Report on

DISTTRUST

Dynamic Security and Distributed Trust

28 April 2006

Barcelona

Table of Contents

1 Summary

The workshop Dynamic Security and Distributed Trust (DISTTRUST) was held on 28 April 2006 and brought together various projects funded by the European Commission. Whereas the overall objectives of these projects lie in different areas, they share common challenges when implementing trust and security in the context of open, heterogeneous and dynamic networks.

While the state of the art in this area is centred on static network topologies, it is clear that networks are becoming increasingly dynamic. Mobility of nodes and code, resource sharing, distributed computation and services, ad-hoc networking and opportunistic forwarding, all have in common that they can no longer rely on placing tight perimeters around fixed structures, nor on a central authority that controls access and transactions.

During the panel discussions it became clear that, although new networks have introduced increasing degrees of autonomy and adaptiveness, security mechanisms are still static. Creating dynamic security mechanisms, in particular in a way that would allow them to adapt to previously unseen attacks, is a great challenge that would require progress at a fundamental level.

Relating to this point, there was an extensive discussion about the clash between the networking community and the security community. These communities had little interaction before, and recent interaction has been characterised by various misunderstandings. The networking community strives to create complex and heterogeneous systems based on the new network infrastructures and massively available devices, but without a full understanding of the requirements from the security perspective. The security community would like to slow down this development, as the lack of understanding of new network infrastructures limits the ability to provide adequate protection.

The question was posed if technological developments are “breaking down walls” that should not be broken down for the sake of security, or whether “walls are falling down by themselves” leaving the security community no choice but to answer to this challenge in as far as possible given the available means. This was also related to economic factors, which may be driving researchers to making a more critical trade-off with regards to security when designing new systems. A better collaboration between the two communities was advocated, in particular with an emphasis on designing well understood systems and proper assessment of the extent to which security can be provided in new systems.

This led to an extensive discussion regarding the term “security” which carries a meaning of guaranteed or proven security to some people, and a meaning of “security with high probability” to others. In fact many projects that are investigating the design of new networking systems look at mechanisms that provide some form of “best effort” security. While it was argued that these projects are just trying to find minimal assumptions to build the best possible applications that can account for security needs, it did also become clear that security can be approached with different modalities (i.e. proven or best effort), depending on the security needs.

This was exemplified by the dichotomy between the security of the device and the security of the individual. In a networking context a device can be given probabilistic security, and some devices may be sacrificed if the whole system prevails, for example in the case of a sensor network. In the context of providing privacy of individual data it is however hardly acceptable to sacrifice a few individuals as part of a security policy.

There was agreement that the increasing number of devices which is being introduced, and will start to form part of emerging networks, creates a need for trust and cooperation mechanisms between ad hoc coalitions to allow the workflow between untrusted entities to be bootstrapped. The role that trust fulfills was also discussed extensively, including the question whether trust-based security is preferable over security-based trust or vice versa, although at the same time there was an agreement that this can not solely form the basis of network and service security.

The need for taking into account legal issues and collaboration with actors from the legal area was emphasised. In particular, it became clear the future systems need to respond to new legal requirements on logging and providing data for auditing. At the same time recent networks and applications (such as peer-to-peer applications that take actions away from traditional juridical persons) create new challenges for juridical frameworks.

There was a discussion about the way the field is covered by the projects, and the synergies that should take place between them given their position as compared to other projects. This turned out to be a good and comprehensive covering, and the discussion described in the remainder of this document shows concrete points where projects can create leverage from their complementarity.

The discussion regarding the challenges that were identified in the coordination action “Beyond the Horizon” as part of the working group Security, Dependability and Trust received support from the participants. Specifically, the first challenge “ambient security, dependability and privacy” could be linked to the above points, reflecting a need for bridging the networking community and the security community, taking into account the need for provable security as well as the protection of open, heterogeneous networks that are not well understood. The second challenge “Dynamicity of Trust” was seen as providing a pivotal element in emerging networks, responding to its ad-hoc nature and the need for social acceptance of new systems. The third challenge “Quantum technology and cryptology for information security” was seen as identifying an important area which is not covered by the projects represented at this workshop.

2 Introduction

The workshop Dynamic Security and Distributed Trust (DISTTRUST) was held on 28 April 2006 and brought together various projects funded by the European Commission. Whereas the overall objectives of these projects lie in different areas, they share common challenges when implementing trust and security in the context of open, heterogeneous and dynamic networks.

While the state of the art in this area is centred on static network topologies, it is clear that this is becoming increasingly dynamic. Mobility of nodes and code, resource sharing, distributed computation and services, ad-hoc networking and opportunistic forwarding, all have in common that they can no longer rely on placing tight perimeters around fixed structures, nor on a central authority that controls access and transactions.

Key aspects in the participating projects are the dynamicity of the network and the distributed nature of services. Under these conditions, where the concepts of location and identity are becoming separate, it is no longer possible to assume that one authority would have knowledge of every node or device in the network, let alone be able to guarantee all the actions that device takes when it is used as a resource or provides a service.

Some of these projects are funded by the Future and Emerging Technologies (FET) Unit, and some of them by the ICT for Trust and Security Unit. Wide Hogenhout (EC, FET unit) and Bart Van Caenegem (EC, Unit D4) opened the workshop with an introduction on the respective unit activities and project portfolios, the objectives of the workshop and a lookout at the FP7 preparation and timeline.

FET (Future and Emerging Technologies) is the IST Programme nursery of novel and emerging scientific ideas. Its mission is to promote research that is of a long-term nature or involves particularly high risks, compensated by the potential of a high societal or industrial impact. FET goals are achieved via an “open” scheme and a “proactive” scheme. While the open scheme is constantly open to proposals for bold and visionary research on any subject related to IST, the proactive scheme is focused on a small number of proactive initiatives. These are strategic areas identified by the EC (in consultation with the main stakeholders) as holding particular promise for the future, in order to open new possibilities and set new trends for future research programmes in Information Society Technologies.

Three projects are from the FET proactive initiative Global Computing (GC), namely AEOLUS, MOBIUS and SENSORIA. These projects focus on common characteristics representing a family of potential or actual global computers described by appropriate abstractions. These abstractions can be thought of as "overlay computers", i.e., abstractions that can be implemented on top of global computers to yield enhanced classes of global computers that are programmable and computationally complete in their application domain.

Four of the projects involved are from the FET proactive initiative Situated and Autonomic Communication (SAC): ANA, BIONETS, CASCADAS and HAGGLE. These work in the area of new paradigms for communication/networking systems that can be characterised as situated (i.e. reacting locally on environment and context changes), autonomously controlled, self-organising, radically distributed, technology independent and scale-free.

The ICT for Trust and Security unit (Unit D4) is a research unit within Directorate D (Networks and Communication Technologies) of DG Information Society and Media. The mission of unit D4 is to support research on the security and dependability of ICT based systems and services. The work supports the development of knowledge and technologies to manage and control complex and interdependent systems, in order to secure modern information systems and networks and to build resilience in the Information Society infrastructure. The activities also cover research and its interplay with policy developments in trust and security, including biometrics, identity and privacy management, authentication, secure digital assets and cyber crime. In addition, work includes cooperation with Member States initiatives to develop the European Research Area dimension in this domain and targeted international co-operations on security research.

The main instruments for the implementation of the mission are the Community Research and Technological Development (RTD) Programme and support activities for the development of policies in the area of ICT for trust and security. Research projects on Security, Dependability and Trust (SDT) were solicited in FP6 in the Strategic Objective “towards a global security and dependability framework”. The unit is funding 35 RTD projects with around 140 M € EU funding. Of particular relevance to the workshop are five projects with main focus on trust and security aspects in open and dynamic environments. These are FIDIS, PRIME, SERENITY, S3MS and UBISEC&SENS.

The goal of the workshop is to foster synergies, avoid duplication, and exploit common approaches across the participating projects. The results of the workshop will also be used as input to the preparation of the FP7 work programme.

The participants were informed on the state of advancement of FP7. The work programme is being drafted at the moment of speaking and a final first draft of the work programme is expected to be ready in July. Adoption and first calls for proposals can be anticipated to happen just before or just after Christmas this year.

Overview of participating projects

From SAC:

BIONETS (BIOlogically-inspired autonomic NETworks and Services) is a project that aims at developing a biologically-inspired approach (from nature and society) to localised autonomic communication services without central control, allowing high-level services to evolve spontaneously. (See www.bionets.org)

ANA (Autonomic Network Architectures) is focused on adaptation and reorganization of the network, and aims to create a novel network architecture (beyond IP) enabling flexible and autonomic formation of network nodes according to working, economic and social needs. (See www.csg.ethz.ch/research/projects/ANA)

HAGGLE (An innovative Paradigm for Autonomic Opportunistic Communication) has as objective a cross-layer network architecture exploiting intermittent connectivity, which supports opportunistic networking paradigm. (See http://www.cambridge.intel-research.net/haggle/eu)

CASCADAS (Componentware for Autonomic, Situation-aware Communications and Dynamically Adaptable Services) is focused on self-similarity, autonomic component-ware, and aims to define a new generation of highly distributed, pervasive, situation-aware, semantically self-organising communication-intensive services. (See www.cascadas-project.org)

From Global Computing:

MOBIUS (Mobility, Ubiquity and Security) will investigate trust and security for small devices which are functioning as a part of global computers. The main focus is on proof carrying code, a paradigm aimed at checking previously created proofs with modest computational resources. (See mobius.inria.fr)

SENSORIA (Software Engineering for Service-Oriented Overlay Computers) will develop a novel methodology for engineering service-oriented overlay computers and for building a framework for context-adaptive, personalisable global services. (See sensoria.fast.de)

AEOLUS (Algorithmic Principles for Building Efficient Overlay Computers) aims at developing algorithmic principles and implementing the basic functionalities (i.e., programming tools, trust management, secure distributed computation) to enable transparent and efficient access to an internet-based global computer. (See www.ceid.upatras.gr/aeolus)

From Unit D4 (ICT for Trust and Security):

FIDIS (Future of IDentity in the Information Society) is a Network of Excellence that develops and promotes new identity and privacy protecting concepts for Europe. (See http://www.fidis.net)

PRIME (Privacy and Identity Management for Europe) is an Integrated Project that develops and demonstrates a novel identity management system with particular attention to European privacy protection requirements. (See http:/www.prime-project.eu)

SERENITY (System Engineering for Security and Dependability) is an Integrated Project that develops a framework to support the automated integration, configuration, monitoring and adaptation of security and dependability mechanisms in services by capturing security expertise in Security and Dependability patterns and integration schemes. (See http://www.serenity-project.org and http://www.serenity-forum.org)

S3MS (Security of Software and Services for Mobile Systems) is a Specific Targeted Research Project (STReP) that develops a framework based on security-by-contract for trusted deployment and execution of mobile applications, in which the contract expresses the security features and requirements. (See http://www.s3ms.org)

UBISEC&SENS (Ubiquitous Sensing and Security in the European Homeland) is a Specific Targeted Research Project (STReP) that addresses security in wireless sensor networks and develops security aware components to facilitate trusted sensor network applications. (See http://www.ist-ubisecsens.org)

3 Session 1

This panel discussed large scale dynamic environments, including distributed systems with self-* properties (self-organisation, selfmanagement, self-healing, self-evolution, etc.), in which the key challenges include the reaction to new threats as they arise.
In these systems the response to threats or attacks may be of a nondeterministic nature, due to the autonomic behaviour of the elements which may influence each other and change over time.

Presentations

Self-organizing Trust establishment and Cooperation Enforcement
Project: HAGGLE
Speaker: Melek Önen, Institut Eurecom.
Melek Önen first described the general setting in Haggle. Based on this environment she explained the general security challenges in this project and outlined the work to be conducted in the security work package.
The talk focused on the trust and cooperation issues, discussing the requirements for possible solutions as well as indicating existing solutions and potential directions of the project. Especially in the latter case trust establishment without identities, optimistic security protocols, symmetric techniques for key management, and hybrid cooperation enforcement schemes were presented.
Finally, general security problems that will be addressed in this project were mentioned. Given the “unorganized” network structure, these include key management and communication security issues.

Autonomic Self-Defence
Project: CASCADAS
Speaker: Erol Gelenbe, Imperial College.
This talk gave an insight in the work conducted in the CASCADAS (Component-ware for Autonomic Situation-aware Communications, and Dynamically Adaptable Services) project. Focusing on security aspects he reviewed the family of denial of service attacks, the currently existing solutions to detect them and methods to automatically respond to them. Several approaches, their advantages as well as their deficiencies were shortly discussed. At the end of the talk Erol Gelenbe presented a new approach investigated by the project which uses Neyman-Pearson Decision Rules. A rough outline of the overall architecture, the decision rules as well as the mathematical model used in this approach concluded the talk.

Securing BIONETS: How can Security Infrastructures Match Autonomically Evolving Networks and Services?
Project: BIONETS
Speaker: Daniel Schreckling, Univ. of Hamburg.
After a short introduction which explained the general concept of BIOlogically-inspired autonomic NETworks and Services (BIONETS), the talk discussed the security issues the members of the security group face in this project. The security challenges of this approach were emphasized through a comparison of the characteristics of BIONETS with traditional security assumptions and requirements. One component in BIONETS is services which autonomously adapt to their environment based on evolutionary mechanisms. Daniel

Schreckling explained some of the security challenges that this concept yields. Apart from service (re-)combination on one or more network nodes he identified service evolution as one of the major security problems to be solved. As a vision, BIONETS proposes the evolution of services which provide security functionalities.

Dynamic Security in Ambient Intelligence Scenarios
Project: SERENITY
Speaker: Antonio Maña, University of Malaga.
The SERENITY (Systems Engineering for Security Dependability) project is concerned with the development of a framework supporting the automated integration, configuration, monitoring and adaptation of security and dependability mechanisms for Ambient Intelligence (AmI) ecosystems. Its work focuses on the development of mechanisms for monitoring and diagnosis of threats and violations of security requirements and recovery from such violations. Antonia Maña presented the four main challenges in this project. The first challenge is the adaptation to dynamic changes to hardware and software, and to the unpredicted and unpredictable appearance and disappearance of devices and software components. The second is the ability to deal with heterogeneous computing and communication infrastructures and devices.
A further challenge is providing some monitoring capabilities. Finally, the fourth challenge is the dynamic application of the expertise of security engineers in the foreseen AmI ecosystems, which is in fact one of the key problems in SERENITY.

Adaptive Security with Aspect Oriented Programming
Project: MOBIUS
Speaker: Maarten Rits, SAP Research.
MOBIUS (Mobility, Ubiquity and Security) aims to develop a technology for establishing trust and security for the next generation of global computers, using the Proof Carrying Code (PCC) paradigm.
In his talk Maarten Rits explained how Aspect Oriented Software Development (AOSD) can be used in the software life-cycle to satisfy new and evolving security requirements. Using a simple example he described fundamental terms in AOSD and showed how it is used in Java environments by annotating code. Maarten Rits further showed a framework which will support the enforcement of high-level security properties. Manually added source code annotations will help to automatically infer secondary annotations. Based on this metadata new code will be written and woven into the actual source code to satisfy the security requirements. A final step will proof the adherence of the modified code to high level security properties. At the end of the presentation this concept was clarified using a case study.

Panel Discussion

Dynamic versus static approaches
The discussion started with the observation that almost all projects presented in this panel are based on distributed environments. A common characteristic of such systems is the absence of a central control, and security mechanisms based on centralized infrastructures may therefore not be feasible. A first question was whether trust and reputation systems would be able to form a possible solution for such systems and provide enough security for the emerging network architectures.

There was a general agreement in the audience as well as in the panel that the answer should be negative, as trust and reputation systems and network and service security are controversial concepts which can not be merged.

Apart from the distributed aspect of the projects, it was observed that the presented systems share their adaptive nature. Current security infrastructures always seem to be one step behind current developments in technology, and security tends to be employed after the actual network infrastructures have already been defined. This led to the question whether it would be possible to develop security and network infrastructures in parallel and what possible consequences this would have on the nature of the developed techniques.

The panel argued that, as current approaches to security do not show any dynamism, it is intrinsic that security is very static and thus also dependant on the underlying architecture and infrastructure of new systems. Even if security shows some dynamism and is able to adapt to changing environments, it still only reacts based on specific patterns. These patterns can either be based on misbehaviour or behaviour that does comply with established security policies, and current systems integrate these patterns during the implementation of the security mechanisms. An example is the source code annotations that are used in MOBIUS. Such a very static measure to prevent a specific security threat illustrates how the reaction of security systems currently only takes place after the discovery of threats, and knowing them beforehand is a matter of preventing attacker technology from being one step ahead.

This raised the question if it is possible to annotate code in a way that it is also resilient against unknown attacks. The panel stated that it is currently not possible but that it may be feasible in the future. This is explicitly valid if more effort is put into the discovery of errors and malicious behaviour during runtime. In the far future such recognition mechanisms may also be possible for static methods analyzing for example binary code before execution.

Another question focused on the dynamic, adaptive, and evolutionary characteristics of some of the presented environments: Is it possible to design a system that accounts for unknown attacks by adopting mechanisms from the immune system? Are there approaches that try to implement similar behaviour?

The panel argued that we currently do not completely understand how the immune system actually works. Furthermore, there is strong evidence that the immune system can not defend against unknown attacks. The system must have seen the virus before it is able to react on it. However, concepts like vaccination which support the immune system to win the fight against unknown viruses may be adopted, and existing approaches such as the IBM decoy “cells” reflect the adaptation of mechanisms similar to the immune system.

Proven security versus best effort

Another interesting comment from the audience concerned the general concept of security as presented in the talks given. It was argued that the term security is not being used in accordance with its original meaning, since security used to mean that it is possible to give some form of guarantee, whereas the presentations reflected a notion of security which does not imply security in that sense. The projects will aim at developing mechanisms which are trying their best to achieve a certain level of security, but rather than aiming to provide a guarantee, they are only designing best effort security.

The discussion centred on the question whether the term security should be associated with best effort. One of the participants argued that if the security mechanisms that emerge from the presented projects have this best effort property, a different terminology should be used rather than redefining security, since making it fit in the described context might lead to misunderstandings.

The general reaction of the panel showed disagreement with this comment. It is true that there is a best effort approach, however this approach only concerns the design phase in which the projects try to find minimal assumptions on which one may build the best possible applications able to account for security needs. Also, it was argued that security has always been an attempt to give as much guarantee as possible, which in many cases is based on assumptions that may turn out to be false.

The discussion continued on the comparison with the term “Quality of Service” (QoS). Twenty years ago the term QoS was coined, and the mechanism used to make it concrete worked fine as long as the hardware was fairly stable. Over the years the number of links and the number of devices involved in the provisioning of QoS increased, elevating the risk of failing hardware, and as a consequence the proposed approaches only worked to a certain degree. The situation is similar for security. For example, in the case of ambient intelligence, the concept of security effectively introduces a completely new paradigm. This will certainly yield interesting problems and new directions of research, but the question is whether this isn’t weakening the concept of security which has been developed over decades. In that case we would face the choice between aiming at a redefinition of current systems to come up with secure systems in the original sense of the term, or consider that the assumptions under which security was defined before simply do not hold any more, as the environment in which the mechanisms are applied have changed completely.

Tearing down walls versus watching them fall down by themselves

Based on the previous discussion the following case was made: from a security perspective it is more reasonable to head towards systems which are well defined, rather than designing increasingly chaotic and complex systems which we do not yet understand. In this case the task is not to find minimal assumptions but to find assumptions needed to approximate a level of security which suffices with a probability that is close to one.

There were two reactions to this, which were guided by two different opinions. The first argument was based on cost which must be drawn into consideration when developing new systems. Recent developments in Europe force universities to work more effectively and more economically, and as a consequence research is sometimes also driven by economical interests. This often leads to a trade-off between developing new systems and deploying security in such systems.

In contrast, the second argument considers this shift as being more fundamental. It was argued that the observed problems arose when two communities that had little interaction before, the network community and the security community, started to cooperate. It appears that the network community aims at deploying everything that is technologically feasible. Afterwards the security community tries to provide security for the new environment without having been integrated in the actual design process.

Based on the last comment the need for well defined systems was emphasized.

Almost all projects presented in this session set up new network infrastructures. These new environments will combine very complex heterogeneous systems which are not well understood – in fact even current systems are not entirely understood yet. Again, the question is whether we are tearing down walls which we should keep up in order to be able to design strong and reliable security mechanisms.

This was mitigated by the argument that we are not tearing down walls, as they collapse by themselves. The trend toward the combination of heterogeneous systems is driven by current network infrastructures and by new technologies. Especially the later one yields many cheap, small, heterogeneous and computationally powerful devices. To keep up with this development we have to develop new solutions that match these new environments, and from

the presentations it was clear that the projects work within these new environments and try to produce solutions that are feasible in future network infrastructures.
Another argument referring to this issue was that these walls have never been up before. This is somehow reflected by one of the main characteristics of security today. Firstly, we are currently only reacting on emerging threats and on attack patterns. Secondly, there are a lot of areas in the realm of security which we simply do not understand. One example may be symmetric cryptography which we rely on but which we can not prove to be secure. Another example is cryptographic hash functions which we only consider to be secure because we were told that they are secure - but if we take a closer look we can not actually prove that they are secure. It is not possible to tear down these walls as we currently only perceive security to be of probability close to one, and therefore one goal of these research projects is to design security mechanisms that provide a level of security and address most of the security threats created by new networks. Ultimately, a major task in all of the presented projects should be to find out to which extent the emerging mechanisms can provide security.

A final comment on the question whether the notion of security changed, emphasized that the risk of a misunderstanding by the public should definitely be avoided. This comment is based on the widely spread concept of security which conflicts with the security systems which are intended to be developed in the presented projects. Yet, nobody knows what type of effects on security may emerge when networks are spread in many places, there is no ultimate answer and it is the focus of current research. As an example, there are occasions in which collateral damage occurs and in which securing one network location may yield damage in another location. The effects are still unknown and the system is not understood completely.

Conclusion

At the end of this panel discussion a comment nicely described the witnessed clash between the network community and the security community. This can be considered to be a good summary of the problems discussed in this panel.
Up to today, the networking groups and the network security groups can be considered as two separate communities with little interaction. Most of the time it is the network community that comes up with new approaches toward architectures and network infrastructures, based on developments which are often driven by technology. The security community always tries to react upon such new developments. In recent years, both communities started to interact, but unfortunately this interaction appears to be characterized by many misunderstandings. Networking groups try to understand what is required from a security perspective, but it is not easy for them to completely understand these requirements, and as a consequence these requirements are most of the time not integrated in the development process of the networking paradigms. The application of security mechanisms ex post often fails or is insufficient.
As a first step the security community will have to try hard to adjust their mutual interests with the networking groups to come up with a secure network structure. Second, the security community will have to limit the activities of the networking community, which will otherwise be tempted to do everything possible. This means that security will have to be integrated in the design process right from the beginning and with the support of the security community.
Finally, the developed solutions will need a proper assessment phase in which the extent to which security can be provided in a new networking paradigm is analyzed.

4 Session 2

The session “privacy and trust for dynamic coalitions” was focused around a number of topics related to privacy and trust such as reputation mechanisms, the definition of requirements for privacy and trust, personal identification information, anonymity and accountability, usability and security for sensor networks. The presentations by the speakers were followed by an animated discussion.

Presentations

Reputation mechanisms for manet and P2P over manet
Project: CASCADAS
Speaker: Pietro Michiardi, Institut Eurecom.
Pietro Michiardi presented the ideas behind the game theory approach to security for autonomic communication that he is currently developing in the framework of the project. This is based on the assumption that a system designer has no control over the behavior of software agents and humans, and that their opportunistic or selfish behavior will prevent the system from conforming to the objectives of the designer. By modeling the behavior of the participants it is possible to create systems in which the selfish behavior of individuals results in the realization of the system’s goals.

Privacy and Trust requirement engineering
Project: SENSORIA
Speaker: Fabio Massacci, Univ. of Trento.
Fabio Massacci presented work on security and trust requirements capture for service oriented engineering that is currently carried out in the framework of the project. He presented an industrial case study on human resource management that was described and analyzed with the new security and trust features added to the requirements engineering process advocated in the project.

Identity management in social networks
Project: FIDIS
Speaker: Andreas Pfitzmann, Technische Universität Dresden, and Marit Hansen, ICPP.
This presentation was on the topic of identity management in social networks. Giving a central role to the phenomena of “distrust” (with one t), it addressed the topic of personal identification information. It was argued that there is a need for mechanisms that provide digital pseudonyms, i.e. a suitable combination of anonymity and accountability. This can be supported through mechanisms that securely transfer authorising signatures between different pseudonyms of the same party.

Usable Security in Dynamic Systems
Project: FIDIS
Speaker: Sebastian Höhn, Albert-Ludwig-University Freiburg.
This talk discussed the key issues in usability showing statistical evidence that a large number of users do not perceive security as a major issue, as 65% of the respondent to a comprehensive survey were not interested in learning about security features available to them. He argued that usability through the study of Human-Computer Interaction should play an essential role in security and be linked to security mechanisms.

Security Challenges and Solutions for Wireless Sensor Networks
Project: UbiSec&Sens

Speaker: Dirk Westhoff, NEC Europe.
Dirk Westhoff presented the latest trends in sensors networks. He argued that there is a need for lightweight security, extremely low-cost mechanisms in terms of CPU power, and a trend towards “probabilistic security” in which an attacker can only obtain a limited gain. In this context the project has formulated a number of the resulting technical challenges such as secure data aggregation and secure routing.

Trust and Reputation Issues in BIONETS
Project: BIONETS
Speaker: Bruno Crispo, Univ. of Trento.
Bruno Crispo presented challenges that have been identified as part of the BIONETS project. In particular he pointed out that many more devices are starting to populate the environment. Specific challenges were adaptiveness to the different capabilities of devices, context awareness and user interfaces, the disappearing security perimeter, self-evolving security, trust and reputation, and privacy.

Panel Discussion

The discussion among the members of the panel was interesting and lively and a number of similarities were pointed out. As opposed to the design issues addressed in the discussion after the first panel, this panel focused more on the difference between device and individual, privacy, and the related legal aspects.

A first topic that was discussed was the dichotomy between probabilistic security versus all-or-nothing security, which can perhaps also be phrased as the security of the device versus the security of the individual.
There is a significant difference between the security of devices and the security of personal information. In the case of devices, it is clear that a policy can be formulated that allows some devices to be sacrificed. When a complete system consisting of many devices is attacked, the whole system survives due to the large number of devices that will not be affected, even if a small number does not receive adequate protection. In the case of a sensor network, one doesn’t care very much about security of a single sensor, but the emergent security of the overall network is important. There was a consensus that in some settings it is appropriate to consider providing probabilistic security.

In contrast, the security and privacy of the data of an individual cannot be sacrificed for the benefit of the group or as part of a security policy. In this case a probabilistic privacy of the overall set of users is not acceptable if it implies risking the privacy of some individuals.

The discussion continued on the topic of privacy. Privacy is sometimes regarded as concerning (only) the protection of data that is provided by some individual, but it was observed that this view is too narrow. This is relevant for a number of proposals in the area of privacy which claim that data should stay with the owner of that data as much as possible. This approach is technically feasible, although the issue of scalability was raised as a serious challenge, and a suitable mechanism of delegation may be the only possible way to address this. However, even if that is resolved, it still misses what emerged as the main point: it is a misplaced belief that the data provided by the individual is the only data that must be protected. In fact some of the most interesting data about an individual is not produced by the individual himself or herself. For instance, the steps of interaction the client with a bank is by far less sensitive than the credit rating about the very same client produced by the bank.

There was an overall agreement that there will be an increasing number and variety of new devices, and that most of these devices will try to interact with other devices, creating a growing need for a way to bootstrap a workflow among untrusted entities. This creates a growing trend towards solutions involving trust and cooperation for ad hoc coalitions. From the presentations it was clear that several projects are recognizing this need and laying the foundations for such interaction in future systems.
Relating to this there was a significant discussion on the role that trust fulfills and whether trust-based security is preferable over security-based trust or vice versa. A key problem is that when the trustworthiness of an actor has to be assessed, there is a need for a reputation mechanism, and it appears that such a mechanism can only function if actors provide honest information.

The discussion also dealt with the topic of monitorability of security. A number of regulations (SOX; Basel-II) are introducing significantly stricter requirements on logging and auditing data. Some of the same ideas are being applied to IT issues and have begun to surface in EU legislation (e.g. the new data retention act). It was observed that certain properties, such as faking integrity, can be monitored whereas others, for example loss of confidentiality, cannot be. This shows how monitorability, while it is of growing importance, has natural limitations as part of a security policy.

Another topic related to legal enforcement, as opposed to the technical side, was the increasing awareness among lawmakers of the potential risks introduced by IT systems and the management of those systems. Lawmakers have significantly increased the amount of legislation in this area, but there is a widening gap between the legal expectations and the technical capabilities that can be offered by current technologies. Consequently not everything that is legal can necessarily be enforced, and not everything that is enforced is necessarily legal.

The discussion also elaborated on another legal debate. Most legislation deals with juridical persons and regulates what the juridical persons should do or may not do with some level of approximation. However, new approaches to networking such as P2P applications introduce an application that behaves like a singleton, which does not map well to the traditional juridical persons. This creates a legal void which is likely to become more prominent when radically distributed systems become widespread.

The discussion also touched on the issue of usability, in particular around the question whether users would be satisfied if current security mechanisms would become easier to use. The observation was made that most users, in their view of what should constitute usable security, would like to be able to control it with a knob that can be rotated left and right to obtain less or more security as they see fit.

5 Session 3

The third session addressed fundamental aspects of distributed security, such as the resilient nature of the network, the security of mobile code, heterogeneous networks, formal modelling languages and cryptographic tools and protocols. This demonstrated the need to combine the design of radically distributed networks with fundamental progress in program analysis and verification techniques.

Presentations

ResiliNets: Multilevel Survivable and Resilient Autonomic Networking
Project: ANA
Speaker: James Sterbenz, University of Lancaster.
This project concerns techniques for making large-scale networks *resilient* towards various challenges to normal operation. These challenges range from stochastic hardware failures, through different forms of concerted malicious attack, to drowning in more legitimate traffic than the network can optimally handle. The current engineering state-of-the-art emphasizes *fault-tolerance*, which only targets naturally occurring component failures. There is various ongoing work on general theories for *survivability* which concerns deliberate attacks and large-scale disasters. Full *resilience* should also ensure graceful degradation of service when the network is drowning in its own success.
The underlying premise of the ResiliNets project is that real resilience can only be achieved by a design that includes *cross-layer optimizations* in which control influence bypasses the rigid layer separations of the ISO stack model of network design. A core problem is to find ways of doing this without regressing to a “spaghetti design” where the architectural benefits of using the layer abstraction have disappeared. Furthermore, rapid response to challenges demands that the cross-layer optimizations are autonomous.

Scenarios for Security based on Proof-Carrying Code
Project: MOBIUS
Speaker: German Puebla, Universidad Politecnica de Madrid.
This talk discussed the option of eliminating the need to trust the supplier of a piece of mobile code (for example, a web applet). This is to be achieved by *proof-carrying code* – i.e., bundling each package of code with a mechanically checkable mathematical proof which shows that the machine code or bytecode in it satisfies a certain security policy. No cryptography is involved, and it is not necessary to trust anything but the mathematics that went into creating the proof checker on the client side.
There are two basic schemes for creating proofs. In the first scheme, the compiler that produces the machine code also automatically creates a proof that the code is safe. This works well in the case of, for example, simple type safety, and the Java bytecode verifier can be seen as an implementation of this idea. In another variant, the proof is produced by a human-guided semi-automatic theorem prover at the supplier end and again checked automatically before the code is allowed to start. This is used for more high-level properties, such as a guaranteed limit to memory usage, or other resource properties (guaranteed running times, liveness).

Security by contract: A manifesto
Project: M3MS
Speaker: Fabio Massacci, University of Trento.
M3MS aims to handle the problem of automatically negotiating a security policy for mobile code. The currently deployed schemes of cryptographic certificates at most tell you *who* wrote a piece of mobile code – and sometimes the only information is that the unknown author has

managed to convince Microsoft or your telco that he or she means no harm. The problem is that even if the user knows who the author is, one does not know what their intentions are, and they are probably not aligned with those of the user.

This would be improved if the supplier of the code could declare explicitly which security properties it has (a "*contract*"), in a formal notation that the end-user device can easily and automatically compare to its particular configured demands. Conversely, the device could transmit a representation of its demands (a "co-contract") to the supplier, and the supplier could use the information to transmit the "best" agent that fits within the boundary conditions.

The project has not fixed on a strategy for enforcing that the contract is adhered to. If the author has been authenticated, one could consider conventional legal liability if the contract turns out to be a lie, but this is not technically desirable due to the possibility of unintended bugs in programs. A more technical measure would be to monitor the actual running of the code ("sandboxing"). It is envisaged that the eventual system will be a backwards compatible extension of the currently deployed cryptographic code authentication systems.

Security and trust issues for services

Project: SENSORIA

Speaker: Fabio Martinelli, IIT–CNR.

The SENSORIA project takes a higher-level view of the *architecture* and *design* of global computing systems from a service-oriented starting point. It views the communication among multiple agents in a system as consisting primarily of *service invocations* and responses, and advocates a "linguistic" approach to services, where an important activity is expressing system designs in formal languages that make them amenable to automatic analysis. The project is researching several such languages that can be used to model systems at different levels. Given a system description, automated tools could either determine whether a given level of trust can establish some desired property, or conversely determine who and what you need to trust in order for a particular property to hold. Analysis techniques include model/module checking, advanced type systems and control flow analysis.

Secure composition of secure protocols

Project: AEOLUS

Speaker: Guiseppe Persiano, University of Salerno.

The AEOLUS project works with deep algorithmic properties that must be mastered in order to build and trust large distributed systems. The focus of the talk was a particular example of the kind of problems that are considered in the project. This concerned a probabilistic cryptographic protocol in which one party seeks to prove to the other that she knows the plaintext corresponding to a given ciphertext, without actually revealing the plaintext to the other party. A particular protocol which works for RSA encryption was presented. Its peculiar property is that it is safe when run between just two participants, but it is *unsafe* to run several sessions of the protocol in parallel. This example highlighted the need to gain a general insight into the question which security properties of a protocol are preserved by various composition operators (such as parallel composition).

Panel Discussion

This panel discussion looked at how the different areas are covered by the different projects, which in particular relate to network resilience (ANA), mobile code (MOBIUS and S3MS), the formal language and service based approach (SENSORIA) and cryptographic protocols

(AEOLUS). In particular in the area of formal verification this demonstrates is a strong coverage of various elements in the area of distributed computing, where services and code are securely exchanged, with strong protocols on resilient networks.

Both MOBIUS and S3MS share related goals, as in fact they attack similar problems using complementary approaches. Both projects deal with security of mobile code in heterogeneous networks. S3MS proposes the use of security-by-contract, while MOBIUS is based on the Proof-Carrying code (PCC) paradigm and proposes longer term research as among its aims it includes improving the state of the art in Program Analysis and Verification.

M3MS does not intend to investigate PCC very deeply, but in this case there is a potential for interesting synergies with the MOBIUS project. For example, even in the presence of PCC the end device would not *necessarily* check proofs itself. One can conceive paying a trusted third-party service to check proofs off-line and transmit cryptographic signatures to small devices that do not have the power to run a proof checker on their own.

PCC, which is central in the MOBIUS project, attracted attention from the audience. During the panel discussion, several questions were related to finding out more details about the kind of safety policies which can be handled using PCC and about why checking certificates is simpler than verifying the code on the producer.

Current research is ongoing about how to combine the two main approaches (proofs automatically created by the compiler and a human-guided semi-automatic theorem prover) more seamlessly, and also on automatic translation of proofs carried out at the Java source code level into proofs that work for the generated bytecode. As designs for PCC currently assume that all parties agree on the security policies, there is an area of future research and potential synergy with other projects related to schemes for negotiating which security policy to prove and check in a given situation.

It became clear that program verification and analysis is at the heart of the technology required for the projects represented in the panel. The discussion around the various aspects of verification brought forward that verification ought to be at least a two-step process: first proving that the algorithm has the right properties; and then proving that the code which has been developed faithfully implements the algorithm.

6 Discussion on Beyond the Horizon

The Future and Emerging Technologies (FET) Programme of the European Commission funds a Coordination Action called "Beyond the Horizon" which aims at defining the major challenges and promising research directions that FET could support in the forthcoming FP7, by organising consultations with leading experts from the field, and through a web forum on which feedback was given by the community on the initial version of the report.

One working group, led by Prof. Michel Riguidel from the Ecole Nationale Supérieure des Télécommunications in Paris, dealt with the topic of "Security, Dependability and Trust". Wide Hogenhout presented the main challenges that were identified in this working group.

The three main challenges are:

1. **Ambient security, dependability and privacy**. The mass diffusion of networked devices must be supported by mechanisms for enhancing confidence on their usage, but traditional perimeter based mechanisms cannot cope with the large numbers of devices and their mobility. Moreover, security and cryptographic mechanisms are needed for very small networked objects, such as RFIDs and sensors, which lack the necessary resources (such as energy) for using existing techniques and protocols. The challenge is to provide security, dependability and privacy in this environment, which may require adaptive or evolvable mechanisms. This will need to be supported by technologies for provable security and dependability, as well as mechanisms for assessing the trustworthiness of an open and complex information and communication infrastructure.

2. **Dynamicity of Trust.** The lack of trust is one of the main barriers for the establishment of a secure and dependable Information Society. Instead of depending on security perimeters, it will require a capability for managing and negotiating trust relationships, adapted to the level of security required in a given situation. The challenge is then to obtain greater understanding of partial trust, security-based trust, and trust-based security (where security is achieved through a trusted partnership), and to use this understanding to realise a high level of trust of the citizen in the deployment, economic viability and social acceptance of systems and services.

3. **Quantum technology and cryptology for information security.** Quantum technology rests on the use of photons rather than electrons, as in the silicon industry. It offers a range of opportunities for future cryptology, for example through the possibility to generate truly random values or to distribute bits of information with absolute security. But at the same time it is well known that quantum technology is a threat to current cryptographic methods, since a quantum computer may break current asymmetric cryptographic techniques in a few seconds 10-15 years from now. We should lay the foundations for new algorithms to effectively resist code-breaking attempts by quantum computers.

This presentation was followed by a lively discussion during which many aspects where clarified. This demonstrated a strong support from the participants for the challenges described in the report, although they way quantum encryption can be combined with mainstream information technology is an open question and a concern to some participants. One new issue that was brought forward during the discussion was the aspect of privacy at the level of the entire network, which is evident in the amount of private information that can be collected about individual users who take many actions (searching, email and other forms of communication, purchasing goods or services), each of which leaves behind small amounts of information that when put together may constitute a frightening loss of privacy.

7 List of Participants

Marcos Alemán García
Atos Origin
mjalemangarcia@yahoo.es

Prof. Georg Carle
University of Tübingen
carle@uni-tuebingen.de

Prof Bruno Crispo
University of Trento
crispo@dit.unitn.it

Dr Alain Esterle
ENISA
alain.esterle@enisa.eu.int

Marit Hansen
Indepedent Centre for Privacy Protection
Schleswig-Holstein
ld10@datenschutzzentrum.de

Wide Hogenhout
European Commission
wide.hogenhout@ec.europa.eu

Prof Iordanis Koutsopoulos
University of Thessaly
jordan@uth.gr

Dr Henning Makholm
IMM, Technical University of Denmark
henning@makholm.net

Dr Oscar Manso
Universitat Politècnica de Catalunya
o.manso@ac.upc.edu

Philippe Massonet
CETIC
phm@cetic.be

Björn Melen
Ericsson
bjorn.melen@ericsson.com

Prof Refik Molva
EURECOM
refik.molva@eurecom.fr

Prof. Jörn Altmann
International University
jorn.altmann@acm.org

Miguel Colomer
EsCERT, Universitat Politècnica de Catalunya
mcolomer@escert.upc.edu

Boris Dragovic
Create-Net
Boris.Dragovic@create-net.it

Prof Erol Gelenbe
Imperial College

Dr Paul Hodgson
British Telecom
paul.w.hodgson@bt.com

Sebastian Höhn
Albert-Ludwigs University Freiburg
hoehn@iig.uni-freiburg.de

Prof Odysseas Koufopavlou
University of Patras
odysseas@ee.upatras.gr

Dr Jesus Luna
Universitat Politècnica de Catalunya
jluna@ac.upc.edu

Prof Antonio Maña
University of Málaga
amg@lcc.uma.es

Dr Fabio Martinelli
IIT-CNR
Fabio.Martinelli@iit.cnr.it

Prof Manel Medina
Universitat Politècnica de Catalunya
medina@ac.upc.edu

Dr Pietro Michiardi
Institut Eurecom
Pietro.Michiardi@eurecom.fr

Dr Melek Önen
EURECOM INSTITUTE

Prof Giuseppe Persiano
Universita' di Salerno
giuper@dia.unisa.it

Dr German Puebla
Universidad Politecnica de Madrid
german@fi.upm.es

Helena Rifà-Pous
Safelayer Secure Communications, S.A.
hrifa@safelayer.com

Daniel Schreckling
University of Hamburg
schreckling@informatik.uni-hamburg.de

Prof Josep Sole Pareta
Universitat Politècnica de Catalunya
pareta@ac.upc.edu

Dr Salvatore Spadaro
Universitat Politècnica de Catalunya
sspadaro@ac.upc.edu

Werner Streitberger
University of Bayreuth
werner.streitberger@uni-bayreuth.de

Bart van Caenegem
European Comission
Bart.VAN-CAENEGEM@cec.eu.int

Dr Dirk Westhoff
NEC Europe Ltd.
dirk.westhoff@netlab.nec.de

melek.onen@eurecom.fr

Prof Andreas Pfitzmann
TU Dresden, Institute for System Architecture
pfitza@inf.tu-dresden.de

Prof Kai Rannenberg
Goethe University Frankfurt
kair@m-lehrstuhl.de

Maarten Rits
SAP Research
maarten.rits@sap.com

Fabrizio Sestini
European Comission
Fabrizio.Sestini@cec.eu.int

Ignacio Soler
Atos Origin
ignacio.solerjubert@atosorigin.com

Prof James Sterbenz
Lancaster University

Esa Turtiainen
Ericsson
esa.turtiainen@ericsson.com

Dr Christos Verikoukis
CTTC
cveri@cttc.es

Annamaria Woerndl
Austrian Research Promotion Agency
annamaria.woerndl@ffg.at

8 Programme

9:00 - 9:15
Wide Hogenhout / Bart Van Caenegem - introduction

9:15 - 11:15
Session "Evolutionary and Adaptive Security"
HAGGLE - Refik Molva, Melek Onen, Abdullatif Shikfa (Institut Eurecom)
"Self-organizing Trust establishment and Cooperation Enforcement"
CASCADAS - E. Gelenbe (Imperial College), "Autonomic Self-Defence"
BIONETS - Daniel Schreckling (Univ. of Hamburg) "Securing BIONETS: How can Security Infrastructures Match Autonomically Evolving Networks and Services?"
SERENITY - Antonio Maña (University of Malaga), "Dynamic Security in Ambient Intelligence Scenarios"
MOBIUS - Maarten Rits (SAP Research) "Adaptive Security with Aspect Oriented Programming"
Moderator: Boris Dragovic (Create-Net)
Rapporteur: Daniel Schreckling (Univ. of Hamburg)

11:30 - 13:30
Session "Privacy and trust for dynamic coalitions"
CASCADAS - Pietro Michiardi (Institut Eurecom), "Reputation mechanisms for manet and P2P over manet"
SENSORIA - Fabio Massacci (Univ. of Trento), "Privacy and Trust requirement engineering"
FIDIS(1) - Andreas Pfitzmann (Technische Universitat Dresden), Marit Hansen (ICPP), "Identity management in social networks"
FIDIS(2) - Sebastian Höhn (Albert-Ludwig-University Freiburg), "Usable Security in Dynamic Systems".
UBISEC & SENS - Dirk Westhoff (NEC Europe), "UbiSec&Sens - Security Challenges and Solutions for Wireless Sensor Networks"
BIONETS - Bruno Crispo (Univ. of Trento) "Trust and Reputation Issues in BIONETS"
Moderator: Kai Rannenberg (Goethe University Frankfurt)
Rapporteur: Fabio Massacci (Univ. of Trento)

13:30 - 14:30
Lunch

14:30 - 16:45
Session "Distributed security"
SENSORIA - Fabio Martinelli (IIT - CNR) "Security and trust issues for services"
MOBIUS - German Puebla (Universidad Politecnica de Madrid) "Possible Future Scenarios for Security based on PCC"
S3MS - Fabio Massacci (Univ. of Trento), "Security by contract for mobile code"
ANA - James Sterbenz (Univ. of Lancaster / Univ. of Kansas), title to be confirmed
AEOLUS - Giuseppe Persiano (Univ. of Salerno), "Security preserving composition of protocols"
Moderator: Fabio Martinelli (IIT - CNR)
Rapporteur: Henning Makholm (IMM - DTU)

17:00 - 17:45
Wide Hogenhout - Presentation about the results of the Beyond the Horizon activity from FET and discussion