our design is useful against an attacker that is trying to disrupt the cluster management by tampering with the clocks. At any time, the attacker is in a particular bounded region that may change over time. Within that region, the attacker may eavesdrop to the local communications, introduce malicious nodes and compromise the system nodes.

Keywords: Security, Fault-Tolerance, Clock Synchronization, Sensor Networks

Joint work of: Larsson, Andreas; Schiller, Elad; Tsigas, Philippas

Full Paper:
http://www.cs.chalmers.se/∼dcs/

Virtual Leashing: Internet-Based Software Piracy Protection

Nir Shavit (Tel Aviv University, IL)

Software-splitting is a technique for protecting software from piracy by removing code fragments from an application and placing them on a remote trusted server.

The server provides the missing functionality but never the missing code. As long as the missing functionality is hard to reverse-engineer, the application cannot run without validating itself to the server.

Current software-splitting techniques scale poorly to the Internet because interactions with the remote server are synchronous: the application must frequently block waiting for a response from the server. Perceptible delays due to network latency are unacceptable for many kinds of highlyreactive applications, such as games or graphics applications.

This paper introduces virtual leashing, the first nonblocking software-splitting technique. Virtual leashing ensures that the application and the server communicate asynchronously, so the applications performance is independent (within reason) of large or variable network latencies.

Experiments show that virtual leashing makes only modest demands on communication bandwidth, space, and computation.

Keywords: Anti piracy, distributed security, software security

Joint work of: Shavit, Nir; Herlihy, Maurice; Dvir, Ori

See also: Ori Dvir, Maurice Herlihy and Nir N. Shavit. Virtual Leashing: Internet-Based Software Piracy Protection. ICDCS '05: Proceedings of the 25th IEEE International Conference on Distributed Computing Systems (ICDCS'05). IEEE Computer Society, Washington, DC, USA.

Intrusion Tolerance: the Road to Automatic Security?

Paulo Verissimo (University of Lisboa, P)

Intrusion Tolerance has become a reference paradigm for dealing with faults and intrusions, achieving security (and dependability) in an automatic way, much along the lines of classical fault tolerance. However, there are issues specific to malicious faults (attacks and intrusions) that bring about some problems and limitations to the paradigm as a basis for designing resilient systems, some of which quite unexpected. We wish to discuss the theoretical underpinnings of intrusion tolerance and its limitations in distributed systems, and evaluate some recent research results that address a few of those limitations.

Full Paper:
http://www.navigators.di.fc.ul.pt/it/index.htm

Secure distributed storage: recent results and open problems

Marko Vukolic (EPFL - Lausanne, CH)

Even after nearly three decades of intensive research on this popular topic, the quest for an ultimate secure distributed storage algorithms is not over.

This talk will go through several crucial dimensions in designing secure distributed storage algorithms. We will discuss major issues underlying these dimensions, as well as techniques used to address these issues. We will summarize some of the most important results of this research and discuss challenging open problems.

Keywords: Distributed storage, fault-tolerance, security, arbitrary failures

Joint work of: Chockler, Gregory; Guerraoui, Rachid; Keidar, Idit; Vukolic, Marko

Synchronous Consensus with Mortal Byzantines

Josef Widder (TU Wien, A)

It is often argued that the Byzantine failure model captures all kinds of reasons for components to fail, even intentional faults like intrusions. An important question is whether this is the only possible model, or whether even more restricted failure models could be suitable in order to address both reliability and safety requirements. In order to give some hint of what can be done from the reliability side, I present some recent work on novel fault models:

I will consider the problem of reaching consensus in synchronous systems under a fault model whose severity lies between Byzantine and crash faults. For these "mortal" Byzantine faults, one assumes that faulty processes take a finite number of arbitrary steps before they eventually crash.

I present a consensus algorithm that tolerates a minority of faulty processes; i.e., more faults can be tolerated compared to classic Byzantine faults. We also show that the algorithm is optimal regarding the required number of processes.

Keywords: Fault models, Consensus, Fault-tolerence

Joint work of: Widder, Josef; Gridling, Günther; Weiss, Bettina; Blanquart, Jean-Paul

Towards bounded wait-free PASIS

Jay Wylie (HP - Palo Alto, USA)

The PASIS read/write protocol implements a Byzantine fault-tolerant erasure-coded atomic register. The prototype PASIS storage system implementation provides excellent best-case performance. Writes require two round trips and contention- and failure-free reads require one. Unfortunately, even though writes and reads are wait-free in PASIS, Byzantine components can induce correct clients to perform an unbounded amount of work.

In this extended abstract, we enumerate the avenues by which Byzantine servers and clients can induce correct clients to perform an unbounded amount

of work in PASIS. We sketch extensions to the PASIS protocol and Lazy Verification that bound the amount of work Byzantine components can induce correct clients to perform. We believe that the extensions provide bounded wait-free reads and writes. We also believe that an implementation that incorporates these extensions will preserve the excellent best-case performance of the original PASIS prototype.

Keywords: Byzantine fault-tolerant, erasure-coded storage, bounded wait-free, non-skipping timestamps

Joint work of: Abd-El-Malek, Michael; Ganger, Gregory R.; Goodson, Garth R.; Reiter, Michael K.; Wylie, Jay J.

Extended Abstract: http://drops.dagstuhl.de/opus/volltexte/2007/848