







# EU-US Summit
# Cyber Trust: System Dependability & Security

## Workshop 1
## Dublin 15-16 November 2006

Author: Zeta Dooly, Waterford Institute of Technology, Ireland

Contributors: Willie Donnelly (WIT), Jim Clarke, (WIT), Thomas Skordas (EC), Molly Tracy (Univ. of Illinois), Workshop Organising committee, Workshop Chairs and Rapporteurs, Workshop attendees.

**Table of contents**

**Table of Contents - Position Paper – Appendix A**

**Panel F**

## Executive Summary

The workshop of the EU/US Summit Series on "Cyber Trust: System Dependability & Security" was held in Dublin, Ireland on November 15th and 16th, 2006. It was attended by 60 delegates from the EU and the US, along with representatives from Canada, Australia and Japan. This event was co-organised and hosted by Waterford Institute of Technology (WIT), the project co-ordinator of the IST-FP6 Co-ordination Action SecurIST, and also co-organised by the US National Science Foundation (NSF), Department of Homeland Security (DHS), University of Illinois, and the European Commission, Directorate General Information Society and Media, Unit F5 "Security".

The aim of this workshop, and a planned follow-up workshop to be held in Illinois in April/May 2007, was to gain a shared understanding of critical issues, identifying promising dependability and security research directions, and also to foster collaboration between EU and US research teams. The organising committee were satisfied at the outcome of this workshop that it had indeed met the objectives and that further collaborative work would ensure to facilitate international co-operation in Trust, Security and Dependability (TSD) research.

The organising committee of the workshop developed the programme around the following themes within Trust, Security and Dependability (TSD) of future networked systems: Architecture and design issues, scalability and context awareness, security and privacy in dynamic wireless networks, TSD predictive evaluation and assessment approaches and future test-beds. The workshop discussions led to identification of a number of research challenges, scoping research priorities and joint actions to address them.

The first workshop themes and their conclusions were the following:

*(1)* ***Architecture and design issues for TSD of Future Networked Systems.*** Future emerging networked ICT systems will be large-scale, complex mixed mode environments consisting of diverse computing, communication & storage capacities. They will be based on the model of service-centric computing, systems of embedded systems and a mix of classical computers and embedded systems on the Internet. The discussion focussed around the new TSD attributes that such future ICT systems should be endowed with. These include trustworthiness and resilience, protocols, languages, metrics, internet routing paradigms, security provision technologies (cryptology, trusted functionality, multi-modal biometry, etc.), adaptive detection, diagnosis, run-time response mechanisms and stochastic security in core/access networks from an end-to-end perspective. For these new systems, there is a need to specify not only the underlying service semantics but also the TSD semantics and metrics for designing resilient architectures and secure network protocols and for detecting and measuring any anomalous behaviour.

*(2)* ***Scalability and context-awareness for TSD of Future Networked Systems.*** Discussions focused on multi-layered, scalable and context-aware approaches to make future networked systems secure and dependable. The main conclusions focussed on the need to extend scalability from all perspectives (hardware, software and systems) through better, realistic

abstractions and by focusing on three phases of a system's lifecycle: (a) capturing network functionality, system performance and end-users requirements (b) System design, and (c) System evaluation and testing. Other discussions focussed on: development of a formal authorization engineering framework to increase the authorization capability limits required in order to support multiple administrative domains; automated fault detection and remediation techniques for application on a massive and growing scale; and, support of health management of autonomic system-of-systems approaches that enable automated fault detection and remediation on a massive scale.

(3) ***Security and privacy in dynamic wireless networks*** of evolving systems composed of *ad hoc* coalitions of large numbers of sensors and devices for new personalized services. The main conclusions focussed on addressing the lack of a security infrastructure, of threat models and of adequate security evaluation techniques for dynamic wireless networks. The main research directions identified to address this challenge were: testing methods and threat models; security infrastructure akin to tethered networks; federation of security policies and mechanisms across multiple domains; adaptive systems based on context; trust management while giving users more control over choosing risk levels and adaptable context; and, usability of security systems, especially in complex heterogeneous sensor systems.

(4) ***Modelling, simulation, predictive evaluation, assurance cases for evaluating the TSD of networked systems.*** The main issues addressed under this theme were verification and evaluation frameworks related to (possibly) Internet-scale applications and to particular networks and networked systems. There is a need to consider the wider socio-technical aspects and interdependencies as well as their semantic learning and understanding dimensions. There is also a need to use assurance cases and claim semantics from and for different stakeholders' viewpoints in order to communicate assumptions and agree on system security. When developing the above further, scenario building and use case generation would enhance understanding and inclusion of test data. There is a need here to develop and use standard metrics for incremental security improvements and probabilistic approaches for radical security improvements and for reducing stakeholders' interdependencies.

(5) ***Monitoring, operational assessment, auditing for evaluating the TSD of Networked Systems.*** Discussions focused on dynamic and online methods of analysis and evaluation and on real time assessment frameworks, including attacks observed, observation mechanisms, audits, measurement and decision making tools, etc. It is imperative to start now with the challenges associated with metrics, measurements and analysis, even with limited systems and goals, to gain a better understanding for threat characterization, prediction, observation, instrumentation and data collection. On-line measurements are needed to control and adapt, in particular, to put in place network information sharing techniques at all levels (including attacks observed, keystrokes of users, network traffic capture in an anonymous fashion and others). There must also be put in place more incentives for the provision and sharing of data, which is needed to ensure sufficient context that would permit replication through

experiments.

*(6)* Establishment of ***interconnected and/or common test-beds.*** Issues discussed include: opportunities for interconnecting existing experimental facilities and building joint benchmarks; test scenarios and interconnected test-beds for supporting the testing and evaluation of new dependability and security architectures; and, technologies, protocols, and privacy protection mechanisms, together with support towards global standards. Examples of identified potential shared test-beds include a test-bed for software and services to allow experimentation at the application and services level or a test-bed for dynamic wireless and sensor networks. The first would open up valuable opportunities for innovative Small and Medium sized Enterprises and Academics to venture into service-oriented solutions. For wireless and sensor network test beds, there are some standalone test beds already available but the issue that must be explored is to how to federate them taking into account cross testing, mobility aspects and security policies as users move in and out of different environments.

The web site for the Workshop is www.securitytaskforce.eu.

# 1 Introduction: Workshop Challenges, Objectives and Format

There is significant investment and activity on both sides of the Atlantic, both academic and industrial, aimed at creating complex networked systems that can be justifiably trusted to perform as expected, i.e. to be resilient, namely to remain dependable and secure, in spite of problems that they may experience, and the requirement to be adaptable and to be able to evolve in response to changed requirements and environments. Such problems can be caused by malicious behaviour on the part of adversaries, or by accidental events, caused, for example, by human mistakes, residual design errors or hardware malfunctions. Failures of such systems can have far-reaching, often international implications.

The aim of this workshop and a planned subsequent workshop to be held in Illinois on 26-27 April 2007, is to gain a shared understanding of priority critical issues and promising dependability and security research directions, and to foster collaboration between EU and US research teams. Application areas and characteristics of interest include critical information infrastructures, e-commerce, e-government and e-voting, the future of the Internet, scalable and self-regulating systems, and ubiquitous and pervasive computing.

In particular, the main objectives were to explore:

- common research themes and challenges in Europe and the US within the Trust, Security and Dependability TSD environment
- how to facilitate research initiatives that are common to both the EU and the US.

In the six months prior to the first workshop, an international organizing committee was formed and tasked with exploring how Europe and the US could best achieve these objectives. **The organising committee of the workshop** comprised of William Donnelly (Waterford IT), Zeta Dooly (Waterford IT), Karl Levitt (US National Science Foundation), Douglas Maughan (US Department of Homeland Security), Jacques Bus (European Commission), Thomas Skordas (European Commission), Molly Tracy (Univ. of Illinois), William Sanders (Univ. of Illinois) Michel Riguidel (ENST) and Brian Randell (Univ. of Newcastle).

The organising committee has defined the research areas of focus of the workshop as follows:

- **Dependability & Security of Future Networked Systems – architecture and design issues**: including new protocols, adaptive detection, diagnosis, and response mechanisms, and topics such as stochastic security, in core and access networks (both wired & wireless), from an end-to-end perspective.
- **Dependability & Security of Future Networked Systems – scalability and context-awareness**: multi-layered and scalable approaches to make networked systems secure and dependable.
- **Security & Privacy in Dynamic Wireless Networks**: of evolving systems composed of ad hoc coalitions of large numbers of sensors and devices for new personalized services.
- **Evaluating the Dependability & Security of Networked Systems – modeling,**

**simulation, predictive evaluation, assurance cases:** related to possibly Internet-scale applications, and to particular networks and networked systems.

- **Evaluating the Dependability & Security of Networked Systems – monitoring, operational assessment, auditing**: dynamic, online, analysis and evaluation methods, real time assessment (observation mechanisms, audits, measurement and decision making tools, etc.).
- **Future Test Beds**: joint benchmarks, test scenarios and interconnected test beds for assessing new dependability and security architectures, technologies, protocols, privacy protection mechanisms, etc., together with support towards global standards.

The workshop was organised around six panel sessions each addressing one of the above research areas. In his introductory remarks, Willie Donnelly presented the focus of the Panel sessions, the format of the sessions and the next steps of the EU-US research co-operation:

- **Focus of the Panel Sessions**: The sessions were designed to focus and gain consensus on difficult but important research questions for joint collaboration in the EU-US research communities working in the TSD area; describing the impact of solving or failure to solve these challenges; identifying the potential mutual research directions and approaches; and, focussing on potential common research resources, e.g., test-beds, common software platforms, data sets, design and validation tools, etc.
- **Format of the sessions**: Each session would be structured as follows: a number of panel speakers providing position statements on the session topics; these would be followed by a general discussion centered on the identification of priority research areas for a joint EU/US and possibly wider international collaboration in the session topics and a methodology to kick-off the process. The Chairs and Rapporteurs of each session would prepare summaries that are then further discussed in a concluding general session, and used in a workshop report.
- **Next steps of the EU-US research co-operation**: The second workshop goals are to incorporate wider views and inputs from external stakeholders including industry, academic, government and users and the establishment of a core steering group (SG), whose objectives are to act as a catalyst in focussing the output from the workshops in defining and facilitating joint activities for the TSD research communities, whilst remaining lightweight in terms of bureaucracy and focus on alignment of potential funding mechanisms.

The remainder of this report is organised as follows:

Section 2 presents the final discussion session presentations and the reports from the chairs of the sessions.

Section 3 outlines the final conclusions and input to the next workshop in the series.

Appendix A contains all position papers.

The workshop agenda is outlined in Appendix B and the full list of participants is outlined in Appendix C.

All the workshop presentations are available on www.securitytaskforce.eu

# 2 Final Discussion Session Presentations and Reports from Chairs

## 2.1 Panel A: Dependability & Security of Future Networked Systems – architecture and design issues – Chairs report

**Co-Chairs**: David Du (NSF, US) and Paulo Verissimo (Univ. of Lisboa, Portugal)
**Rapporteur**: Jim Clarke (Waterford Institute of Technology, Ireland)

**Speakers**:
Ravishankar Iyer (Univ. of Illinois, US)
Michel Riguidel (ENST, France)
Felix Wu (UC Davis, US)
Bart Preneel (Katholieke Univertsiteit Leuven, Belgium)
Yair Amir (JHU, US)
Neeraj Suri (TU Darmstadt, Germany)

### Main topics debated

- Need for new attributes for trust, security and dependability (TSD) including design for trustworthiness (defined as security and dependability) and resilient designs as the key to reconciling uncertainty with predictability brought about by new architectures, which are characterised by large-scale, complex and networked, pervasive, mixed mode, systems of embedded systems and a mix of classical computers and embedded systems on the Internet.
- Application-aware detection and application-centric trusted computing platforms.
- Security in the digital convergence scene: moving Sec&Dep to the application layer; protecting while evolving from small to large scale complex systems, promoting security of virtual entities.
- Future networked systems will be large-scale and complex systems-of-embedded-systems: future Internet as a blend of conspicuous (PCs, PDAs, etc.) and inconspicuous devices (gadgets, embedded controllers, etc.): potentially harmful if not controlled adequately.
- New paradigms for scalability and security of Internet routing: a fundamental change is required to the way we handle routing; need for real time traffic analysis and protection schemes.

- Crypto challenges: protocols and algorithms securing long-term security (50-100 years); ultra-low footprint architectures; multi-party computations; threshold crypto; location/ context dependent decryption.
- Dealing with the insider threat in large-scale systems: taking into account network infrastructure and distributed information systems.
- Fluid networked environments or mixed-mode environments.
- Long-term data security: taking into account data continuity and availability requirements; long-term key management.

## Prominent challenges

As difficult as it may be to single out challenges from the stimulating scenario just described, one or two research issues come out as forward-looking and general enough to also include some other more concrete challenges. The first is system-oriented and broad in scope. The second is methodological, more focused in scope, but with very large impact potential in a number of computer science and engineering fields. Both are important with regard to TSD, and capable of motivating a large set of researchers in both sides of the ocean.

### (1) Resilient design of new generation networked systems

These systems will be large-scale, complex and networked, mixed mode, a mix of classical computers and embedded systems on the Internet.

Addressing this challenge will put us in the way of achieving trust in the forthcoming ambient intelligence world (the pervasive and ubiquitous world of conspicuous and “disappearing” computers that will surround us).

On the other hand, if we are not prepared, such Internet-born embedded+ pervasive+ ubiquitous systems, if not working properly, can do more damage than good. In such a new reality, doing Control-Alt-Del may not be enough for mending something that went wrong, because behind an IP address may be an embedded system or gadget capable of damaging something if not properly controlled. If an ICT-based society will not be able to provide trustable services, services that are trusted because justifiably they rely on trustworthy components and infrastructure, then, such services, which nevertheless will be deployed due to market pressure: (i) will be perceived with suspicion by large mass of users; (ii) information will be managed by a restricted group of "experts", increasing info-exclusion; (iii) as it is managed, it may very well be mismanaged, becoming a big opportunity for cyber-crime, e-frauds, cyber terrorism and sabotage.

The key to this challenge lies on investigating new distributed paradigms and mechanisms fusing the best from autonomous, peer-to-peer, pervasive, and embedded systems. There will be many more computing devices in embedded, sometimes function-specific, inconspicuous devices, pervading the societal tissue than there will be 'computers' as we know them today. Most, if not practically all of them, will be networked, interconnected in some form. One of the things that make it a grand challenge is exactly the difficulty of understanding and tackling all the new problems brought by this confluence, especially in what concerns TSD. A hacker could usually

deface a web page, bring a server down, modify a disk's content, but in the future he/she may remotely wipe out the firmware from car controllers, bring down a power station, or set fire to a home microwave.

As for the research infrastructure, if we consider as a starting point a typical Internet-based infrastructure for distributed systems research, then additional means may include: sensor networks; RFID systems; wearable computers; control area networks such as field-buses; emulators of real control or cooperative systems made of embedded devices in several areas (e.g. robots; cars or planes; electrical power or process control plants; home systems). Example would-be prototypes for systems research in this area might be: cooperative enhanced human and/or autonomous robot teams; cooperative cars (e.g. platoons); assisted road-side and Internet-born navigation and driving; Internet-born car maintenance, diagnostic upload and firmware download; remote operation of critical infrastructures through Internet-born secure links; Internet-enabled safe and secure home control and information sharing.

We can anticipate that the challenge will be met when we see prototypes of these mixed-mode systems being deployed over the Internet, in a performing, user-friendly, safe, dependable and secure way. Along this road, relevant milestones may consist of proof-of-concept prototypes showing partial results in the research topics enumerated above.

**(2) Methods and paradigms for reconciling uncertainty with predictability**

Complexity of systems is growing out of control and more and more "always-on" complex systems are being deployed or planned, especially by governments. This inevitably creates growing interdependencies between systems, services and humans. This situation opens doors to new forms of threats or amplifies existing ones, specially those threats that take advantage from complexity and uncertainty. There is no current definite solution to achieving predictable behaviour in face of these threats, because most design approaches and paradigms for TSD rely on reasonably complete aprioristic assumptions about the problem and the environment at stake. Since this problem cannot be removed and will be amplified in the future, then, like in nature, solving it requires adaptation and one good enough for survivability.

This problem is important and hard enough that a coherent research effort should be promoted, on innovative approaches promoting trustworthy adaptability: proactive-reactive design under uncertainty; adapting functional and non-functional properties while providing guarantees on adaptation result; autonomous, decentralised, self-stabilising system algorithmics; runtime stochastic modelling; trustworthy monitoring, update and self-healing.

Solving this challenge will contribute towards the widespread deployment of very dynamic and evolvable systems that can be trusted in spite of their complexity. In other words, where 'complexity' is not an excuse for 'undependability' or 'insecurity' as is the case today. Moreover, this challenge has transversal impact in practically all sectors depending on ICT, and if not addressed, there is the risk of a serious fallback on plans to attain an advanced, widely connected information society. Keeping with traditional design, deployment and operation methods is likely to cause serious hazards in the operation of systems that alternate between states of fossilised dependability/security and periods of undependability/ insecurity during and after system changes.

## 2.2 Panel B: Scalability and Context Awareness – Chairs report

**Co-Chairs**: John Knight (Univ. of Virginia, US) and Brian Randell (Univ. of Newcastle upon Tyne, UK)
**Rapporteur**: Jim Clarke (Waterford Institute of Technology, Ireland)

**Speakers**:
Jean-Claude Laprie (LAAS, France)
Nick Weaver (Univ. of Berkeley, US)
Christof Fetzer (TU Dresden, Germany)
George Kesidis (Penn. State Univ., US)
Gerard LeLann (INRIA, France)
Ming-Yuh Huang (Boeing, US)

### Introduction

The size and complexity of networked information systems have increased dramatically in recent years, and there is every indication that this trend will continue. These systems are supporting more and more, and larger and larger applications, many of which are critical in nature. Criticality might concern: (1) a requirement for service *availability* in which users can expect to receive service whenever it is needed; (2) a requirement for *security* in one or more of several forms including confidentially and data integrity; (3) a requirement for *reliability* in which trustworthy service has to be provided for some specified time; (4) a requirement for *safety*, i.e. an absence of failures that could lead to loss of life, in situations where the system is controlling a set of potentially-dangerous physical devices (e.g., a transportation system or an automated manufacturing system); and/or (5) a requirement for *survivability* in which the networked system provides acceptable service, albeit in various reduced or different specified forms, even if it suffers physical damage or security attacks.

Society's dependence on applications implemented by large, networked information systems is increasing faster than the systems themselves are being built. This is because new applications are being created and existing ones expanded or enhanced on existing as well as new network equipment.

These circumstances present significant challenges for which there are no known solutions and which must be addressed by the research community. In particular, the research community needs to concentrate on the issues of *scalability* and *context awareness*. The meanings of these terms are as follows:

**Scalability:**

As a network's size parameters are increased by significant amounts, a network that is scalable will continue to provide acceptable service. (Size parameters will include number of nodes, computing power and storage capacity per node, or communications bandwidth between nodes.)

**Context awareness:**

The networked systems that we build have to operate in a context. This context includes all of the environmental factors that might impact the network, operator interaction, data sources and data sinks, and a lengthy list of adversaries. A network that is context aware is capable of accommodating anticipated changes in its operating context.

With these definitions in mind and with the demand for coping with these challenges, the panel discussion led to the detailed observations that are listed in the next two sections.

## Prominent Challenges

The challenges that were identified as requiring attention were:

- Scaling up of dependability and security to address the dynamics of environmental and operational issues.
- Technologies for ensuring the resilience of large systems and for maintaining these properties as systems scale up.
- General techniques and technologies for building and scaling networked information systems. In particular, four areas needing development were identified:
  - Techniques for exploiting the very significant parallelism that is becoming available as additional processing, storage, and communications facilities are deployed.
  - Techniques for dealing with the growing number of significant file and application servers and the protocols that they use.
  - Techniques for supporting novel applications such as Voice over IP (VoIP) traffic.
  - Techniques for dealing with the advent of significant broadcast traffic.
- General techniques for dealing with the resources being used in developing networked information systems and their role in scalable dependable systems:
  - Techniques for the use of commodity, less reliable, hardware/software to build mixed mode dependable and secure systems.
  - Basic computer hardware is increasing in integration level (more devices per unit area) and this is being accompanied by an increase in defects during operation. These are usually intermittent failures, but they present a major difficulty for system designers. Such failures can cause difficulties for an application or for system software, and these failures will not be dealt with by hardware changes. Techniques must be developed for tolerating such failures, something that will have to be done in software.
  - Approaches to software containment need to be developed to allow system architects to partition systems so that the effects of software defects can be limited.

## Potential Areas of Cooperative Research

The opportunity to develop a cooperative program in which researchers funded by the EU and the USA tackle these challenges is very appealing to the research community. The topics identified by the panel as potential areas to pursue for cooperative research were:

- Development of large-scale routing strategies for the next generation Internet:
  - The Border Gateway Protocol is a widely distributed and widely used protocol, and it is important to determine its limitations in the face of network scale and context changes.
  - Determination of the extent to which other current techniques will scale with the network.
  - Development of superior network communication, management and security approaches.
  - Development of better, more powerful, realistic abstractions for scalability.
- Research concentrating on the three phases of the development lifecycle:
  - The necessary network functionality and performance requirements need to be captured, including input from involved citizens.
  - Technology needs to be developed to better support system design and evaluation.
  - Technology needs to be developed to better support system testing. It should be noted that this is a very difficult task. Assessment by testing is important yet testing a large network is almost impossible since development of a realistic laboratory model is infeasible.
- Techniques need to be developed to permit systems-level analysis in order to provide an analytic capability of network characteristics.
- Technology is required to support formal authorization engineering and a federated authorization framework. A general lack of authorization capabilities limits the scope of systems that can be considered because typical systems cannot easily support multiple administrative domains.
- Support for health management of autonomic system-of-systems is required because: (a) systems of systems are a realistic form of modern, large networked systems; and (b) it will be essentially impossible to manage large-scale systems manually in the future (it is virtually impossible now). Automated techniques will be/are needed to effect this management.
- Automated fault detection and remediation techniques for application on a massive and growing scale are needed.
- New security paradigms are required that will provide strong guarantees of security properties in large networked information systems. Such systems will be used to store more and more sensitive information, and protecting that information and operations upon it is crucial to the future effective deployment of such systems.

- The possibility of security in context needs to be examined:
  - Significantly different directions are possible including approaching the problem via the network infrastructure, the supported applications, the supported services, or via human processes.
  - Enabling individual citizens to have control and awareness of dependability and security issues might result in significantly enhanced trust in networked systems.
- Possible EU/US areas of cooperation:
  - All dimensions of scalability. The community needs to work together on the scalability agenda.
  - The community needs to try to capitalize on the complementary research and personal cultures of EU and US. For example, the two communities have significantly different approaches to privacy in health care.
- Approaches need to be developed to facilitate and encourage all stakeholders to discuss jointly why or how a given system failure has happened. Specifically, there need to be approaches that:
  - allow the determination of the extent to which the fault was human, some aspect of the system, or a component;
  - enable a significant reduction in the probability that similar incidents will occur again;
  - allow cross disciplinary investigation of incidents, including sociologists and other related disciplines;
  - involve setting up a moderated global forum including all communities for sharing information of all aspects of network dependability and security.
- There is a new world of the Internet, boundaries are no longer clear and anyone can set up a network. To begin to deal with this situation, we need to build intelligence into networks so that rogue/malicious elements can be recognized.
- Citizens must have this new Internet at their disposal with guarantees of privacy and confidentiality.
- A multi-dimensional technique needs to be developed for dealing with intellectual problems and developing solutions without making the systems or infrastructure easier or more valuable for adversaries.

## Conclusion

The challenges identified by the panel are vital. If they are not dealt with in cost-effective and technically acceptable ways, they might limit our ability to exploit the technology of networked information systems for the benefit of society. If we fail to recognize this situation, we might

deploy systems that are not engineered with the properties that we need, and this could lead to failures with societal impacts far more serious than the absence of the networked systems.

The above is the bad news. The good news is that we have identified the challenges, and we have identified much of the detail of the problems that have to be addressed. The opportunity exists to attack these problems with carefully conducted research programs. This research would profit considerably from the potential synergy of cooperation between researchers in the United States and the European Union.

## 2.3 Panel C - Security & Privacy in Dynamic Wireless Networks – Chairs report

**Co-Chairs**: Gene Tsudik (UC Irvine, US) and Roberto Baldoni (Univ. of Roma, Italy)
**Rapporteur**: Mike Bailey (Univ. of Michigan, US)

**Speakers**:
David Kotz (Dartmouth, US)
Reijo Savola (VTT, Finland)
Joe Evans (KU, US)
Stephan Engberg (Priway, Denmark)
Wenke Lee (Georgia Inst. of Technology, Atlanta, US)
Paddy Nixon (UCD, Ireland)

### Introduction

Dynamic wireless networks include multi-hop ad hoc networks (MANETs), single-hop residential or hot-spot-type wireless networks, as well as sensor networks and other wireless devices. The main conclusions of the discussions held under this panel session are as follows:

- Dynamic wireless networks face many security and privacy challenges. There is a lack of a security infrastructure, threat models and adequate security evaluation techniques. There is an evident risk of loss of privacy when using such networks. Networks need to effectively perform basic operations (e.g., routing, membership) in adaptive contexts. There is also a need to consider risk management approaches as well as address control, configuration, and usage of ubiquitous devices.
- The panel suggested to "break out of the loop" of working on small problems and within partial models. The main issues concern the protection of sensitive information, the network topology, the biometric information, the identity, and, the ability to manage the risk associated with their loss. The core challenge is about the creation of secure, efficient and usable systems. The main research directions agreed to address this challenge were the following: testing methods and threat models; security infrastructure akin to tethered networks; adaptive systems based on context; flexible trust management (giving users control over balancing risk levels and security costs); and, usability of security/privacy measures.

### Prominent Challenges

#### (1) Network membership

A dynamic wireless network, by definition, needs to cope with a dynamic topology and dynamic membership. Without secure membership techniques (i.e., knowing who is in the network), secure communication is difficult. Since such a network typically lacks any central point of trust, membership decisions must be made in a decentralized (distributed), yet secure, manner.

However, admitting and evicting members in a distributed fashion is a notoriously difficult problem due mostly to the communication and computation complexity of the cryptographic protocols required for these tasks.

- A number of methods have been proposed to-date. However, some are insecure, others are too costly and the remainder have never been tested (experimented with) in realistic network settings.
- Solutions to the network membership problem can benefit greatly from the availability of a comprehensive large-scale wireless network test-bed that can be used for experiments with new membership control techniques. Realistic scenarios can be played out to test the viability (delay, fault-tolerance, etc.) of various techniques.

**(2) Privacy in wireless networks**

Current wireless networks suffer from many privacy-related problems. One-hop wireless (such as cellular and 802.11-type networks) allows tracking of nodes/devices. More sophisticated multi-hop dynamic wireless networks (MANETs) leak information about the topology and node movements (typically, through routing protocols).

- A few techniques have been proposed for privacy in one-hop wireless networks, however, they often make unrealistic assumptions. For the multi-hop case, little (if any) work has been done to-date. The main challenge lies in designing privacy-preserving routing protocols. This is, in and of itself, a departure from traditional routing protocols that view dissemination of topology/location information as their main goal.
- Once again, wireless network test-beds can be of tremendous help in assessing effectiveness of methods and techniques for privacy preservation in wireless routing and communication, in general.

**(3) Usability of Security/Privacy Techniques**

There exists a big rift between techniques proposed by the research community (typically, cryptography and security experts) and those acceptable by the public at large. Many cryptographic techniques allegedly oriented towards human use are in fact unusable by (or poorly suited for) the majority of the user population. This applies not only to obvious applications, such as secure pairing of personal devices (phones, PDA-s, headsets, etc.), but also to usability of privacy techniques for RFID tags.

- Although numerous techniques have been proposed, usability community has been doing the work separately from the cryptography/security research community. Consequently, the former has sometimes produced insecure techniques, while the latter – poorly usable ones.
- Experimental test-beds and standardized platforms are needed for evaluation usability factors of security and privacy methods and techniques. A set of common usability scenarios (and metrics) is also necessary.

**(4) Wireless Intrusion Detection**

Wireless devices and networks are subject to attack by worms and botnets, unless network administrators are diligent about management. To assist operators in this task, we need to

construct security infrastructures that help support this diligent monitoring and evaluation of network risk. Fortunately, many of the lessons we learned in the wired networks can provide a good starting point for this type of infrastructure in wireless networks (e.g., wireless honeypots).

## Recommendations

- Facilitate inter-disciplinary research blending elements of computer science, sociology, psychology, law, ethics, and economics in order to better understand tradeoffs in security and privacy of dynamic wireless networks. Likewise, encourage interdisciplinary research into usability of security and privacy techniques for wireless devices and networks. (This especially applies to ubiquitous personal devices, such as cell-phones, as well as RFID tags.) Furthermore, provide a set of joint US/EU test-beds that model several common types of dynamic wireless networks, such as: 1) a multi-hop configurable MANET, 2) a large-scale wireless sensor network (WSN), 3) a retail- or warehouse-style RFID setting, and 4) a home (or small office) one-hop wireless network environment.
- Define threat models and security requirements that guide research security and privacy in wireless networks, so that there exists clear agenda for progress.
- Provide additional tools, models and test-beds for evaluating new techniques in order to foster measured and demonstrable progress and better disseminate research results.

## 2.4 Panel D: Modelling, simulation, predictive evaluation, assurance cases for evaluating the TSD of networked systems – Chairs report

**Co-Chairs**: William Sanders (Univ. of Illinois, Urbana-Champaign, US) and Dieter Gollmann, (TU Hamburg-Harburg, Germany)
**Rapporteur**: Stephan Engberg (Priway, DK)

**Speakers**:

John Rushby (SRI International, California, US)
Bev Littlewood (City Univ., UK)
John McHugh (DAL, Canada)
Aad van Moorsel (Newcastle Univ., UK)
O. Sami Saydjari (Cyber Defense Agency, LLC, US)
Robin Bloomfield (City Univ. + Adelard, UK)

The objective of this panel session was to investigate how to create a scientific foundation, methods, and tools for quantitative assessment of metrics of combined dependability & security (i.e., trust) that can be applied to enterprise-level and Internet-scale systems. The main issues addressed under this theme were verification and evaluation frameworks related to (possibly) Internet-scale applications and to particular networks and networked systems.

### Prominent Challenges

The main challenges are as follows:

- That a gradual approach to semantic learning and understanding is required to avoid taking in of unproductive direction;
- Consideration of wider socio-technical aspects and interdependencies is required;
- Make subjective assumptions, reasoning and evidence explicit and visible as part of a probabilistic approach;
- Use economic theory and security valuation to make multi-objective trade-off decisions;
- A requirement was identified to use assurance cases and claim semantics from and for different stakeholders viewpoint to communicate assumptions and agree on system security levels. When developing these further building scenarios and generating use cases may assist with understanding and populate of test data;
- Need to enable corporate and government determining accurate, quantifiable, security & dependability metrics, establishing quantifiable, repeatable metrics, monitoring and updating metrics where required. The signals in the metrics require measurements to be very detailed and specific – agreement on metrics is not enough;

– Using the MILS approach of certified components building more complex systems may assist further development of this area;

– Models are required to quantify and analyse the business case and adversary attacks probability, these will assist development of solutions and also need to be able to adapt to changing environments.

Related research priorities that were highlighted include: Develop methods to deal with the intelligent and innovative attacker: in particular, how to deal with the unknown and perhaps even the "unknowable unknown". Investigate mechanisms to incorporate the different valuation systems of different stakeholders including both normal players and a range of adversaries, as well as methods to build trustworthiness and metrics into vendor products. Devise metrics to make intelligent engineering decisions based on probabilistic and Bayesian analysis. There is a need here to develop and use standard metrics for incremental security improvements and probabilistic approaches for radical security improvements and for reducing stakeholders' interdependencies.

During the discussion session, the following points were raised:

- Existing methods focus narrowly and exclusively on a single aspect of security;
- Most technical metrics are not quantitative but focus on process;
- Models and metrics on multiple levels must be integrated;
- There is a need to determine methods for estimating metrics;
- Supporting tools are required to ensure that the process of collecting and analysing metrics is timely and practitioners can analyse results efficiently;
- Must put the methods in the hands of the practitioners;
- Traditional certification regimes consider only complete systems;
- Need Bayesian subjective probability to handle "confidence".

## 2.5 Panel E: Evaluating the Dependability & Security of Networked Systems – monitoring, operational assessment, auditing – Chairs report

**Co-Chairs**: David M. Nicol (Univ. of Illinois, US) and Marcelo Masera (JRC, Italy)
**Rapporteur**: Jim Just (Global Info Tek Inc., US)

**Speakers**:
Roy Maxion (CMU, US)
Evangelos Markatos (FORTH-Creta, Greece)
Alfonso Valdes (SRI, US)
Fabio Martinelli (IIT-CNR, Italy)
Todd Heberlein (NetSQ, US)
András Pataricza (Univ. of Budapest, Hungary)

### Introduction

During this panel session, discussions focused on dynamic and online analysis and evaluation methods and on real time assessment frameworks, including attacks observed, observation mechanisms, audits, measurement and decision making tools, etc.

### Prominent Challenges

The main research challenges identified were as follows: Exhortation with respect to metrics, measurements, analysis: it is imperative to start now, with limited systems and goals, to gain basic understanding. Many particular problems are induced by the inherent dynamicism of systems, users and threats. There is need for threat characterization, prediction and observation. Instrumentation and data collection play a key role for diagnosis, whether manual or automatic. There is a need to use on-line measurements to control and adapt, in particular, to put in place network information sharing techniques at all levels (including attacks observed, keystrokes of users, network traffic capture in an anonymous fashion and others). Headway is being made in the legal framework but there is lack of incentives for providing and sharing data. While there is need to have more data, it is also indispensable to ensure sufficient context that would permit replication through experiments.

Some difficult Research Challenges include:

- Developing useful metrics and testing them
- Scale
- Data collection

Any quantitative approach that can transform raw measurements into meaningful analysis requires well-defined metrics. Failure to make headway on the metrics problem will undermine the widely recognized need to develop an engineering approach to building and monitoring secure systems. Specific challenges that make the problem difficult include the diversity of applications, the complexity of systems, and the scale of the systems needing this technology. The only way to make headway is to start now, on a small number of highly focused examples. This ground-up approach will lead to understanding of deep issues that (hopefully) find expression in abstractions that are useful in more contexts. In this approach we need to develop:

- Clear, testable hypotheses and problem statements;
- Operationalized definitions;
- Principled, well-described experimental methods;
- Sharable and validated data sets;
- Ground truth;
- Calibration – benchmark gold standards;
- Repeatable measures and results;
- Well measured detected coverage;
- Tools tailed to the measurement/modelling/analysis needs.

A promising avenue for success is to develop multidisciplinary teams whose skill sets span application domain, system design, and system analysis, and begin work on a test bed with instrumentation capabilities.

The problem of scale arises because the systems of interest are large, as may be the data sets to be gathered and analyzed. Prevailing thought is that measurements capable of yielding insight will have to be detailed, and voluminous. Failure to deal with the problem of scale will limit the ability of any methodologies that are developed. One problem of scale is that "interesting" data may lay hidden in volumes of uninteresting data, and causal relationships between interesting data points are likewise hidden. Approaches for dealing with scale problems involve automation. Machine learning techniques might be developed to be coherence and a higher level of abstraction to large noisy data sets. Automation may also be used in developing models from the data, generate and test hypotheses, and develop causal analyses. Instrumented test beds and tools for generating attacks in the midst of background activity will be useful for creating controlled data sets that serve as a basis for experimental evaluation of the types of technologies mentioned.

Data collection is at the heart of a quantitative approach, and is rife with problems. Appropriate meta-data must be identified and recorded as well just to make sense of the raw measurements. The problem is that meta-data collection and dissemination runs into significant privacy and legal issues. However, on the legal front there are some promising developments that might be emulated. DHS has worked for 2.5 years to establish agreements between DoJ, ATT, universities, and DHS that will enable release of data and meta-data to researchers. There is hope for success as data providers seem to agree that old data has lost its business value and may be released for research.

## 2.6 Panel F – Future Test Beds – Chairs report

**Co-Chairs**: Doug Maughan (DHS, US) and Jim Clarke (Waterford Institute of Technology, Ireland)

**Rapporteur**: Zeta Dooly (Waterford Institute of Technology, Ireland)

**Speakers**:
Mike Bailey (Univ. of Michigan, US)
Pekka Nikander (Ericsson NomadicLab, Finland)
Anthony Joseph (Univ. of Berkeley, US)
Ollie Whitehouse (Symantec, UK)
Jim Just (Global Info Tek, US)
Henrique Madeira (Univ. of Coimbra, Portugal)

Test beds take many diverse forms and the spectrum of testing scenarios and requirements are diverse across US and Europe. As demand for services evolve, technology advancements and supporting infrastructures needs to meet this demand to provide adequate testing environments to deliver the desired level of quality services and products by all stakeholders including citizens, government, industry and academia. During this panel, challenges and opportunities in the Future Test Beds area were explored and discussed with the objective to identify the future requirements of a global testing environment that will complement existing work in this area.

### Prominent Challenges

From the presentations made and discussions held in this session, it was concluded that there is already a significant amount of work in future test beds in both the US and EU. However, there should be a stronger link and collaboration between the test bed communities beginning with a clear consolidation of what is available and experiences thus far so that both sides can have a real appreciation of what is happening and the best way to move forward.

An example of potential collaboration stems from the fact that the US side seems to have more facilities concentrating on the lower level (network) whereas in the EU, there is more concentration on the higher levels (services, applications). Therefore, the communities should leverage this synergy, for example, by establishing a software and services test bed that could be built on top of GENI with a number of application level experiments. Planetlab runs a production level services platform and University of Berkeley is also building an experiments cluster, which may also be used for this purpose. It was recognised that there would be a significant amount of collaboration required as the middle layers, which provides the OS would need to be jointly addressed also. It was decided that a good start would be to establish a top ten experiments list, including a requirements list for each of the experiments.

A number of potential examples were discussed including an application and service level test bed and a test beds for wireless or sensor networks. The rationale for the former is for running experiments that would open up valuable opportunities for innovative SMEs and Academics to venture into service-oriented solutions, without having to endure the expense and complexities of setting up a full scale testbed. For wireless and sensor network test beds, there are some standalone test beds already available but the issue that must be explored is to how to federate them taking into account cross testing, mobility aspects and security policy's as user's move in and out of different environments.

The EU is also pushing to ensure the inclusive participation of all the stakeholders including technology providers, service providers, academia, citizens, policy makers, and standards bodies.

In conclusion, the ability of a global test bed to facilitate diverse technologies will be key to the success of such an initiative and a number of challenges were highlighted during this session on future test beds, including:

- Enterprise Test beds, in which the edge is different than the core; Need enterprise "at scale" testing to enable previously unaddressed enterprise research;
- Test beds of the future need to consider the business issues of fair markets and ensure the balance of privacy & accountability;
- Combination of physical and cyber test beds, which must be open to all researchers
- The use of a data warehousing approach to analyze, cross-exploit and share test bed results enabling knowledge collection and management;
- Need predictive science for networks
  - No methodologies for comparative assessments – Test beds (including simulation/emulation) can possibly help with assessment
  - Automation required for testing
  - Need to be able to remotely manage testing process
  - Traffic generation is a constant problem;
- There is a trade-off which must be addressed – details/rigor vs. speed/slop;
- Application and Service level testing needs to be supported, especially for potential consumers/providers that ordinarily wouldn't be able to gain access to these technologies underlying infrastructure e.g. SMEs;
- Trade-off: Generalization vs. Specificity and the utility of the results;
- Wireless (and other application areas) – Needs to be federated connected to other test beds in a secure fashion.

# 3 Final Conclusions and input to the next EU-US Cyber Trust: System Dependability and Security Summit workshop

One of the objectives of this summit was to explore common research themes and challenges in Europe and the US within the Security, Trust and Dependability environment and how to facilitate research initiatives that are common to both the EU and the US.

Through a number of panel sessions on these topics, participants gained a shared understanding of critical issues, identifying some initial dependability and security research directions. Key contacts were made that will facilitate the emergence of future collaboration initiatives. Different types of collaboration techniques were discussed, such as joint projects where funding was identified as the main barrier to joint initiatives as both sides have different funding mechanisms that do not easily facilitate cross-Atlantic research projects. Identification of common objectives and priorities for research in the TSD was also cited as a hurdle to joint collaboration but this workshop and other similar events were a step in the right direction to establishing these common research objectives and priorities.

While the original concept of holding the Cyber summit workshops was initiated between the European Union and the United States, participants from Australia, Canada and Japan were invited to explore areas of potential collaboration. These highlighted international, both industrial and academia, collaboration towards global efforts for establishing benchmarks and identification of a simple[1] set of metrics, methods of measurement, calculation, validation, and evaluation, but also interdisciplinary approaches where the evaluation results are made visible to all the stakeholders and get mapped to economic values. In addition, another area of collaboration highlighted by the Japanese representative was the exploration of new theories and technologies to handle and filter the enormous amount of information produced ("information explosion") and recorded on global networks everyday in the world to extract only useful and validated information for the society, businesses and international security. This would involve exploring methodologies to verify the trustworthiness and evaluate the dependability of the summarized and most likely prioritized information.

Other collaboration techniques include student and faculty exchange and shared creation and/or use of research resources, e.g., test beds, data warehouses and knowledge bases. The workshop concluded that collaboration would be beneficial to both the EU and the US due to efficiency, systems are international in scope and, thus, there are opportunities in size and volume of research projects.

[1] simple in the sense that the metrics can represent the degree of dependability for services provided by networked systems understandable from the user's point of view.

Barriers to collaboration include the differences between different countries, specifically, legal issues and attitudes, and lack of just one governing and financial bodies.

At a more detailed level, research collaboration could include mechanisms for Government agencies (e.g. by re-sourcing a facilitating organisation or cooperating pair of organisations) to encourage and assist international collaboration e.g.

- Provide support function to facilitate collaboration
  - Point Principal Investigators (PIs) to consumers of research results;
  - Help PIs prioritize research ideas;
  - Help navigate multiple government offices, rules, etc.
- Facilitate PI gatherings
- Provide funding in a form yet to be defined.

Other collaboration could materialize if calls are aligned in content and in time; if both are available and selected – we would ensure that there is a work-package for international co-operation.

A longer term joint programme could be investigated using current science and technology agreements – for example, extend this for DHS and European Commission Framework Programme (e.g. ICT-FP7). The relevant participants attending the workshop agreed to explore these possibilities further.

Consultation meetings in the EU could link with PI meetings in the US. Moreover, EU and US funding agencies could exchange more information regarding ongoing research activities, scientific and technological objectives, plans, and new program announcements and calls for proposals and may be influenced more proactively by each other in this kind of context.

The participants were asked for their feedback on this workshop to help with planning for the second workshop within 2007 to include: Technical Content; Objectives; and Organization.

It was concluded that the first EU-US Workshop on "Secure, dependable and trusted ICT infrastructures" had indeed succeeded in achieving the core objectives of focussing and gaining consensus on difficult but important research challenges for joint collaboration in the EU and US (and potential for Canadian, Australian and Japanese) research communities working in the TSD area. It was agreed that the second workshop goals are to incorporate wider views and inputs from external stakeholders including industry, academic, government and users and the establishment of a core steering group (SG). The objectives of this SG would be to act as a catalyst in focussing the output from the workshops, in defining and facilitating joint activities for the TSD research communities, and in promoting an alignment of potential funding mechanisms. It was also agreed that we need to define a common vocabulary to ensure there is consensus within this community when addressing these mutual issues and challenges.

# 4 Appendix A - Position Papers

The following section contains all the position papers received from participants in advance of the event; these were used as input to the preparation for the panel sessions and the presentations. A separate table of contents is presented on page 3 of this report.

# Panel A

## Dependability & Security of Future Networked Systems – architecture and design issues

# Security Challenges on Long-Term Data Preservation- David H.D. Du

David H.D. Du

National Science Foundation, USA

Concurrent with the exploding of Internet, businesses of all sizes also have demanded scalable, reliable, and secure storage. Fixed-content data, such as check images, medical imaging scans, video/audio files, and email records, for example, may be infrequently accessed but usually must be retained for long periods of time and must be readily accessible when needed. Retaining and protecting an enterprise's information assets is vital for conducting business and often is required by law and government mandates, such as the Sarbanes-Oxley Act, HIPAA, DOD 5015.2-STD, and the European Data Privacy Directive. With the globalization of the economy, business data needs to be available twenty-four hours a day, seven days a week. Furthermore, in the event of a disaster, the data must be restored as quickly as possible to minimize the business' financial loss.

Current storage and data archiving technologies are simply incapable of satisfying the increasing demands for long-term data security. The need to protect and preserve cryptographic keys in a long-term archive also presents a major challenge since many unforeseen changes can occur during the data's lifetime. For instance, the user who originally encrypted the data may be unknown or unavailable when the data is to be decrypted. Furthermore, keys can become compromised as can the cryptographic algorithms themselves. All of these issues require fundamentally new approaches to develop a storage archive capable of securely preserving data for thirty years or more.

There is a definite need for a long-term archival data management system that will enable high-performance while maintaining appropriate levels of security throughout the data's lifecycle. The following desired properties of such a system are identified. We are assuming data integrity and protection can be done by a key management. In fact, this is also a challenge issue.

- **Backup and efficient retrieval of keys**: Key backup is one of the fundamental requirements of key management for long-term storage. The cryptographic keys should be secured as long as the data is considered to be existent. It is almost impossible to recover encrypted data if the associated keys are lost. This implies: 1) If the data is backed up, keys must be backed up, 2) If the data is deleted, keys should be deleted, and 3) Keys must be backed up securely. In order to retrieve data that was encrypted or signed several years ago, the system should also retrieve the keys that were used to encrypt or sign the data. Considering the large number of keys that will be created in an organization over time, this can be a daunting task. As a result, *efficient search and retrieval* of these keys is important.
- **Key recovery:** In practice, just like the data, keys can be lost or accidentally deleted. Loss of keys will result in loss of data. Therefore, the key management mechanisms must ensure that lost keys can be easily recovered. Further, it can happen that the user who encrypted the file may not be available at the time of decryption (e.g., the user no longer works for the organization). Under such circumstances some designated agents of the organization should be able to retrieve the keys. Key recovery mechanisms must be secure and efficient. In order to ensure that the organization will be able to retrieve keys during

emergency situations, the system should also ensure that the encryption keys used by the user are valid. For example, a disgruntled employee can encrypt all of the data and destroy the keys, which will result in destruction of the data. To prevent an employee from encrypting data without archiving keys, the system must validate before storing the data that the encrypted data is actually encrypted by a key that is guaranteed to be backed up.

- **Long-term Management:** Typically, enterprises are periodically reorganized. Reorganization can merge or split existing groups. As a result, the keys should be reorganized as well. Therefore, key management mechanisms should efficiently handle such reorganization activities. Since the data and the keys should be protected for long durations of time, many unforeseen events can happen. For example, the keys can be compromised, or a cryptographic algorithm can be compromised since an attacker has more time to attempt to break a key or an algorithm. If an attacker can compromise a key, s/he could read encrypted data or even write data without being detected since the attacker can generate a valid signature. Therefore, the system must be able to gracefully handle these disasters without compromising security.
- **Usability:** Strong security with poor usability and/or performance is not sufficient to make the system practical. This is the reason that secure systems and cryptographic software are used less than we would expect due to their lack of consideration for usability of their products. A recent survey performed by Sun Microsystems shows that 45% of the surveyed companies find storage management to be the most expensive task, and 95% of the surveyed companies state that it is the most difficult. Considering these, usability (by the administrator, and, naturally, by the client) is one of the major concerns for this project besides the traditional performance, security, scalability, and availability requirements. A user should be able to use the file system that provides long-term key management without having to know any underlying cryptographic operations. Transparency should be provided as much as it is necessary.
- **Scalability:** The scalability is required in several aspects. The archiving storage capacity and bandwidth should be able to scale up as the data volume increases. Due to the huge amount of data generated for long term, the number of keys will be increased with respect to the data. The archival capacity and bandwidth should be increased with data volume. The key management system must scale well with respect to the number of keys. Further, the number of keys that must be secured should be minimal.

# Embedded systems pave the road to the future Internet – Paulo Verissimo

**Paulo Esteves Verissimo**

**University of Lisboa Faculty of Sciences**

**pjv@di.fc.ul.pt , www.di.fc.ul.pt/~pjv**

***A research grand challenge***

'The future of Internet lies with embedded systems'. The key to this challenge lies not in software for peer-to-peer autonomous systems, and/or software for embedded systems, and/or software for pervasive systems. It is the whole of them, but not the sum of them. The key to the challenge lies in the confluence of these problems, what I might equate, in systems terms, as learning how to design new generation *large-scale, complex and networked systems-of-embedded-systems*, because this is how the future Internet will look like.

The opening statement encapsulates a powerful and multi-disciplinary challenge, which I will try to explain in these brief words. There will be many more computing devices in embedded, sometimes function-specific, inconspicuous devices, pervading the societal tissue (business, architectonic, even human) than there will be 'computers' as we know them today (stuff that people program and interact directly with, even if as small as hand-held). Most, if not practically all of them, will be networked, interconnected in some form. Software, architecture and protocols animating the systems under this vision will not have the same characteristics as we are used to in classical Internet-based systems. However, they can not either be designed under the assumptions and techniques prevailing in the embedded systems arena (hard real-time, static, closed and integrated systems).

***Why this is a grand challenge***

As with any other grand challenge, the impact of addressing it is estimated as commensurate, though of opposite sign, to that of not addressing it. One of the things that make it a grand challenge is exactly the difficulty of understanding and tackling all these problems simultaneously, since they have been addressed by different research communities so far. One of the arts of the research sponsors will be identifying it as a multi-disciplinary challenge whose roots are centred in, but not confined to, distributed systems theory and practice.

Why is it crucial to the future Internet? In concise terms for an abstract, the message is: Internet-oriented researchers, embedded systems came into your life (as I might also say: embedded-systems-oriented researchers, Internet came into your life). Addressing and solving this challenge will contribute to keep the future Internet as a useful asset of society, because the risks are manifold. And time is not on our side. Whether we want it or not, the future will bring systems described in this vision into societies' daily life. If we are not prepared, such Internet-born embedded+ pervasive+ ubiquitous systems, if not working properly, can do more damage than good. And the future has started...

Imagine cars or planes connected to the Internet for entertainment or assistance, having their by-wire control systems interfered with. SCADA embedded systems controlling electrical utilities, once designed as closed systems, being now intruded upon from the Internet-connected corporate network, and sabotaged.

Your home's electric oven setting fire to your house, instead of pre-cooking your meals controlled by you from the Internet, all because the oven computer wrongly joined some peer-to-peer computation, or fell victim to a stupid script-kid. And so forth.
If I'd dare use a "sound byte" to equate how relevant this challenge is, I would say that in such a new reality (an Internet made of large-scale systems-of-embedded-systems) doing Control-Alt-Del may not be enough for mending something that went wrong. And in an era of attacks and intrusions, things will definitely go wrong, sometime, in some way.

***A bit of roadmap***

I will briefly enumerate some systems research topics required to address the problem, all under a distributed systems context:

- Reference architectures for pervasive/ubiquitous/embedded systems --- since existing architectures normally specialize in one facet.
- Event-based middleware --- encompassing a mixed software+devices reality, that is, events generated indistinguishably by software and by hardware devices, such as gadgets, sensors and actuators, or networks thereof.
- Sentient object-oriented computing models --- a form of encapsulating and easing programming under this vision, where objects, besides the usual method computations and interactions with other objects, are capable of directly consuming and producing events to/from devices.
- Component-based software architecting --- where object composition should obey, whenever relevant, the structure of sub-systems and hardware devices populating our visionary future Internet.
- System evaluation --- as a form of determining whether the properties of these new systems meet the desired levels of QoS, dependability and security.

***Some research resources***

The research infrastructure required need not be impressive. If we consider as a starting point a typical Internet-based infrastructure for distributed systems research, then additional means may include: sensor networks; RFID systems; wearable computers; control area networks such as field-buses; emulators of real control or cooperative systems made of embedded devices in several areas (e.g. robots; cars or planes; electrical power or process control plants; home systems). Example would-be prototypes for systems research in this area might be: cooperative enhanced human and/or autonomous robot teams; cooperative cars (e.g. platoons); assisted road-side and Internet-born navigation and driving; Internet-born car maintenance, diagnostic upload and firmware download; remote operation of critical infrastructures through Internet-born secure links; Internet-enabled safe and secure home control and information sharing.
When will the challenge be met? We can anticipate that the challenge will be met when we see prototypes of systems-of-embedded-systems being deployed over the Internet, in a performing, user-friendly, safe, dependable and secure way. Then we will know that the path to the future Internet is clear. Along this road, relevant milestones may consist of proof-of-concept prototypes showing partial results in the research topics enumerated above.
(1) An earlier version was presented at the NSF Workshop on Grand Challenges on Distributed Systems, MIT Cambridge USA, September 2005.

# Application-Aware Reliability and Security: The Trusted ILLIAC Approach- Ravishankar Iyer

**Ravishankar K. Iyer**

**Coordinated Science Laboratory and the Information Trust Institute**

**University of Illinois at Urbana-Champaign**

**1308 W. Main, St, Urbana, IL 61801**
**rkiyer@uiuc.edu**

Security and reliability are the key attributes in building highly trusted systems. System security violations (e.g., unauthorized privileged access or the compromising of data integrity) and reliability failures can be caused by hardware problems (transient or intermittent), software bugs, resource exhaustion, environmental conditions, or any complex interaction among these factors. To build a truly trustworthy system, the designer must find ways to mitigate (avoid and tolerate) against accidental errors and malicious attacks.

***Trusted ILLIAC*** **is a reliable and secure cluster-computing platform being built at the University of Illinois Coordinated Science Laboratory (CSL) and Information Trust Institute (ITI), involving faculty from Electrical and Computer Engineering and Computer Science Departments.** ***Trusted ILLIAC*** **is intended to be a large, demonstrably trustworthy cluster-computing system to support what is variously referred to as** ***on-demand/utility computing*** **or** ***adaptive enterprise computing*****. Such systems require that a significant number of applications co-exist and share hardware/software resources using a variety of containment boundaries. Current solutions aim at providing hardware and software solutions that can only be described as a one-size-fits-all approaches. Today's environments are complex, expensive to implement, and nearly impossible to validate**

**The challenge is to provide an application-specific level of reliability and security in a totally transparent manner, while delivering optimal performance. A promising approach lies in developing a new set of application-aware methods that provide customized levels of trust (specified by the application) enforced using an integrated approach involving reprogrammable hardware, enhanced compiler methods to extract security and reliability properties, and the support of configurable operating system and middleware.** Our approach is to demonstrate such a set of integrated techniques that span entire system hierarchy: processor hardware, operating system, middleware, and application.

**At the processor level,** a Reliability and Security Engine (RSE) provides a hardware framework that enables embedding low-cost, programmable hardware modules to provide application-aware error detection and security services (e.g., process hang detection, selective duplication of the instruction stream, and detection of memory-corruption attacks).

**At the operating system level,** we propose a trusted microkernel that is a software framework deployed as a loadable kernel driver in Linux. This framework provides support for rapid detection of failures and attacks, including transparent application checkpointing and recovery.

**At the middleware level,** we employ a self-checking, run-time framework to provide detection and recovery to applications using flexible and configurable software solutions. A set of well-defined communication gateways facilitates the robust, low-overhead flow of information between the system and applications, while providing trust guarantees.

**At the application level,** we plan to enhance the Illinois COMPACT compiler to support the automated generation of assertions for runtime detection of accidental failures and malicious attacks, while accelerating application performance. The idea is to use information generated during the compilation process and at run-time to identify patterns in data variables and compose compact signatures of correct application behavior. Assertions for runtime signature checking can be integrated within the application or implemented directly into the RSE hardware.

**A distinctive, integral concept of the *Trusted ILLIAC* is that of a *validation framework*, which constitutes a cornerstone for quantitative assessment of alternative designs and solutions.** Such evaluation is crucial in making design decisions that require the management of tradeoffs, e.g., between cost (in terms of complexity and overhead) and efficiency of proposed mechanisms. **The framework leverages years of experience in analytical and experimental evaluation of highly reliable and secure systems, on comprehensive fault and attack injection technology, and on modeling and simulation tools developed in Illinois.**

*Trusted ILLIAC* is based on research and support provided by, among others, The National Science Foundation, MARCO/GSRC (SRC and DARPA), IBM, HP, AT&T, AMD, Intel, Motorola, XILINX, Nallatech, and the University of Illinois.

# Advanced cryptology: overkill or essential tool? - Bart Preneel

**Bart Preneel,**

**COSIC, K.U.Leuven, Belgium, bart.preneel(AT)esat.kuleuven.be, http://www.ecrypt.eu.org/**

**<u>Contribution to Panel A:</u> Dependability & Security of Future Networked Systems – architecture and design issues**

Cryptology has played an essential role in securing electronic communications. After the foundational work of Shannon in the 1940s and Diffie and Hellman in the 1970s, it has taken several decades to develop crypto as a scientific discipline and to bring the solutions to the mass market. We have witnessed the development of open standards for cryptographic algorithms and protocols (AES, SSL/TLS, IPsec). More recent successes are the understanding of the concept of "authenticated encryption" and efficient primitives to realize it.

It is clear that in a networked system, the connection itself is no longer the weakest link. Most of us are familiar with the quote from Prof. Gene Spafford ("Using encryption on the Internet is the equivalent of arranging an armored car to deliver credit card information from someone living in a cardboard box to someone living on a park bench"). From this one could conclude that the answer to the question in the title is: advanced crypto is overkill since the basic crypto we are currently using is already too strong.

Developing more robust and secure networked systems is definitely a very important research challenge, and I am convinced that distributed crypto can play an important role here as well. Any large scale networked system needs strong crypto (for example, to authenticate address records); one the other hand, it seems realistic to assume that – even with improved baseline security – individual components of the infrastructure will be penetrated. Multi-party computation techniques allow to distribute critical operations and to refresh secrets shared between multiple nodes without a single long-term secret being present in one node; the main requirement for these protocols is that a (weighted) majority of nodes is honest. This can clearly be an important building block in a robust infrastructure – but also for an end user, who could share his critical secrets over multiple devices (PC, laptop, PDA, mobile phone); this could help address the problem of the insecurity of the endpoints. While the theoretical foundations for multi-party computation have been laid in the 1980s, there is still a large need for research on how one can implement these basic primitives efficiently under realistic circumstances and which more complex primitives one can realize in practice.

The topic is privacy has been included in the title of Panel C (and not in the one of Panel A). However, I would like to mention it here too, I believe that privacy is an integral part of security and that is the biggest challenge facing security researchers in the next decade. This may seem surprising: technological developments, market forces and national security concerns seem to lead in an unavoidable way towards increased surveillance and erosion of privacy, to the extent that some people believe that we should give up this battle. Given the title of this panel, I would like to stress here the importance of the protection of

privacy at all communication layers: an anonymous e-cash or e-voting scheme does not make sense if every packet contains in its header the unique address of the sender. The development of efficient and robust anonymous communications and the protection against traffic analysis is clearly a very important research challenge; we have made just the first steps in this area and we are very far from any large scale deployment.

At the level of cryptographic protocols and algorithms, there are also several important challenges:

- the design of an efficient and secure cryptographic hash function; many protocols and applications depend on hash functions, but the recent attacks on MD4, MD5 and SHA-1 have shown us once more how shaky the foundations of cryptographic primitives are;
- the design of cryptographic protocols that are more algorithm-agile; in the past years we have learned that even if our protocols include algorithm negotiation (as introduced in the mid 1990s), this does not mean in practice that upgrading algorithms (such as currently the hash functions) is easy;
- the development of secure implementations of cryptographic algorithms, that offer protection against side channel attacks (e.g., cache-based timing attacks on cryptographic operations that use tables);
- in the long run we need to improve the available trade-offs for cost/security/performance. The main challenges relevant for fixed networks are how to achieve high speed encryption (Terabit/s) and long term security (50 years or more).

For some of the challenges discussed in this note, there exist solutions at various levels of maturity, while other areas are completely open and need long-term fundamental research. In the area of cryptographic hash functions, an open international competition (similar to AES, NESSIE) seems appropriate. In some cases strong benefits could be obtained from developing solutions and testing first versions in realistic environments in order to gain a better understanding of what the current limits are and in which direction solutions should evolve.

# The Insider Threat: A Challenge to Scalable Networked Systems – Yair Amir

Yair Amir
Department of Computer Science
Johns Hopkins University
yairamir@cs.jhu.edu

**Introduction**

As network environments become increasingly hostile, even well protected systems constructed with security in mind may become compromised. There is a significant need for systems that can achieve their main goal even while compromised by a malicious adversary. This challenge manifests itself in two related domains: the network infrastructure itself (compromised or malicious routers and overlay routers) and information access systems (compromised servers).

To date, most practical efforts addressing the insider threat are designed with an actual insider - a person, in mind. With the advent of malicious hackers and the knowledge we already have regarding networks with 100,000s of compromised zombie nodes, it seems very likely that some of these zombies serve important functions in critical systems. Malicious hackers currently tend to use these zombies as generic networking and computing resources. However, it is only a matter of time before they will start examining what kind of zombies they have more closely and exploit them in their original roles as attackers from the inside.

As almost all systems are connected to the Internet these days, this is a real threat to any networked system. Architecting systems that can function correctly after parts of them are compromised will probably be one of the most important challenges for large-scale dependable systems in the future.

**Potential Research Directions**

While there are theoretical answers to some aspects of this problem, they are mostly limited to small-scale systems that are confined to one room. The greater challenge is to design systems that are practical and scalable and yet able to withstand such attacks.

There is a need to study what are the parameters under which such a system can achieve its goals and what is the cost (performance, monetary, system complexity) systems will need to pay for this added resiliency.

There seem to be two relevant domains:

- Network infrastructure resilient to some of the routers acting maliciously. This includes messaging systems, overlay routers, wireless mesh networks, etc.

  The research community is mostly starting to look at hardening specific common routing protocols to some (usually limited) router adversarial behavior. Perhaps a more fundamental approach can be taken to examine what is and is not possible to achieve for any messaging system under a strong insider attack, devising models, metrics and algorithms that can then be adapted to the common routing protocols.

- Distributed information systems resilient to some of the servers being compromised.

The research community investigates different flavors of Byzantine coordination and replication algorithms in the last few years. Most of these solutions are confined to relatively small systems. Broadly speaking, current approaches tend to make one of two assumptions: they either assume that at most some percentage of the system (e.g. less than a third) will be compromised, or they assume that some specific parts of the system will not be compromised. Both of these assumptions seems too rigid and (may I say) a bit contrived and probably do not reflect reality well.

Perhaps a new approach can provide the best of both worlds by architecting hybrid approaches that, for example, provide stronger guarantees for critical parts (beyond just assuming they will not be compromised) while allowing broader overall or localized compromises. Better ways need to be devised to scale such systems to Internet size (at least geographically), possibly using careful hierarchical approaches.

**Common Research Resources**

It seems that the above challenge is especially hard for large systems that are geographically distributed. Having a global test bed can go a long way toward reducing the cost of such research. Planet Lab is providing an adequate solution to many research endeavours, however, as a shared, always heavily used, resource it does not fit such research because it lacks reservation mechanisms that allow access (for a limited time) to the full power of a large-scale networked environment. This is required to allow meaningful experiments addressing important performance issues (e.g. latency).

**Concluding Remarks**

Architecting a large-scale system that is resilient to insider threats and is able to achieve its main goals even after being compromised is a formidable problem that may require research involving distributed algorithms, networks, system security, cryptography and policy. There currently exist few practical solutions, while among theoretical approaches, each has its own assumptions that are too firm and therefore may not be totally practical. Moreover, the significant cost associate with existing approaches is most likely too prohibitive to be adopted by system builders.

This problem domain is potentially a good fit for US-EU collaboration as there are several groups working on it on both sides of the Atlantic, and interestingly, they bring different approaches to the table that can point to very useful potential synergy.

# Position Thoughts – Neeraj Suri

**Neeraj Suri**

Define 1 or 2 difficult problems (research challenges)

- It is the networked computing/communication environment that increasingly determines what is or isn't viable from the viewpoint of providing a trusted service. More often than not we direct our focus on the trust properties we would like to achieve than detailing the environment on which to obtain the property. We really need a proper specification & semantics of the overall networked, heterogeneous (mixed-mode) environment along side the specification & semantics of "trust" relevant to that MM environment (MME). We usually like to deal with a sub-set environment though the security relevant interactions and implications come from the collective inter-connected environment.

- We make assumptions about availability of resources (higher bit crypto, more computing power etc) when it comes to provisioning of security. The real-world environment (MME) doesn't allow for this convenience. We really need to think of provisioning of security from a technology-invariant viewpoint

For each problem, describe the impact of solving or failing to solve the problem and a description of specific challenges making the problem difficult.

Provide potential research directions for handling the problem, and metrics for measuring progress against the problem. Possibly, describe also the limitations of current approaches, promising approaches not yet fully explored and desirable approaches where relatively few ideas have been proposed.

- I will present a "data centric" viewpoint on the mixed-mode computational and communication environment
- Will present viewpoints on providing run-time security: issues of run-time security profiling and run-time security enhancers approach

# Panel B

## Dependability & Security of Future Networked Systems – scalability and context-awareness

# Dependability & Security of Future Networked Systems - John Knight

John Knight
University of Virginia

Define 1 or 2 difficult problems (research challenges)

(1) **Scalability of network management mechanisms**: Large networks are designed to provide applications with rich *semantics* so as to enhance their utility and *flexibility* in the face of environmental changes, malicious and non-malicious damage, and application changes. Both semantics and flexibility rely on network management—placement of data, software and computation demand, and organization of the target resources. Some of the necessary management is automatic, but much of it presently requires human intervention. Future networks will continue with the provision of management services but future management will have to be much more automated and will have to scale as network sizes increase. The provision of automatic, efficient and effective management in very large networks is an unsolved problem. As an example, consider the problem of merely ascertaining the current value of some aspect of the network state for a network with a million nodes in order to be able to make a management decision.

(2) **Implementation assurance**: The vast majority of the dependability problems that arise and have arisen in large networked information systems are the result of defective software. The operating systems, the network systems, and the applications themselves all host vulnerabilities that can be the subject of security attacks, losses of availability, and so on. Despite this, the demands imposed on system developers are considerable and getting worse. More software is needed to provide more functionality on larger and larger networks in less time. Clearly, this is a recipe for disaster since this is a problem that grows rapidly as the size of the network increases.

For each problem, describe the impact of solving or failing to solve the problem and a description of specific challenges making the problem difficult.

(1) Failing to develop effective approaches to the proper management of very large networks will inhibit our progress in development of next generation information systems. Crude, ad hoc solutions in which coarse-grained techniques are used will not allow the diverse spectrum of anticipated applications to be developed. Fundamental advances are needed that will permit what amount to complete “programs” to be built for entire networks so that network-wide management algorithms can be described and implemented on networks of arbitrary scale.

(2) No matter how novel and effective the algorithms are that are developed for network management (or anything else for that matter), they will be useless unless they can be implemented correctly. The track record of failures in current systems shows clearly that the weak link is defective implementations. Unless we deal with the problem, we will be held ransom to uncertainty about exactly what our networked information systems will do. For security, we will never know whether the systems have been compromised. For availability, we will never know whether the systems will operate as desired when we desire it.

Provide potential research directions for handling the problem, and metrics for measuring progress against the problem. Possibly, describe also the limitations of current approaches, promising approaches not yet fully explored and desirable approaches where relatively few ideas have been proposed.

(1) The most promising direction of which I am aware for network management is to use content and attribute addressing. This approach allows network resources to be selected as a group based on properties that are desired. It has two other major benefits: (a) there is no need to maintain a record of network state that is separate from the network elements themselves; and (b) the state record of the network is current even if nodes join or leave the network.

(2) The software problem remains with us despite years and years of research. The first crucial step in the solution is to make radical changes to the current system of higher education. Most graduates in computer science and computer engineering in the United States do not have an adequate training in software engineering, even less in the relevant material needed for systems requiring significant levels of dependability.
Continuing research is clearly important but I am not convinced that concepts such as self-healing are appropriate for software. Design faults are, by definition, faults that the designers were unable to avoid or detect. To expect such faults to be routinely either avoided or corrected by automatic means is a laudable but very ambitious goal.

# Dependability & Security of Future Networked Systems – scalability and context-awareness – Brian Randell

The systems that most interest and concern me are large critical systems-of-(computer-based) systems. By "computer-based system" I mean a system that is in essence composed of hardware, software, and people. Typically these systems involve issues of scalability, and often - especially at their edges – of mobility and context-awareness.

A fearsome example of the sort of system I have in mind at present is the UK National Health Service's "National Programme for Information Technology" (NPfIT). This system-of-systems is intended to provide a country-wide integrated information service, to be directly used by over 200,000 doctors, hospital consultants, pharmacists, nurses, medical secretaries, etc. in over 300 hospitals, 30,000 doctors' offices, etc. The information to be held and shared includes electronic patient records, digitised X-rays, hospital management data, appointment bookings, prescriptions, laboratory results, etc. (It is in fact claimed to be the world's largest-ever civil IT procurement, with a total cost now admitted to be £20-30B, i.e. €30-45B. And it is the subject of much media concern – see http://nhs-it.info/ – given its costs and delays, and the necessity of it being highly reliable, available and secure.) However, there are many other critical systems-of-systems under development or being planned, even if not on quite such scale, which also merit our attention.

It is clear that the problem of achieving adequate levels of dependability and security from such systems, *from the time of initial deployment* (as opposed to only after years of painful evolution – something that is quite unacceptable for highly-critical systems such as NPfIT) is an extremely difficult one. Descriptions of the impact of not solving such problems abound, in various catalogues of major computer project disasters.

Despite many attempts to the contrary, it seems evident to me that the only successful way of building and deploying such systems-of-systems is incrementally, starting with component systems that have already gained the trust and confidence of their users. (John Gall, in his book *Systemantics*, puts it this way: "*A complex system that works evolved from a simple system that worked*".) However, even with such an approach (as opposed to a single big-bang, top-down, huge centralized implementation effort) there is the problem, in large part one of scalability, of how to assure (i) that the component systems remain adequately dependable and secure, as the overall system-of-system comes into being, and (ii) that this overall system itself has adequate dependability and security, especially as the systems evolve and grow to meet the needs of changing environments and user requirements. Consequently the research problem that I and my colleagues, in for example ResIST, the EU Network of Excellence on the Resilience of Information Society Technologies, are pursuing is that of *how to express the dependability and security-relevant characteristics of systems and components, and of system design and evolution processes, and then use this "meta-information" to guide, and predict the outcome of, these processes.* (The term we use for this approach is "resilience-explicit computing", after the original LAAS proposal of some years ago of "dependability-explicit computing". This work is linked to work in ResIST on a "Resilience Knowledge Base", which in its present prototype form already holds a very large amount of dependability and security research-related information on people, institutions, projects and publications, and which we in due course envision evolving so as also to hold designs and operational data and thus become the support platform for resilient-explicit development and evolution activities.)

The second point I wish to make is that the design and implementation of such systems is evidently as much a socio-technical task as a technical one. It thus really needs an interdisciplinary team, including experts in the dependability and security problems (and solutions) that arise from the various different types of system "component" – hardware, software *and* human. I and my colleagues in the six-year five-university "Interdisciplinary Research Collaboration on Computer-Based systems" (DIRC) have found, as have others, that (i) building an effective such team not only takes a long time, but (ii) is best achieved by having the team collaborate closely on the actual problems of particular real systems. A second research problem I therefore wish to highlight is that of *how to target and carry out such interdisciplinary application-specific research so that the results are likely to be of general applicability*, i.e., of evident relevance in a number of different application domains and system environments. Failing to achieve this, by producing results that are at best of interest only in the immediate application domains and environments, has obvious consequences – other potential beneficiaries lose out, and the overall impact of the research may be rather low relative to its costs, even if people and organizations in the immediate application domains and environments express satisfaction. Such a problem is not unique to computer systems research, leave alone unique to research that is specifically aimed at the problems of developing large critical systems-of-(computer-based) systems, but it is especially challenging in this particular arena.

It is, as is so often the case in research related to large and complex systems, difficult to define meaningful direct metrics for measuring progress on either of the research challenges that I've discussed – however the more important measure is that of the extent to which the results are taken up by other projects, and in particular by government and industrial development projects.

# Two Scaling Problems for Network Intrusion Detection - Nicholas Weaver

International Computer Science Institute

a: Network intrusion detection and response is a key component of any computer security infrastructure. A NIDS system offers an otherwise unavailable global vantage point of the network state, allowing it to observe attacks which are undetectable from the end host. By being in the network and therefore separate from the end hosts, NIDS systems can contain attacks as well as block attacks. Finally, NIDS can protect otherwise unprotectable yet extremely security critical devices, such as high-volume printers, control systems, and legacy systems.

In order to be resist evasion and produce useful high-level analysis, NIDS systems need to perform detailed semantic reconstruction of the network stream, beginning with TCP stream reassembly and protocol-specific parsing, before applying context dependent policies on the traffic. These tasks are very intensive, requiring large amounts of memory and processing power.

The sheer computational requirements results in two critical scaling problems. First, we must improve the cost/performance of NIDS by one to two orders of magnitude if we wish to meaningfully protect today's systems. Second, we must develop techniques which can scale NIDS on parallel hardware if we want to be effective in the future.

The first requirement, a one-time-only improvement of 10x to 100x in cost/performance is necessary because the traditional NIDS installation, at the institutional perimeter, is ineffective against worms, insiders, and many other contemporary threats. In particular, once an attacker or worm gains control of a ***single*** system within the firewall, the attacker effectively becomes an insider. We thus believe [hardlans] that NIDS needs to be ubiquitously deployed throughout the local network.

Yet the LAN vantage point imposes a huge challenge on price/performance. Because LAN bandwidth is vastly cheaper than WAN bandwidth, and the number of devices needed to create a defense must be significantly higher, we believe that LAN based defenses need to be vastly cheaper in the amount of bits of traffic processed per dollar, requiring a one to two order of magnitude improvement in cost/performance. Otherwise, hardening the LAN against attacks would be economically intractable for most networks.

The second, scaling through parallelism, is a consequence of the Moore's Law curves for network bandwidth, CPU performance, and memory latency. In the past, CPU performance and network bandwidth have followed similar exponential curves, with single-threaded CPU performance continuing to scale and larger caches used to compensate for increasing DRAM latency. Thus as network traffic grew in volume, the IDS developer could count on CPU scaling to match.

Unfortunately, over the past 3 years, the performance growth single-thread CPU performance has slowed rapidly. Processor vendors have responded with multicore systems, which offer the promise of scaling overall throughput even if single-thread performance no longer improves. It requires a very different IDS structure to take advantage of parallelism. Yet without the ability to construct a parallel IDS, we will be left behind as ever increasing network traffic and attack sophistication leave our IDSs obsolete.

b: Both these problems are critical if we wish to develop robust computing systems. Although there is a significant commercial interest in Intrusion Detection and Response systems, these are mostly "checkbox security", a device placed on the network perimeter primarily for compliance rather than actual utility. As such, there is little economic motive for deep stateful analysis when simple signature detection and incomplete stream reassembly suffices to meet current market demand.

Additionally, it is not rational for a current player to undercut the existing market by the amount we believe is necessary to enable LAN-centric defenses, and the significant risks involved in exploiting the available parallelism forces the research community to step up.

Finally, much of these defenses are relevant to emerging attacks. Traditionally, there is a huge social an economic focus on preventing previous attacks, not future threats. Thus overall, we believe that these market failures force the research community to invest significant time and effort. Without the ability to reliably protect our network devices, we will remain vulnerable to future attacks.

c: Both objectives, improving cost/performance by 10x-100x and enabling IDS systems to scale with improvements in networking, require a careful hardware/software co-design.

For cost/performance scaling, we are already experimenting with a technique called shunting, which allows a low-cost hardware accelerator to be coupled to a conventional IDS. The IDS can then use this hardware for fine-grained (flow by flow) control over all traffic on the link to focus its effort on the traffic of interest, while efficiently passing the less-interesting traffic, such as encrypted SSH sessions or bulk file transmission. In practice, this allows the IDS to ignore 80\% of the traffic observed on a typical WAN link, greatly reducing the load on the peripheral bus and memory bus, loads which would otherwise prevent the IDS from operating at full rate on Gigabit networks.

We are extending the hardware mechanism to support VLAN-rewriting, which will allow our system to isolate, monitor, and control all the traffic on the local network. There are many open questions we need to explore, including whether the shunting hypothesis holds in a LAN

environment, how to rewrite broadcast and multicast traffic to enable location services without enabling worms, what form of monitoring and analysis of file server traffic is appropriate and effective in detecting misuse and insiders, and how to develop an appropriate policy for a secure LAN.

The second problem, scaling IDS with network capability, requires a complete parallel restructuring of how intrusion detection systems operate. Fortunately, the IDS can introduce considerable latency by hardware standards (up to 1-2 microseconds) which are practically unnoticeable by networking standards. This mismatch enables "wide, slow, and parallel" hardware designs which are optimized for many-task throughput rather than single-process latency.

This will require recasting the IDS pipeline as a parallel computation, with many hundreds or thousands of simultaneous instances which perform work on a per packet, per flow, or per host basis. By reordering how these operations proceed, it should be possible to divide the working set across multiple memories and provide for better memory access, allowing the very stateful computations to efficiently utilize many high-bandwidth, high-latency memories. This will require both new software (to automatically parallelize the IDS analysis pipeline) and new hardware.

If successful, this would enable the IDS system to scale as networks scale, by providing more processing and memory as networks continue to increase in bandwidth. Without the ability to parallelize the IDS, IDS evolution will hit an insurmountable barrier, for both traffic volumes and attack sophistication will greatly surpass the available single-thread processing of commercial microprocessors.

d: The biggest problem in evaluating intrusion detection systems is not testbeds but developing good test data. Furthermore, the problem is usually not in capturing attack traffic, as attack traffic can be generated and injected, but capturing realistic background traffic, ***with payload information***.

The simple approach, simply recording all traffic on a major network, is politically impossible. The privacy implications often prevent doing this for internal data, and it is usually impossible to export this data. This prevents reproducibility, apples-to-apples comparisons, and introduces significant biases in the results.

What is needed is exportable datasets of LAN and WAN traffic, complete with full payload information. There appear to be two possible strategies: synthetic networks with synthetic users (like the Lincoln Labs dataset) or tools which can, by using full semantic analysis, anonomyze the payload data while still preserving its utility in IDS evaluation.

Although subject to much debate, a Lincoln Labs style dataset, where a large number of real systems are coupled together and driven by synthetic models of user behavior, can create a reasonably effective background data set. It would require considerable effort to construct a new dataset in this manner, because of all the myriad user-behavior, P2P systems, file sharing protocols, and non-traditional computing devices such as VOIP phones, SCADA systems, and workgroup printer/copier/scanners which form the modern enterprise network.

Techniques to anonymize payloads from real traffic, although more difficult, may prove in the end more valuable. With synthetic datasets, there is always a question of how realistic the traffic model really is. With tools to anonymize real traffic, there is no longer a question of network realism, only representativeness.

# Research Challenges in Engineering Scalable FIT Systems- Christof Fetzer

Prof. Christof Fetzer, Ph.D.

TU Dresden, Germany

One of the most important goals in the domain of dependable systems research is to be able to engineer systems with a system dependability of 1 FIT, i.e., systems with at most one system failure in about $10^9$ hours. To be able to achieve FIT systems, we need to partition the problem such that each sub-problem becomes solvable, i.e., we need to use a divide and conquer approach. A good top-level systems engineering approach is to make sure that a component, which is implemented by some program P, satisfies the following properties:

I. The program P is correct

II. The program P is executed correctly

III. The correct program, i.e., P, is executed

IV. Our (environment) assumptions are correct

Such a division of responsibilities identifies a set of research challenges that need to be solved. First, we need to make sure than all four properties are probabilistically guaranteed, i.e., we can neglect the probability that any of these properties are violated more than once in $10^9$ hours of operations. I will address three concrete challenges in this context:

1) Guaranteed failure virtualization

2) Guaranteed sandboxing

3) Guaranteed patching of code

**Guaranteed failure virtualization** ensures property (II), i.e. the correct execution of programs. The shrinking feature sizes of future chips will result in (a) an increasing variability of transistors, (b) increasing soft error rate and (c) accelerated aging of transistors. Therefore, one has to expect that for feature sizes of 22nm about 30% of the transistors might not operate as designed. The likelihood of undetected failures can therefore increase above 1 FIT even if standard techniques like lock-stepping of cores are used. For FIT systems we need to detect all wrong executions with a false negative rate below 1 FIT.

Guaranteed failure virtualization is important for the scalability of systems in the sense that it permits the use commodity hardware for 1 FIT systems. For geographically distributed systems without physical access control, the guaranteed failure virtualization must be able to detect the wrong execution of programs executed on modified hardware, i.e., hardware modified by some attacker of the system.

**Guaranteed sandboxing** addresses requirement (III), i.e., it helps to ensure that the correct program is executed. One problem is that some software bugs in a program can permit attackers to inject code in a running program. Some of these issues are caused by the need to reuse existing code written in unsafe programming languages.

However, even safe languages permit intrusions, e.g., hardware failures like bit flips can be used to circumvent type safety.

Therefore, we need a guaranteed way to detect and terminate injected code before this leads to a system failure.

**Guaranteed patching** addresses requirement (I) and helps to make sure that the executed program is correct. Having guaranteed sandboxing permits to detect all code injections (modulo some negligible probability). While this prevents the propagation of malicious failures, it still can be used to decrease the availability of components and potentially, of the system. Therefore, we need to make sure that bugs are patched immediately to prevent repeated intrusions (that will be detected by reduce the availability). The probability that one cannot patch a bug needs to be negligible and the probability that we prevent correct requests from being processed needs to be negligible too.

# Large Scale Routing Experimentation for Next-Generation Networking - George Kesidis

George Kesidis (speaker) and Glenn Carl

Pennsylvania State University

{gik2,gmc102}@psu.edu

Supported by the DHS/NSF EMIST project, http://www.isi.edu/deter

Dependability and Security of Future Networked Systems Panel (B):

Scalability and Context-Awareness

1. Motivation for large-scale testing

Both real and hypothetical problems exist for BGPv4. These include slow convergence, policy-induced instabilities, ease of misconfiguration, and security vulnerabilities. Significant research has already been conducted exploring solutions. These include recommendations for BGP parameter settings (e.g, MRAI), new protocol attributes (e.g., communities) and routing policies, BGP overlays, and clean-slate control plane architectures. Much of this research has not been widely deployed because of unknowns that such techniques may induce at large-scales. This was seen with route flap dampening, as there were increases in convergence time with larger scale deployments. Performance improvements of a given technique need to be significant enough, and worthy enough under increased deployment, to fuel further deployment. The techniques under evaluation can be slight modifications of existing routing protocols or completely different alternatives for a possible next-generation (not necessarily IP-based) inter-network. In today's sober economic climate, there is a need to test (i.e., "engineer") these approaches at scale, prior to even just an incremental deployment, and to test the strategy of incremental deployment itself.

2. Challenges

BGP is a globally distributed protocol operating between more than 20,000 autonomous systems, each having a diverse policy framework governing them. BGP testing at such large scales, subject to routing policies, is difficult. Most straightforward techniques, using laboratory testbeds, simulation or analytical theory, are either incapable or highly inefficient under these conditions. New techniques are needed to reduce the BGP testing problem down to a reasonable scale and to improve the power and flexibility of testing facilities.

3. Existing Testbeds

There has been significant BGP experimentation through general simulation tools like SSFNet, NS-2, OpNet, as well as the special-purpose C-BGP platform. NS-2's add-on BGP support does not support both routing policy and IP forwarding table population concurrently, restricting its use for combined data-plane and control-plane studies. Furthermore, NS-2 is generally cumbersome when attempting to increase its

scale through parallel simulation. A new NS-2 module for a router is now under development by Fahmy et al. of the DHS/NSF EMIST project team.

SSFNet is not architected for large sized networks, and does not have distributed simulation support for BGP. However, SSFNet can be modified to pre-calculate its BGP routes, thus increasing its scalability by reducing BGP run-time efforts. Such an approach does not allow for detailed visibility of the protocol in action (i.e., no exchange of routing messages), thereby limiting its use for studying protocol dynamics such as responsiveness and stability.

C-BGP specializes in solving for BGP routes, thus avoiding the burdens of processing simulation events. It can achieve Internet scales, while including routing policy, but may lack in its extensibility to integrate new routing protocols.

Similarity, there are several options for BGP emulation. The VINI project plans to develop light-weight router emulation software suitable for time-slices of a PlanetLab workstation. There also exists sophisticated hybrid simulation, emulation, and prototyping frameworks, such as Emulab-based DETER and WAIL, that can include actual routers as well as XORP/Zebra software routers running on dedicated, multi-NIC machines.

Many tools also exist for processing of measured BGP data (e.g., Routeviews and RIPE), and for other experimental specification and visualization functions, in different stages of development. Current versions of these tools are available from the DETER experiments workbench at http://www.isi.edu/deter

## 4. Preliminary and Current Work

Realistic scaled-down models are needed to increase simulation efficiencies and make hybrid simulation-emulation experiments feasible. The fidelity of scale-down could consider both graphical and dynamical metrics. These include power-law distribution parameters (e.g., for node degree) and those specific to BGP (e.g., MOAS and convergence dynamics) respectively. However, care must be taken so that when scaling-down so that one does not lose the intent of the routing policies.

We are currently exploring scale-down techniques that preserve policy-influenced BGP routes via Thevinin/Norton-theorem analogies from linear circuits. Simple experiments have been conducted where, beginning with a Routeview-inspired topology, we study the “natural” response of BGP’s route creation. For example, at the start of the experiment, advertisements are inserted to create a MOAS problem, and then observations are made regarding BGP created routes assuming no other “external stimulus” to the protocol during its course of execution. We compare these responses between networks of many reduced sizes and find similar traffic engineering.

Currently, we are also developing “forced response” experiments, i.e., joint data-plane and control-plane simulation, similar to the SSFNet experiment jointly simulating worm and BGP activity. Specifically, we are synchronizing a time-driven scaled-down data-plane simulation derived from a trace of the SQL Slammer worm to a correspondingly scaled-down event-driven BGP simulation.

# Dependability & Security of Future Networked Systems – scalability and context-awareness – Gérard Le Lann

Dr. Gérard Le Lann, Team NOVALTIS, INRIA, Rocquencourt, France

http://www.inria.fr/recherche/equipes/novaltis.fr.html

Define 1 or 2 difficult problems (research challenges)

**Scalability and context-awareness** (of our future networked systems that shall be dependable and secure) can be achieved if and only if such **systems are *designed* properly, i.e. in a *provably correct* manner.** Rationale is as follows.

Context-awareness implies:

- The ability to maintain (on-line) knowledge of "current" state(s) of (1) a system, (2) its environment,[2]
- The ability to make (on-line) correct decisions based upon such knowledge. Almost inevitably, correctness implies timeliness ("real-time"), "short" reaction/response times more precisely. Hence, in many real settings, context-awareness implies autonomy. Since the complexity of our networked systems has long ago exceeded the capabilities of human brains, human-based on-line diagnoses (bound to be correct) are out of question. As an example, our networked systems must be equipped with the necessary level of intelligence that would make their components able to, e.g., reach time-bounded agreement (within dynamically created groups of components). Proving this type of capability, as well as gaining a priori knowledge of correct upper bounds for reaction/response times implies meeting proof obligations.[3]

**Scalability** implies the ability to perform on-line and/or off-line reconfiguration, re-dimensioning, etc. of a system/solution which has been performing satisfactorily with $n$ components, so as to encompass $N$ components, $N >> n$. Without analytical expressions of what is called "feasibility conditions" or "worst-case lower bounds", as functions of $n$, it is simply impossible to demonstrate scalability properties a priori. A posteriori "verification", i.e. based upon observing a system which has been fielded and in use "long enough", is not acceptable in many real settings. Analytical expressions imply meeting proof obligations.[4]

At the dawn of the 21st century, we know how to provide ourselves with two different kinds of systems (simplified view):

[2] Related problems are studied in the area of distributed dependable networking/computing—see, e.g., the proceedings of the ACM Conference on Principles of Distributed Computing (PODC), the IEEE Intl. Conference on Distributed Computing Systems (ICDCS).

[3] Related problems are studied in the area of real-time networking/computing—see, e.g., the ACM/IEEE Journal on Networking, the proceedings of the IEEE Real-Time Systems Symposium (RTSS).

[4] Comments under footnotes 1 and 2 apply here.

- General-purpose systems, developed and maintained by multiple actors, used by an unknown numbers of people and/or devices, which systems are bound/expected to deliver "sufficiently good" services along with some QoS "most of the time"; almost inevitably, design and validation/verification techniques that apply are drawn from stochastic disciplines,

- Specific systems, notably those carrying the attribute "critical", developed and maintained by actors selected out of a small set (of eligible actors), used by entities (humans, devices) belonging to a well defined/authorized set (albeit subject to e.g. intrusions); design and validation/verification techniques that apply are drawn from "deterministic" disciplines (e.g., logic, combinatorial optimization) as well as from stochastic disciplines (e.g., coding theory).

We know reasonably well how to validate/verify *software programs* when considered in isolation from each other, as well as some program compositions, assuming almost "ideal" operational conditions (zero-time or constant-time step abstractions, no "serious" failures/attacks, no queuing, etc.); model checking, "synchrony" paradigms (languages, computational/system models, tools) belong to this category.

**The difficult problem we have been facing for years, since the inception of "formal" methods—essentially directed at software concerns so far, without making tangible progress, is related to validation/verification issues as they arise in more general, hence more realistic, settings, i.e. how to validate/verify:**

**-- *(sets of) software programs* under general assumptions, in particular considering asynchronous paradigms (further "away" from synchronous models than the GALS paradigm),**

**-- *(general-purpose, specific) systems*, that is:**

*** Proving that a collection of components (a component being a system whenever "system" means "system of systems"), in the presence of varying computational and communication delays, queuing phenomena, varying densities of "loads", failures/attacks occurrence, etc., behaves globally as specified, under specified operational/environmental assumptions which meet some imposed coverage figure (usually bound to be (very) close to 1),**

*** Such proofs being established irrespective of choices (which choices may evolve with time) made regarding which technology is resorted to for implementing each of those components under consideration (software, hardware, optronics, mechanics, etc.).**

The following is becoming clearer every day:

- With few exceptions, we have been overlooking the "system" dimension in many disciplines, dependable/secure computing/networking included (!),

- We have developed an extraordinary infatuation with software.

For example, as for dependability, most studies published on causes of failures or quasi-failures—and their remedies— focus on software (number 1) and humans (number 2). Unfortunately, albeit non intentionally, too often, such analyses reveal a confusion between consequences and causes. Granted, when "something bad" happens with or within a system, some software is running. It is then way too easy to designate "the software" as the culprit. But quite often, "the software" is just a "victim" of faults rooted into lifecycle phases that precede an implementation/deployment phase, typically the "requirements capture" and the "system design & validation" phases. This is demonstrated by a growing number of analyses of project cancellations and/or operational system failures in various domains, such as, e.g., the

US public telephone network (NIST), AAS (FAA), some Mars missions (NASA), On-line Libraries or Horse Race Betting (France), Ariane 5 Flight 501 (ESA).

Therefore, any attempt at "improving" the trustability/dependability of our systems by "improving" **solely** software engineering methods, tools, as well as their theoretical foundations, is inappropriate or, at best, marginally interesting, *from a practical angle*, since the major roadblocks lie with system engineering.

An analogy with Civil Engineering may be appropriate here. Consider a building. So far, we have concerned ourselves with how to master the development of blocks of reinforced concrete (hardware components) as well as satisfactory customized apartments/offices (software components). But in our field, where do we see the equivalent of overall building blueprints (overall specifications of systems), which must be validated (tools, experts) prior to starting costly construction activities (implementation, development, fielding)?

**System engineering (of networked systems) has become the weakest link in projects lifecycles or undertakings directed at such systems. Since proof obligations are central to the security/dependability of such systems, one of our very major and difficult challenges lies with proof-based system engineering**.

*For each problem, describe the impact of solving or failing to solve the problem and a description of specific challenges making the problem difficult.*

Unless we start an international effort directed at proof-based system engineering (of networked systems), we will keep experiencing disillusions, and maybe some very severe damages that might be irreversible, due to the intrinsic weaknesses "built in" the designs of our systems.

Such an effort would echo the effort introduced and assessed by Tony Hoare in the Journal of the ACM, January 2003 ("The Verifying Compiler: A Grand Challenge for Computing Research" (≈ 1000 man-years of research effort, spread over 10 years or more)). Tony Hoare has stressed the importance of addressing *formal software engineering* issues better than done so far. We should provide ourselves with an agenda for addressing ***proof-based system engineering*** issues (much) better than done so far. The level of effort should be in the same range as the level quoted by Tony Hoare.

Why "formal" with software engineering and "proof-based" with system engineering? Simple:

- models needed to correctly "capture" the true semantics of many real networked systems (as well as applications they support), especially those deployed in the future,
- proof techniques that are applicable to such models,

are not tractable with existing "formal" methods/tools. Hence, proofs are established by resorting to Mathematics and appropriate "hard sciences". Their "formalization" (their mechanical verification) can occur only when tools have been developed. When this happens, then we will be allowed to use the term "formal system engineering".

Also, there is a fundamental reason why formalism/formalization cannot encompass fully a project/system lifecycle, and that is the requirements capture phase. By definition, the inputs of such a phase are documents written, interviews conducted, in some natural language. Therefore, the work involved with arriving at a formal specification at the end of such a phase inevitably starts with a translation into a semi-formal specification, in some restricted vocabulary of a natural language (only terms that have well known and/or formal definitions are eligible). Existing "formal" methods do not address this type of work.

Some of the specific challenges are as follows:

- Formal definitions/semantics of models matching (a) desired functional as well as *non functional* properties (timeliness, dependability) as they arise with real-world applications, (b) the characteristics or those technological components (processors, memory/storage, communication channels, etc.) used to build our networked systems,
- Construction of the knowledge bases (dictionaries, validated solutions, etc.) that would encapsulate scientific results (meaningful to networked systems),
- Development of tools that would make such knowledge bases one click away from engineers (scientists and teachers as well, for educational purposes),
- Orientating the research work conducted by some communities so as to address currently open problems, namely composite problems in distributed real-time dependable/secure networking/computing (see further).
- Abstractions and operational paradigms that would permit to synergistically combine achievements in the AI domain (application-level semantics, goal-based designs) and achievements in the distributed real-time dependable/secure computing/networking domains (see further).

*Provide potential research directions for handling the problem, and metrics for measuring progress against the problem. Possibly, describe also the limitations of current approaches, promising approaches not yet fully explored and desirable approaches where relatively few ideas have been proposed.*

A very first and much needed step would consist in fostering cross-fertilization between scientific communities that tend to be working in isolation from each other (for good reasons, most often, such as lack of time). A simple example may be useful.

In many instances, the time dimension (either global time or durations (of time intervals)) plays a central role in solutions directed at security and/or dependability issues. Then the following question: uncertainty being an inherent attribute of operational conditions under which our future systems will operate (to the exception of a few very specific systems),[5] how can we predict "correct" values for timers, distributed clocks, etc. in the presence of uncertainty? **IOW, how can we prove timeliness properties in distributed synchronous, partially synchronous, asynchronous, systems prone to partial failures/intrusions and/or transient overloads?**

The **queuing theory** community (ACM/IEEE Networking) has responses to this question for some problems arising with communication networks. The **scheduling theory** community (IEEE RTSS) has responses to this question, mostly for computing systems that are not distributed, nor prone to failures/attacks, nor prone to overloads.

Conversely, the **distributed computing** community (ACM PODC, IEEE ICDCS) and the **secure computing** community have responses to questions arising with distribution, failures, intrusions, etc. in

[5] US GPS or EU Galileo are not subject to uncertainties such as unknown "loads". Nevertheless, according to a recent study from Cornell University, solar activity in 2011-2012 shall reach intensity levels unknown so far, and radiations will inevitably disrupt or disturb very seriously the functioning of all systems that rest on satellite positioning.

synchronous, partially synchronous as well as asynchronous systems, under the form of proofs of safety/liveness properties (a dependability/security property being a combination of such properties), but they ignore the fact that such proofs (e.g., optimal lower bounds) are invalidated whenever scheduling problems arise—inevitably the case with real networked systems.

Not much is known in the area of **algorithms (and proofs) for distributed *and* real-time *and* dependable/secure computing/networking under non restrictive assumptions**.

The "blending" of these diverse scientific cultures is much needed.

In particular, **synchronous assumptions/paradigms have very limited applicability** in this area (their coverage factors cannot be high enough), given that our future networked systems must be *dynamically* scalable, self-adaptable, self-regulating, in the face of constantly changing operational and/or environmental conditions, while still being predictable.

Hence, it would be extremely useful to stimulate some fundamental research work (aimed at these problems) targeting **partially synchronous or asynchronous system/computational models (those matching the very nature of our future networked systems) along with realistic failure semantics**. For example, so far, with very few exceptions, the ACM PODC and the IEEE ICDCS communities have been studying *permanent* failure models (once erroneous, a component is considered failed forever). In many real settings, *transient* failures (crash, omission, timing, value, failures) are the dominating type. Not much is known regarding *system-level* algorithms, optimality proofs, impossibility results, for such failures in partially synchronous or asynchronous system/computational models.

Adversaries of our future networked systems are extremely powerful: constantly changing operational and/or environmental conditions translate in multiple possible patterns/scenarios, which is why context-awareness is necessary. Our traditional approaches to designing networks/systems can only fail, since basically they boil down to endowing each of our networks/systems with a single set of protocols/algorithms. In so doing, we can only prove that our networks/systems "win against" a limited subset of the patterns/scenarios that can be deployed by an adversary. Our future networks/systems must be equipped with *families* of protocols/algorithms, as well as with the ability to dynamically select in a family which set of protocols/algorithms shall be enacted at any time, according to currently observed context.

Goal-based designs are necessary. Also, we shall establish the theoretical frameworks—still missing—for viewing and reasoning about processes as "socially intelligent" entities, i.e. that may be endowed new roles and "higher" responsibilities (within networks/systems bound to be autonomous) as time elapses.

The "blending" of the AI culture and the Computer Science/Control Theory culture—a goal that has been around for a while, without widespread tangible success yet—is needed.

One last remark is related to languages. Languages are not a panacea for the most fundamental problems and challenges that we are faced with, such as those presented above. Nowadays, UML and some other model-oriented or architecture-oriented description languages are being marketed (in some circles) as "the" solution for system engineering concerns. This cannot be true (and has never been in the past, for any language one may think of). No language can eliminate the necessity of conducting those mathematical exercises involved with proving that an algorithm/protocol specification solves a certain problem in distributed real-time dependable/secure networking/computing, or with establishing worst-case feasibility conditions. For example, in the case of distributed systems, distributed processes modelled as finite directed graphs, every possible execution of any such process being bound to terminate within a strict deadline despite failures, one may find it appropriate to use matrix calculus in (max,+) algebra in order to prove timeliness properties. Clearly, this is something totally foreign to UML or any other language.

*Describe common research resources that the research community could create to facilitate the research described above*

- A common method directed at proof-based system engineering (PBSE), i.e. theoretical foundations of PBSE, a PBSE-compatible lifecycle model, …
- Knowledge bases (dictionaries, validated solutions, etc.) that would encapsulate scientific results:
    - models of systems/networks, failures, event arrivals, processes, …
    - specifications of problems and solutions, proofs, feasibility conditions, …
- Development of tools that would make such knowledge bases one click away from engineers (scientists and teachers as well, for educational purposes) as well as from "intelligent processes" in an autonomous system. For example, dependability and timeliness feasibility conditions for a given pair {problem, solution} being instantiated as a set of analytical constraints, it would be necessary (and relatively easy) to implement such constraints as a program, to be run by a tool. This is how provably correct scalability can be automated, either off-line or on-line—the program and the tool being an integral part of an autonomous system in the latter case.
- A common repository and/or a moderated forum for in-depth analyses of mishaps, accidents, project cancellations or difficulties related to critical information infrastructures, embedded systems, and so on, the goal being to gain a common understanding of the real nature of the dominant causes of "problems" experienced with our networked systems, now and in the future. Such an undertaking would resemble the existing RISK forum (SRI) and the Safety-Critical forum (York University), although the focus would be put essentially on Formal Software Engineering vs. Proof-Based System Engineering, in order to "prove" or "disprove" the conjecture according to which faulty System Engineering has become the dominant cause of "problems".

# 21st Century Information Security R&D Challenges: A Large Global Enterprise's Perspectives - Ming-Yuh Huang

Ming-Yuh Huang

The Boeing Company

ming-yuh.huang@boeing.com

The Enabling Picture

As a large global enterprise, The Boeing Company operates one of the largest computing infrastructures in the world executing global manufacturing, distributed collaborative engineering, and massive virtual enterprise integration. Additionally, Boeing also builds highly complex large-scale defense and government System-of-Systems (SoS). Within this context, information/cyber security is more than just securing assets and keeping people out; it's very much about letting people in - the right people, the right time, to the right resources. Modern social and business practices require us to work much more closely than before. Consequently, information security has become the key "business-enabler" to propel the next-generation paradigm shift. Without this enabling security technology, tight integration demanded by the new business models will be impossible. Traditional view of treating information security as a protecting and prohibiting technology is out of date.

Formal Authorization

Under this enabling concept, one of the most critical issues lies in the area of entitlement and authorization. Increased connectivity inevitably leads to combinatory explosion of mapping from large users/applications permutation sets to even larger set of data and resources. MAC (Mandatory Access Control, e.g. multi-level security) is critical to military/high-security systems. Additionally, there is also an urgent need for fine-grain, scalable, dynamic, and manageable DAC (Discretionary Access Control) technologies such as policy-enabled Role-Based Access Control (RBAC) to enforce the complex operation/business/IP need-to-know basis authorization logics. As in the software engineering life cycle, there is a great need for formal authorization requirement modeling, automated authorization policy generation (from the requirement models), and consistent policy application across multiple heterogeneous platforms within SoS. Today's systems encode authorization logic in low level "implicit" bits-and-bytes. Such infrastructures do not scale and cannot respond to the highly dynamic requirements within the next-generation integrated systems. Furthermore, the concept of a federated authorization framework is also needed where authorization ownership & logic can be managed and enforced in a realistic & distributed manner.

Autonomic System Health/Trust Management

There is a great need of "system-level" technology where security requirement, design, implementation and operation are managed at the overall level. It fuses data from local sensors, monitors the overall system health and performance; and with the aids of modeling and simulation, performs intelligent

diagnosis, prognosis, preventive maintenance, and eventually system self-healing. Today, many stove-pipe technologies exist but they do not come together within the overall system framework. And, as a result of more and more system integration, today's large system security often consists of merely aggregation of sub-system security. The reality is the sum of parts is not equal to the whole in this case. System level autonomic system health/trust management technology is critical for future large systems.

Considering & Protecting the New Paradigms

Security has always been a catch-up game – people build, people break, and then people fix. It has always been a passive and reactive. There has not been sufficient effort or deep thoughts on where computing is leading our societies and our industries toward and the security implications come along with it, as well as the gap analysis to reach there. To break up today's "catch-up" cycle, information security needs to have a sense of the future of computing and then articulate/innovate security options to prevent the pitfall and lead the road of safe, trusted computing and systems. Fixing yesterday's security challenges will not prevent future challenges from coming. It is important to be vigilant on current challenges. However, we need to spend more effort investigating future computing/architecture/social/business paradigms and their relevant dangers and risks.

Security in Context

Finally, today's information assurance approaches lacks the distinctions of abstraction layers. Real-life systems contain at least four layers of abstractions – infrastructure (computer & networks), applications (software), services (grouped and purposed applications), and business processes (operation & transactions). There are many "intra" and "inter" layer relationships (e.g. hosting of an application on a server, dependency of a messaging-backbone service on an anti-virus application) amongst various components in this hierarchy. Faults and vulnerability could reside in any of the layers. Thus, overall system health management cannot be achieve without good understanding layers and intra/inter relationships. Today's security technology focuses on low layer protection (network & systems) while the society is moving toward higher level processes (transaction). Local security measurement can only achieve limited result. A comprehensive approach encompassing all levels of hierarchy is the only solution to achieve totally system trust.

# Dependability & Security of Future Networked Systems – scalability and context-awareness

**Sandro Etalle**

**sandro.etalle@utwente.nl**

**Define 1 or 2 difficult problems (research challenges)**

The problem is to enforce end-to-end security policies in heterogeneous systems effectively and efficiently.

In networked systems data is manipulated (generated, composed, altered etc) by different peers often belonging to different security domains (e.g., deployed to different service providers). When such information is sensitive, the illegitimate alteration of the information flow - whether intentional or not - is an urgent problem, raising the necessity of establishing and enforcing information *flow policies* which determine who is entitled to do what on which piece of information. Unfortunately, the enforcement of such policies is extremely problematic, as heterogeneous systems often do not offer a security monitor which is trusted by all peers. The challenge is thus finding methods giving reasonable (if not absolute) assurance that the expected security policies are actually enforced. This may be achieved by using *detective* methods in addition to the standard *preventive ones.*

The concrete technical challenge is *integrating preventive and detective methods for the enforcement of end-to-end security policies in collaborative environments.*

**For each problem, describe the impact of solving or failing to solve the problem and a description of specific challenges making the problem difficult.**

This problem of end-to-end enforcement of security policies arises in various contexts. Here are two examples.

1. Networked embedded system need to exchange information to function properly. This information needs to be reliable: incorrect information is often worse than no information. The problem is exacerbated when the information flow needs to be dynamic as well as ad-hoc, being composed as and when required. Indeed, information acquires value through the transformation process. The transformation nodes may be people, hardware devices such as sensors, or even software agents. To guarantee interoperability without having to resort to proprietary solutions we need a way to enforce end to end security policies also when the information moves *across* security domains.

2. In an inter-organizational cooperation, data has to be moved across security perimeters, i.e. from one organization to another. To protect the intellectual property of the companies involved, special agreements determine which data can be exchanged, who are the people entitled to access it and for which purposes the data should be used. Yet, there exists no automatic mechanism for checking whether the data is actually used in accordance with the agreements.

Present trust-management systems allow to the *trustworthiness of peers*, yet they cannot provide any *degree of assurance* that the policies are actually enforced and do not offer a way of detecting nor reacting to policy infringements.

**Provide potential research directions for handling the problem, and metrics for measuring progress against the problem. Possibly, describe also the limitations of current approaches, promising approaches not yet fully explored and desirable approaches where relatively few ideas have been proposed.**

To meet this challenge, we believe that we need to integrate a-priori methods (e.g., standard access control techniques) with *a-posteriori measures.* In other words, the innovation aimed at by this challenge is integrating the problem of preventing illegitimate use (which in situations such as collaborative environments is often unsolvable) with the problem of *deterring illegitimate* use (which is a major challenge to solve, and which is very worthwhile). The result should be a uniform framework for compliance control, in which e.g. access control and accountability-based measures are seamlessly integrated.

This challenge is primarily of technical nature, but it has a multidisciplinary character; first because deterring measures do not rule out illegitimate behaviour. It is therefore important to link these measures with an appropriate *risk management and risk mitigation strategies*. Secondly, deterring measures cannot omit taking into consideration the human and organizational factors.

# Evaluating the Dependability & Security of Networked Systems – modelling, simulation, predictive evaluation, assurance cases – Andrea Bondavalli

**Andrea Bondavalli, University of Firenze**

Dependability and security have become a crucial requirement of current computer and information systems, and of large-scale infrastructures, the construction thereof are witnessing at an always faster and faster rhythm. These infrastructures integrate previously disjoint systems to provide services, which have become critical in our everyday life. They need to survive failures of components or subsystems, as well as attacks and intrusions and under no circumstances a complete denial of service is acceptable for them. Instead, some level of (reduced) functionality must be provided also in the event of faults. On the other hand, the massive deployment of low-cost, relatively unstable COTS (Commercial Off-The-Shelf) hardware and software components, as well as the integration of old legacy subsystems, seriously undermines the possibility of meeting stringent dependability requirements, especially for long running applications.

These factors have extended the field of application of dependability computing much beyond the traditional one (i.e., critical applications, such as nuclear, space, transport, telecommunications, etc.). Without efficient architectural solutions, dependable COTS and legacy-based applications, as well as large survivable infrastructures are virtually impossible to obtain. In this context, the capability to detect errors or component failures, the ability to correctly diagnose the status or health of components (including entire legacy subsystems), and the availability of quick and intelligent reconfiguration and fault treatment strategies are key system features.

Here, we focus on the diagnosis step, essential in the definition of fault tolerance strategies, aiming at determining the cause of errors, both in location and in nature. The growing complexity of modern and future systems exacerbates the necessity of diagnosis mechanisms to cope with system and environment faults. Because of the intrinsic characteristics of such modern infrastructures, often including COTS and legacy components, diagnosis is a complex and delicate activity, which requires a careful assessment of the status or the extent of the damage in individual components.

Issues to be tackled, that make diagnosis a critical task, include:

i) the limited knowledge and control over the system as a whole, as well as over individual components, by the system designer,

ii) the large grained peculiarity of components on which diagnostic activities must be conducted, whereas traditional applications typically consist of relatively fine grained components;

iii) heterogeneity of the environment under analysis, whereas the targets of traditional diagnosis are - to a large extent - homogeneous.

Thus, it is not practical, as soon as an error is observed, to declare the entire component failed and proceed to repair and replacement. Traditional one-shot diagnosis is thus inadequate: an approach is needed, which collects streams of data about error symptoms and failure modes and filters them by observing component behavior over-time.

In a wider perspective, diagnosis needs to assess the suitability of component or subsystems or infrastructure to provide services with adequate quality of service, which may dynamically change over time. In this view, the goodness of a component is not strictly tied to the absence or presence of faults which may impair its functionality; rather, it is the overall quality of service which determines whether a component is useful and contributes to the system activities or it is better to keep it out.

Actually, this more general framework allows capturing several possible scenarios, such as:

i) the component QoS decreases because of malfunctions affecting the component itself,

ii) the application using the component changes the QoS requirements in such a way that they do not match anymore with the specification of the component under utilization, or

iii) changes in the environment (e.g., system load) may lead to a change in the QoS provided.

Component obsolescence is a typical example of case ii), while classical examples of case i) can be taken from the system fault tolerance area. Current and future systems and infrastructures are going to be more and more characterized by heterogeneous components, used for disparate applications which may change their requirements during the lifetime, thus posing the problem of monitoring system components to assess when services with unsatisfactory QoS is going to be released, and for how long.

An over-time diagnosis is required, to cautiously understand when a component is no more beneficial against withdrawing still useful ones. Such dynamic analysis must be performed on line to provide systems and infrastructures with means to flexibly and dynamically react to adverse conditions.

Several diagnosis methods have been proposed, which are mainly based on heuristic approaches. However more general frameworks and approaches are desirable to be able to provably assess the diagnostic accuracy and completeness. A few formalisms such as Bayesian theories, Hidden Markov Models, Causal Probabilistic Networks appear promising for such goal and worth investigation on how well they can model such over-time diagnosis

# Panel C

## Security & Privacy in Dynamic Wireless Networks

# Some Security and Privacy Challenges in Wireless & Mobile Networks- Gene Tsudik

**Gene Tsudik**[6]

Security in wireless & mobile networks has been a popular area of research in both security and network research communities. Among more mature security research topics are: secure routing, distributed certification, DoS attack countermeasures and secure location verification. Privacy pops up in a number of different aspects some of which have been recognized and studied extensively while others are only now moving into the foreground. We will discuss a few of these below.

Privacy and Anonymity in MANET Routing Key features desired from routing include

effectiveness, efficiency, robustness and security. Privacy is not a typical concern in secure routing. This is natural since routing is a fundamental network service and privacy seems to contradict its very purpose. However, in some settings (e.g., law enforcement) privacy in routing is an important issue[7].

Why is confidentiality in routing important? First, because mutually suspicious MANET entities might take part in a common routing protocol. For example, a single network may include nodes from different countries or with different privileges. Second, some nodes maybe captured or compromised and it might be very hard to quickly identify them. It is thus important to minimize the information available to all nodes, while enabling effective routing. Sensitive information includes long-term identities of nodes as well as their movement patterns. A number of novel and challenging research issues arise in the context of privacy for MANET routing protocols.

There have been no solutions proposed thus far. A major stumbling block is the need to modify routing protocols and to amend the notion of node identities. Privacy in Last-HopWireless Networks A similar (but much more obvious) privacy problem occurs in last-hop wireless networks such as cellular (notably, GSM) and 802.11-style hotspots. As a node migrates from one cell (hotspot) to another – whether in real-time or not – its movements are easily traceable since it, as a rule, identifies itself via some permanent means, e.g., IMSI or MAC address. This problem is not new; it has been studies since mid-90s in the context of GSM; however, no practical impact has been felt. It was re-discovered in the 802.11 context and some solutions have been recently proposed. However, they all require the ability to mask or change node identities. One solution is to pressure the manufacturers of network interfaces to abandon permanent interface/node identities. Another (less likely to succeed) is to develop better work-arounds. Privacy-Preserving Authentication In some suspicious (or hostile) ad hoc dynamic network scenarios, two or more nodes needs to establish secure and private (with respect to both outsiders and other network nodes) means of communication. One network may include nodes corresponding to all types of groups and roles and

[6] Department of Computer Science, University of California, Irvine. Email: gts@ics.uci.edu

[7]Note that the challenge is not how to keep routing information private from outsiders; this is easy enough to achieve with well-known cryptographic techniques. Rather, the challenge is the privacy of routing informationvis-a-vis network nodes themselves.

some nodes may belong to a group of secret agents who need to be aware of, and communicate with, each other. This must be done without anyone else (including other such nodes) recognizing their presence.

The main idea is the need to authenticate members of a certain group in a way that is unobservable by anyone else. Standard authentication techniques are unfortunately unsuitable since they all require either the existence of pre-shared secrets or open exchange of credentials, such as Public Key Certificates or PKC[8]s. There is, therefore, a need for a special type of privacy-preserving authentication, also known as ”secret handshakes”. Some viable approaches have been suggested. However, most rely on the use of one-time credentials which makes them clumsy or, at least, unscalable.

Secure Membership control is an essential and fundamental security service in a decentralized network settings, such as a MANET. It is needed to securely cope with dynamic membership and topology as well as to bootstrap other security primitives (such as key management) and services (such as secure routing) without the assistance of any on-line centralized trusted authority. An ideal admission protocol must involve minimal interaction among nodes, since connectivity can be unstable. Also, since MANETs are often composed of resource-limited devices, admission control must be particularly efficient in terms of computation and communication.

There are differences between long-term and short-term groupings (or associations) of nodes. The former require membership certification and may justify more expensive admission techniques, whereas, the latter might not require the same heavy-weight measures. Some prior research considered both long- and short-lived networks, with mixed success. Current techniques are particularly inefficient for long-lived networks. Also, a missing component in prior work is secure distributed node eviction. Admission and eviction are (surprisingly) not very similar and call for different techniques. Node admission is prompted by a prospective ”member” who triggers the admission process. Eviction, in contrast, is triggered by an existing node(s) which needs to find enough other

nodes to sponsor eviction. A newly-admitted node can willingly prove its membership by producing a membership proof, whereas, an evicted node has no incentive to prove that it has been evicted. The burden of propagating information about evicted nodes falls upon nodes that have taken part in the eviction process. This presents some challenges, e.g., how to to distribute eviction information and how to minimize its size.

[8] Other more sophisticated approaches, such as enveloping an authentication session with an ephemeral key, also fail, since an active adversary can always violate confidientiality of authentication.

# Security & Privacy in Dynamic Wireless Networks – Roberto Baldoni

*Define 1 or 2 difficult problems (research challenges)*

**Privacy and trust in dynamic coalitions**. Pervasive and ambient ICT has great potential in areas such as health, finance, education, defence, entertainment, and transport. The basic building block of this development is wireless communication which allows the formation of constantly changing coalitions including a multitude of elements, e.g., *devices, nodes, sensors, protocols, services* etc. Coalitions could be autonomous in the sense that there cannot be any central authority or central control, they can emerge spontaneously according to some specificity of the coalition itself (physical vicinity, context, common interests etc.). Autonomous elements continuously join and leave the coalitions creating the conditions of an environment in which ensuring dependable and secure behaviors become quite complex. Nevertheless, the modelling and the analysis of dynamic coalitions is vital to master the development of *trustworthy* and predictable interactions among elements while preserving at the same time *privacy* (personal data, communication location etc) and freedom of mobility. Solution to this problem is not only a matter of computer scientists (from software architects to security experts) but it is a typical multidisciplinary area where competences in legal aspects, maths, telco, economics and control theory are paramount.

*For each problem, describe the impact of solving or failing to solve the problem and a description of specific challenges making the problem difficult.*

If a user does not trust its computing environment, no applications for that environment will ever be purchased by the user. For example, the first question answered by people wishing to join large projects that aim to share user's resources, e.g. Fon, is: "how can we trust people that will go to use my resources (e.g., wi-fi connections)?" This is why the volano of economy of applications based on a pervasive and intelligent environment will leverage on the degree of trustworthiness the environment will be able to achieve. Therefore, managing coalitions on which trustworthy and predictable interactions can be carried out means creating the conditions for a myriads of future applications that, in my view, is also difficult just to imagine today.

Descriptions of a few specific challenges

***Privacy-trust metrics and trade-offs***. Due to the inherent dynamism of the system in space and time that can lead the coalition to work in a disconnected way, there is no assurance to rely on any central authority while at the same time there is the need of cooperation among the nodes of the coalition to share resources, services and data. Cooperation is necessary because the coalition is indeed expected to be formed by resource constrained nodes while the latter are expected to execute distributed workflows (e.g., in battlefield, natural disaster scenario etc). This implies the definition of metrics for privacy and trust at the different levels of the protocol stack (physical, network, service) and the formulation of tradeoffs between privacy and trust. Trust, needed to enable cooperation, have to be gained also at the cost of disclosing some private data. The ultimate challenge could be therefore "trust preserving as many privacy as possible".

***Selfish and malicious nodes***. Solution of any problem on dynamic coalitions (including trust) needs to cope with the presence of selfish and malicious nodes. Let us note that selfishness in dynamic coalitions can be a natural act due to resource constraints and that the difficulty lies in detecting malicious nodes without false positive. Mechanisms guaranteeing element and information provenance should be exploited both to give some initial level of trust and to achieve this detection rapidly.

***Predictive Autonomic Systems***. Dynamic coalitions have to respond in an autonomous way to attacks and failures of parts of its components and have to recover to the required level of trust. This passes through techniques aimed at the isolation of malicious nodes and the bypassing of failed ones. Moreover, the designing of algorithms that constantly try to augment the level of trust in the coalitions while preserving privacy is also a priority. Further, improving trust has not to be done by sacrificing the scalability of the coalition.

***Composable and adaptable secure services***. Dynamic coalitions have to be able to establish on-the-fly end-to-end security assurance by composing and orchestrating services. This implies advances in securing all the protocols stacks and the components of a coalitions including OSs, middleware layer etc. Designers have therefore to architect systems, platforms and frameworks in which security and trust mechanisms are first class citizen (at any level of the protocol stack) in order to compose functional and non-functional properties.

*Provide potential research directions for handling the problem, and metrics for measuring progress against the problem. Possibly, describe also the limitations of current approaches, promising approaches not yet fully explored and desirable approaches where relatively few ideas have been proposed.*

Classical security mechanisms based on PKI or other central authorities do not seem to be flexible enough to provide security and trust in dynamic coalitions while preserving privacy at the same time. It is not certain that a central authority is available in space and time in a dynamic coalition (a coalition could indeed work in a disconnected way with respect to a central authority for a long time), so the usage of security based on trust becomes a necessary condition. The problem is that the study of trust is yet at an empirical stage. There are a few ideas that have emerged recently for developing a theory of trust in dynamic networks. For example, game theory adapted to the wireless environment for improving trust gained some attention (a few papers have appeared in 06). Also bioinspired theory is starting to attract researchers even though, to my knowledge, at this time no theory or valuable result is available.

Let us remark that privacy and trust in dynamic coalitions is a vertical problem in the sense that it affects all the sotfware/service/protocol chain till the final human user. Therefore, this problem can be instantiated to any layer and to any component. Of course each layer will have its definition of trust and privacy. The notion of privacy and trust for a microprocessor is indeed very different (and simpler) from the one of a human. Therefore a metric to measure progress with respect to the problem will be the progressive availability of trust frameworks preserving privacy starting from the bottom layers.

# Security & Privacy in Dynamic Wireless Networks – David Kotz

Institute for Security Technology Studies at Dartmouth College

**Position paper for Panel C: Security & Privacy in Dynamic Wireless Networks**: of evolving systems composed of ad hoc coalitions of large numbers of sensors and devices for new personalized services.

Digital technology plays an increasing role in everyday life, and this trend is only accelerating. Consider daily life a few years from now, in 2010: we will each be surrounded by far more digital devices, mediating far more activities in our work, home, and play; the boundary between cyberspace and physical space will fade as sensors and actuators allow computers to be aware of, and control, the physical environment; and the devices in our life become increasingly (and often invisibly) interconnected with each other and with the Internet. Today, typical home users struggle to maintain the security of their home computer, and have difficulty managing their privacy online. Tomorrow, these challenges may become unimaginably complex. As a research community it is critical for us to address the security and privacy challenges involved in developing this world of Digital Living in 2010. Without a focused effort on security and privacy, today, the technology developed and deployed tomorrow will be exceedingly frail and untrustworthy.

Consider, in particular, the advent of sensor networks. Although sensor networks have been an active area of academic research, and are becoming commercially available for deployment in industrial settings, sensor networks will soon have many uses in enterprise and residential settings. People will live in spaces, or work with devices, that have embedded sensing capability. For people to accept this new technology into their lives, they must be able to have confidence that the systems work as expected, and do not pose unreasonable threats to personal privacy. How can people be confident that the sensor network, and associated services and applications, will behave correctly– including, among other properties, protecting their privacy? This confidence builds on fundamental technical properties of the sensors, network, and services; on the organizations that operate and monitor the system; and on technical mechanisms to support monitoring and auditing. All of this effort must be informed by research that provides a deep understanding of the sociological underpinnings of privacy and trust in digital living, of the technological foundations for secure and robust sensor networks, and of technical and organizational mechanisms for users to express control over information about their activity.

There are many fascinating research questions to explore. How can the network securely sense, aggregate, and carry sensitive information (such as the location, activity, or health of occupants)? Can the network produce valid results, or at least indicate when its results are uncertain? Can the network detect or resist attacks that attempt to discover its secrets or disrupt its behavior, even when the adversary is an insider with access to a sensor device? How do we support strong security and privacy guarantees on low-capability embedded platforms? What mechanisms support adequate oversight by individuals and organizations? What trade-offs are occupants willing to make, between the conveniences provided by the sensor-based applications and the potential risk to their privacy?

What privacy policies and controls do users need to control the use of information about themselves, and effect these tradeoffs? What abstractions, interfaces, and mechanisms can people use to control the privacy of information sensed about them, and how can we make them accessible to the non-technical user?

These questions are difficult to resolve even in a controlled setting, where the sensor network is developed and deployed by a single administrative organization, such as a university campus or a corporate enterprise. The organization provides a policy framework for the collection and use of sensor information, social mechanisms for resolving abuse or conflict, and (presumably) thoughtful and thorough monitoring and management of the infrastructure and its applications. Although this centralized model reflects the computing environment of yesterday (where computer systems were managed by professional staff), it barely applies today (where users carry laptops, personal digital assistants, music players, telephones, and other "personal" devices) and will most certainly not apply tomorrow (when every user wears or carries numerous devices, some personally owned, all networked, and all interacting with each other and with the physical world). How can we develop protocols, infrastructures, and policy frameworks that work well in an environment that mixes networked sensing devices owned by multiple organizations (e.g., an large city building with embedded devices installed by the city, the building owner, multiple corporate tenants and private residences)? How do the interests of occupants and their organizations (employer, or building owner, for example) relate, and how does that affect the design or management of technology?

In short, there are countless challenges ahead of us. In my opinion, the research world has yet to focus on some of the most important challenges; too many in the sensor-network world are focused on narrow problems (such as key distribution) and on relatively benign application domains (agriculture, wildlife monitoring). We need to move now to look hard look at sensor networks that support pervasive computing– that sense *us* and our activity– because these will bring us the most critical security and privacy challenges tomorrow.

So what can we do to encourage and support the research community? First, we should encourage pervasive-computing researchers to engage with security researchers, to develop a common understanding of the problems and potential solutions. One approach is to encourage and support the creation of suitable international workshops. Second, we need to build pervasive-sensor testbeds, and support research that brings together groups of researchers that include expertise in pervasive computing (applications and systems), sensor networks, security, and sociology. Each testbed should explore a different range of potential application domains. For example, one testbed should leverage a decade of work in "smart homes," in which a model home is studded with sensors, actuators, and ambient interaction devices; now, the research focus is on security and privacy. Another testbed should explore the "smart office," enriching offices, conference rooms, and semi-public spaces. Finally, another testbed should explore the outdoor spaces of a town, where public and private meet in complex ways. All testbeds should include a mixture of interesting sensor devices, wireless networks, trusted-computing platforms, and be driven by compelling applications; all should be structured to support interdisciplinary investigation into the research challenges above.

# Towards Requirement Centric Security Evaluation and Testing – Reijo Savola

Reijo Savola, VTT Technical Research Centre of Finland

## 1 Introduction

Information security evaluation of telecommunication or software intensive systems typically relies heavily on the experience of the security professionals. Obviously, automated approaches are needed in this field. Unfortunately, there is no practical approach to carrying out security evaluation in a systematic way.

Security evaluation, testing and assessment techniques are needed to be able find adequate solutions. Seeking *evidence* of the actual information security level or performance of systems still remains an undeveloped field. To make progress in the field there is a need to focus on the development of better experimental techniques, better security metrics and models with practical predictive power. The goal of defining security requirements for a system is to map the results of risk and threat analysis to practical security requirement statements that manage (cancel, mitigate or maintain) the security risks of the system under investigation. The requirements guide the whole process of security evidence collection. For example, security metrics can be developed based on requirements: If we want to measure security behavior of an entity in the system, we can compare it with the explicit security requirements, which act as a "measuring rod" [1].

## 2 Challenges

In particular, there are following challenges:

- expression and modeling of security requirements at the appropriate level based on threat and vulnerability analysis and use cases,
- development of suitable security metrics for evaluation and testing,
- development of holistic security evidence collection methods, and
- in general: development of industrial-strength methods and tools for security evaluation and testing.

These challenges are obvious if we investigate an example process for security evaluation that could include the following iterative stages [1]:

- **Risk and threat analysis**. Carry out risk and threat analysis of the system and its use environment if not carried out before. These are lacking in many practical systems.
- **Define and prioritize security requirements** in a way that they can be compared with the security actions of the system. Based on the threat analysis, define the security requirements for the system, if not yet defined. The most critical and security requirements should be paid the most attention. Remember that the weakest links of the system are critical too.
- **Model the security behavior.** Based on the prioritized security requirements, identify the functionality of the system that forms the security actions and their dependencies in a priority order.
- **Gather evidence** from measured, reputation and tacit security information. Use suitable evidence collection tools like vulnerability identification and assessment tools.

- **Estimate the probabilities and impacts of security actions** based on the evidence. Aggregate the results to form a clear picture of whether or not the system fulfils the security requirements.

## 3 Impact

If there was a systematic way to build security requirements, a collection of security metrics, and industrial-strength methods for security evaluation and security testing, security engineering work would be affected to a great extent. These kind of goals would enable building of complete automated tool environments. Industrial-strength solutions in this area would be of great benefit to wide variety of companies producing software-intensive and telecommunication products.

## 4 Directions for Research

A practical security evaluation framework requires a lot of future development. In the following we list some goals for the future work. A suitable language needs to be developed to formalize and express security behavior and their cross-dependencies, as well as security requirements. Both the system security behavior and requirements need to be expressed in a way that it is possible to compare them. For example, a pattern language [3], with added details, could be used to describe the security actions. Attack trees [2] could be used for estimating especially the negative impacts of security actions. Development of a formal language to express security properties is a very challenging task and requires cross-disciplinary effort. In addition, semantics of the transformation from architectural aspects to the behavioral aspects requires future work.

A knowledge base of typical security constructs should be established to offer pattern information on their security behavior. The information needs to be collected experimentally to enable development of the knowledge base. Security evaluation or testing can be done in practice if this kind of support system could be used for security behavior modeling and suitable security requirement documentation is available. As a long-time goal, general-level statistical knowledge has to be collected on: security algorithms, network products, user behavior, applications, experiences from virus and worm attacks, etc. – about all critical issues contributing to the aggregate level of security.

## References


[1] Savola Reijo and Röning Juha (2006): Towards Security Evaluation based on Evidence Information Collection and Impact Analysis. In: Supplemental Proceedings of the 2006 International Conference on Dependable Systems and Networks (DSN), Workshop on Empirical Evaluation of Dependability and Security (WEEDS), June 25-28, 2006, Philadelphia, PA, pp. 113-118.

[2] Schneier B. (1999) Attack Trees. In: Doctor Dobb's Journal, December, pp. 21-29.

[3] Schumacher M. and Roedig U. (2001) Security Engineering with Patterns. In: Pattern Languages of Programs, September 11-15, Monticello, Illinois.

# Security Paradigm – from Assets Protection to People Protecting & Empowerment – Stephan Engberg

**Stephan J. Engberg**

**Define 1 or 2 difficult problems (research challenges)**

Security is mostly seen from the perspective of protecting assets from people. However, in this paradigm, we start by weakening person security and thereby create even more vulnerabilities both to people and through criminal actions such as identity theft also to assets. Any system can be attacked by stealing the identity of security-cleared users and only the user himself can know of an identity theft attack is ongoing.

The main problem is that we take offset in the assumption that Identification is pro-security but ignoring that Identification involves reuse of keys across different purposes, creating risk of fraud using data outside context, dis-empower the user from self-protection and focus on WHO the user (or device) is rather than WHAT the actual security properties are. Identification inherently involves the use of keys that can be used to attack the person being identified.

As such the existing security paradigm are caught in a self-destructive arms race to the bottom between attackers and those trying to protect ever more insecure systems. Bad examples are multi-model biometrics, surveillance, general logging, profiling, automated identification in as diverse technologies as passports, RFID, biometrics, communication protocols, payments etc.

These problems are already escalating out of control further fuelled by with ubiquitous and heavily integrated or even grid-type networking, Risk and security failure accumulate almost everywhere we look – because of the security paradigm itself

**For each problem, describe the impact of solving or failing to solve the problem and a description of specific challenges making the problem difficult.**

The complexity of security resolution and automatic reasoning across systems, protocols and domains are rapidly growing and as such also the mistakes, errors and barriers for growth. If the vision of an ubiquitous world where all devices are able to talk seamlessly with each other without collapsing into ever larger security breaches.

The main problem is the multiple facettes of the problem. Ranging from poor understanding of perception, trust and risk socio/economics over problems related to usability, resource constraints and control balancing to the fact that most existing protocols and technologies are unsecure. Three major value drivers are at stake:

a) Preventing instead of creating profitable cyber-crime such as Identity theft.

b) Trust towards the entire digitalization process which goes beyond the crime experienced

c) Enabling interoperability, new unpredicted functionality and limiting the number of interfaces a user has to deal with.

A vital aspect of the security and empowerment problem is realisation that biometric identification and surveillance is security-destructive and must be replaced with better balanced tools. Biometrics is integrated part of all identity schemes to establish basic root identity, accountability and for forensics. But for instance the existence of third-party certified biometrics outside absolute citizen control is an open invitation for criminal abuse and fraud.

We are going to have to deal with a range of serious threats from both abuse of biometric surveillance created as part of hastily deployed assumed pro-security infrastructure but even more on how to deal with the rapidly growing threat from biometrics surveillance as spyware, industrial espionage, wiretapping of political oppents and general surveillance to collect information that make people, systems and assets vulnerable for attack.

**Provide potential research directions for handling the problem, and metrics for measuring progress against the problem. Possibly, describe also the limitations of current approaches, promising approaches not yet fully explored and desirable approaches where relatively few ideas have been proposed.**

The security & dependability focus is in principle to eliminate the need for identification of people or devices. Instead we move towards a technology paradigm, where people are always in control and we are able to resolve and deal with security properties on a structured basis.

Basic questions raised are issues such as making security models and wireless protocols interoperable and dependable. Identification merely leads to the mess of access control cluttering with impossible rules as complication grows with the number of combinations. Permission management is turning into a nearly impossible issue leading to weak fallback solutions such as logging assuming abuse can be prosecuted instead of real security.

A system has no usage of the information that some entity is James Lucas, if it is not able to reason based on WHAT James Lucas is, how certain we can be on root identification, present authentication and means for accountability. If identity is made purpose-specific we can both protect assets from attacks from James Lucas and from attacks ON James Lucas Identity in order to secondarily attack systems and assets.

We need to depart with the assumption that we can protect data and especially identity credentials (such as credit card details) AFTER they have been collected. In a heavily digital integrated and all-ubiquitous world, primary security has to be preventive focussing on not creating the vulnerabilities in the first place and especially ensuring that security breaches do not scale due to reuse of keys across difference purposes.

This openly raise issues related to the police investigation interest of getting total access to information. However, we need to being with the fact that if we cannot design ICT systems for police investigation that ends up creating and worsening the crimes we are basically trying to prevent by making systems more vulnerable to both internal and external attacks.

The rule to apply is that criminals can do anything that government can do. So we better design for self-protection only limited by the insurance of accountability in case of abuse. The principle is using least invasive technology design to avoid weakning the wictim fearing the victim is an attacker and escalate the

security alertness as indicators show cause. E.g. Surveillance cameras cannot be protected from fraudulent abuse, but we can have physical shades for the cameas that move aside when a user fail to authenticate and thus can be assumed to be a threat. In ubiquitous computing this raises awareness towards how we can ensure that owners and especially temporary users can CONTROL any device despite the challenges related to usability, interoperability.

# Security Infrastructures for Mobile Devices - Wenke Lee

Wenke Lee

College of Computing

Georgia Institute of Technology

wenke@cc.gatech.edu

As wireless and mobile devices becoming ubiquitous and essential in our society, they will increasingly become the targets of malicious attacks. Although many existing security technologies can be used to thwart attacks on wireless/mobile devices/networks, there will be new attacks that require new security research.

An example is attacks that propagate via the Bluetooth interface of the mobile devices. Since these attacks do not go through the providers' core/backbone networks, traditional network security mechanisms, particularly those based on perimeter defense, are not effective. On the other hand, since the mobile devices will have limited resources (with respect to the demand for feature-rich OS/applications), host-based security mechanisms may not be practical either.

We proposed to study "security infrastructures" for mobile/wireless networks to counter attacks such as the above. The security infrastructures will enable research into: 1) prevention of attacks, such as "network admission control", where a mobile device is first checked and cleaned/patched when it connects to the network; 2) detection of attacks, such as capturing attack samples at the early stage and automatically analyze them and extract filtering rules or even generate patches; and 3) mitigation of attacks, such as automatically pushing filtering rules to network access points and patches to mobile devices.

In this research, we will build on our considerable experience in both wired and wireless network security, including vulnerability analysis, intrusion detection, and malware analysis. Specifically, we propose to build a network of smart Access Points and Honeypots as a concrete example of "security infrastructures" for wireless and mobile devices.

The wireless Access Points will be Linux (or other open source system) based systems that will allow us to implement security mechanisms such as firewalls and intrusion detection systems. The honeypots are machines that have multiple wireless interfaces such as Bluetooth, WiFi, WiMAX, EVDO, etc. Many laptops already come with several interfaces. In addition, these machines must support virtualization such that each machine can emulate multiple (kinds) of mobile devices. For example, three virtual machines on a laptop can emulate a WiFi laptop, an EVDO PDA, a WiMAX laptop. In other words, each machine hosts multiple (kinds) of honeypots. The honeypot machines are placed throughout the protection "space",

e.g., main buildings and gathering places on campus, to capture attacks. For example, a Bluetooth honeypot can be “infected” by a Bluetooth based attack. The machine that hosts the Bluetooth honeypot must have a robust VMM (virtual machine manager, or Hypervisor) that will not be compromised by the attacks to the virtual machine (i.e., the one that runs the Bluetooth honeypot). The VMM needs to have a security management module (e.g., at Domain 0, a special/privilege virtual machine) that can monitor events at the honeypots running on the virtual machines, detect attacks, and send the attack code and system/network log to a remote central server. The central servers run automatic analysis engines to generate filters and patches, and push them to the access points. When a wireless device connects to the wireless network, the access point automatically scans and patches the device so that it has the up to date system and security software. This also prevents an infected device from passing known attacks into the networks.

In summary, in the above infrastructures, the access points serve the purposes of intrusion prevention and mitigation, and the honeypots and back-end analysis servers serve the purpose of intrusion detection.

We are looking for partners in industry and government for this research, in particular for funding/donations of devices/equipment, and in the near future, for collaboration in testing/deploying of the infrastructure.

# Panel D

## Evaluating the Dependability & Security of Networked Systems – modelling, simulation, predictive evaluation, assurance cases

# Compositional Development and Assurance for Secure Systems – John Rushby

John Rushby

Computer Science Laboratory

SRI International

Menlo Park CA 94025 USA

## 1 Research Challenge: High Assurance Secure Systems

Despite substantial and sustained efforts since the 1970s, construction of distributed systems that require very strong security guarantees (e.g., EAL 6+) has not been reduced to engineering practice. Many efforts to construct and provide assurance for such systems fail or are abandoned, and most of those that reach deployment have either been very simple, or ruinously expensive. An effective engineering approach to the design, construction, and assurance of these systems is urgently needed.

## 2 Approach: The MILS Idea

I advocate a research direction known as MILS, which originally stood for "Multiple Independent Levels of Security," though my interpretation is more general.

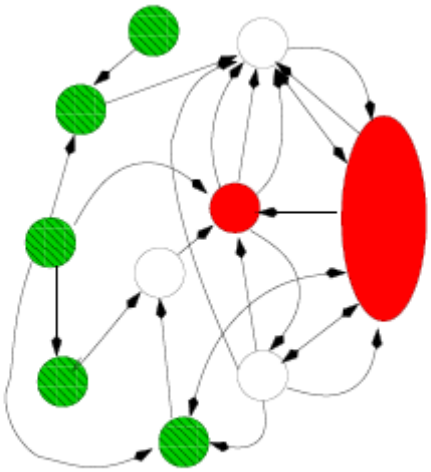

Figure 1: Example Security Architecture

Discussion of the design for a secure system or application inevitably leads to diagrams such as the one shown in Figure 1. These diagrams reflect a powerful intuition: security requires encapsulation of information and control (the circles) and explicit control over their transfer (the arrows). Implicit in the diagram is the expectation that there are no paths for information or control flow save those shown by the arrows[9] (i.e., the circles do not "leak" and there are no hidden or unknown arrows). Some of the circles encapsulate functions critical to the security of the overall system (in Figure 1, red is critical and green is semicritical); the implementations that correspond to those circles must be provided with strong evidence that they perform those functions correctly. It is a sound principle that the overall design should be developed in a way that minimizes the complexity and number of these "trusted" components.

One approach to simplifying the trusted components does so by increasing the total number of components: rather than a single trusted component that implements, say, a multilevel secure file system, we may prefer to have many untrusted single-level file systems and a trusted mediator that controls access to them. This is the essence of the MILS approach: assume that circles and arrows are free, and seek a system design that requires the simplest and fewest trusted functions (multiple instances of the same trusted function count only once).

When implementing the security architecture, several circles and arrows may share common hardware resources. The MILS approach thus entails two classes of trusted components: those corresponding to circles that perform security-relevant functions, which I call trusted operational components, and those responsible for realizing the security architecture (i.e., the circles and arrows) on shared hardware, which I call trusted foundational components.

The feasibility of the MILS approach rests on the observation that modern hardware and modern approaches to operating systems and other basic software (e.g., paravirtualization) indeed make it

[9] Research funded by AFRL through a contract to Raytheon

inexpensive to provide many circles and arrows, and can do so securely (i.e., without introducing leaky circles or unintended arrows).

## 3 Research Directions

The main issues are at the intersection of design and assurance. We need compositional methods for calculating the security attributes of a complete system from those of its security architecture and the attributes of its trusted operational components. This is difficult because many classical security attributes (e.g., non interference) are not preserved under composition—indeed, they are not even properties in the computer science sense (i.e., sets of traces). In addition, the trusted foundational components need to provide a composable foundation (i.e., one that preserves the attributes of the systems implemented on them), which is related to partitioning or separation, and these need to be additive (i.e., preserved under composition). Formal explications of these concepts need to take into account that certification is not the same as formal verification, so a related issue is development of a "science of certification" based on compositional principles.

## 4 Common Research Resources

These include formal verification tools, an evidential bus to link them together, a repository of worked examples and standard components; ontological and SOA tools for automated composition of operational components and for synthesis of protection profiles.

# Evaluating the Dependability and Security of Networked Systems – modeling, simulation, predictive evaluation, assurance cases – Bev Littlewood

Bev Littlewood, Centre for Software Reliability, City University, London

b.littlewood@csr.city.ac.uk

## 1 'Confidence' in probability-based dependability/security cases

Society is increasingly requiring quantitative assessment of risk (e.g. Basel II accords in banking). This poses some hard problems when the risks arise because software-based systems are not sufficiently dependable and secure; their solution requires a better understanding of *uncertainty* in system behaviour.

There are two types of uncertainty involved in making dependability claims about systems. The first concerns the actual behaviour of the system, in particular the process of its failures. The second concerns our assessment of this behaviour.

Informally, a dependability case comprises some *reasoning*, based on *assumptions* and *evidence*, that supports a dependability *claim* at a particular level of *confidence*. Often the claim is a 'given' – e.g., a protection system may be assigned a reliability requirement by the safety case of the wider system (e.g. that the probability of failure on demand is no worse than $10^{-3}$). 'Confidence' in the claim can be expressed as the probability that it is true, and will depend upon: the strength of the evidence, our confidence in the truth of the assumptions and the correctness of the reasoning.

Thus probability plays two roles:

- it is (often) used in the formulation of the claim: here it concerns uncertainty about the world;
- it is used to express confidence in the claim: here it concerns uncertainty about the argument used to support the claim.

For dependability cases involving *safety and reliability* claims, the first of these roles for probability has long been widely accepted. On the other hand, whilst notions of claim confidence are implicit in much practice (e.g. as expressed in standards), the probabilistic representation of confidence is more controversial (but see recent attempts at formalisation, e.g. (Littlewood and Wright 2006)).

For security, even the use of probability to express claims is still not widely accepted. There is not space here to preach at length about the necessity for probabilistic formulation of security claims (see (Littlewood and Strigini 2004) for a discussion on this). Suffice it to say that there is inherent uncertainty here (e.g. in the process of attacks upon a system), just as there is in reliability and safety. Whilst there are alternatives to probability as means of expressing this uncertainty, such as fuzzy/possibility theory, these do not have the power of probability, nor do they easily fit into a wider engineering framework, in particular for enabling quantitative risk assessment.

Probabilistic confidence-based security cases pose real challenges. But the difficulties are not, I believe, primarily in the mathematics and probability; rather they lie in our acquiring a better understanding of the variables upon which the modelling can be based. An important example is the notion of exposure, or 'effort' (e.g. by attackers), which plays the role that 'time' plays in reliability modelling. In general, and speaking as a modeller, I find the dearth of quantitative data the major obstacle to progress in this area.

A particular difficulty concerns issues of *validation*: how can we be assured that the numbers coming from our evaluations can be trusted?

I think there are useful lessons to be learned from earlier work on modelling reliability and safety. For example, there was a realisation that problems in *safety* (e.g. the '$10^{-9}$' problem) are *very* hard, and most progress has been on problems where the goals are more modest. Perhaps in security cases, also, we should learn to walk before we run: i.e. tackle first the problems of assessing and predicting modest security levels, rather than those involved in, e.g., national security.

## 2 Diversity and security

There has been a growing interest in the application of diversity in security in recent years. Many people think, for example, that the use of diverse intrusion detection systems may be useful, but others disagree. The issues here are, of course, empirical ones concerning efficacy: e.g. is diversity cost-effective in competition with other approaches?

So what is diversity? Informally, it involves doing things in *deliberately different* ways (and comparing the results). The simplest way just involves *separation*, e.g. simply isolating software development teams to prevent communication. More constructively, diversity can be *forced* by requiring the use of different development processes, tools, etc.

Whilst there are plausible grounds for believing that, in some sense, diversity is a good thing, some claims are not believable: e.g. the *impossibility* of common failures; e.g. *statistical independence* between failures. So the important question is 'how good is it?' Some insights into the efficacy of diversity to achieve *reliability* have come from experiments, and from probability modelling. An important question is whether such approaches, particularly the probability models, could be used in security, e.g to determine whether (and how much) diversity of protection against attack is effective. Other areas where diversity might be studied include:

*Diversity of intruders*. Clearly diversity will be useful in a team of intruders: how do you pick the best red team of size *m* from a population of size *n*?

*Diversity of intrusion procedures*. What is the best mix of attack procedures to use to find vulnerabilities? (cf. diverse fault-finding in software (Littlewood, Popov et al. 2000)).

*Diverse intruders, diverse sensors*. This introduces a further element of diversity – between the attackers and the sensors. Interestingly, this has not been addressed in reliability: there seems to be a view that 'nature' does not mount diverse threats – is this true?

## 3 References

Littlewood, B., P. Popov, et al. (2000). "Modelling the effects of combining diverse software fault removal techniques." IEEE Trans Software Engineering **26**(12): 1157-1167.

Littlewood, B. and L. Strigini (2004). Redundancy and diversity in security. ESORICS (European Symposium on Research in Computer Security), Sophia Antipolis, Springer.

Littlewood, B. and D. Wright (2006). "The use of multi-legged arguments to increase confidence in safety claims for software-based systems: a study based on a BBN of an idealised example," http://www.dirc.org.uk/research/DIRCResults/DiverseDependCases.html

# Thoughts on Evaluating the Dependability & Security of Networked Systems – John McHugh

**John McHugh**

Despite a number of attempts, we have had relatively little success in evaluating large scale networked systems for Security and dependability properties in artificial or test bed environments. There are a number of reasons for this, starting with a woefully inadequate understanding of the salient aspects of the real environment that must be captured and replicated in the artificial setting and continuing through the difficulty of provisioning a test bed to provide a test facility of the needed scale. Recently, I have been trying to approach the problem from two directions. the first is attempting to capture sufficient data so as to characterize network activity on scales that range from small enterprises to extremely large aggregations of of large enterprises (from a single /24 or less up to the equivalent of multiple /8s). The other is to develop methods for carrying out meaningful evaluations in situ. While these approaches are not without problems, results to date appear promising.

Computer Security researchers, especially those involved in the development of IDS or similar systems, frequently complain that there is no data suitable for evaluating their systems. They often rationalize their usage of the data created by Lincoln Laboratory for the 1998 and subsequent DARPA evaluations by claiming that primary data collection (or generation) is too difficult, despite the fact that this data has never been appropriate for IDS evaluation and becomes less so with the passage of time. The problems of using real data for evaluation are well known, and good, publicly available, data sets might be of substantial use to the community. The Lincoln approach was to study real networks and, then attempt to produce artificial data that was similar to real data in the dimensions that were critical to the intended use of the data. Many of the inadequacies in the Lincoln data were probably due to the fact that, at the time the data was created, the team had no expertise in either network data analysis in general or in intrusion detection.

For the last several years that I was at CERT, I was fortunate to have access to NetFlow data from the border of a very large network. Data collection started prior to 2003 and continues to this day. This data is useful in characterizing network traffic and has been used to inform synthetic data generation for at least one evaluation project. It is used operationally to provide situational awareness products (TOP-N reports, etc.) to the operators of the network and to analyze incidents such as worm and virus outbreaks.

As far as I know, no long term characterization of the data has been done. I believe that this would be extremely useful as it might allow us to develop some useful notions of “normal” behavior and its variation over time and network space. This, in turn, would help to determine training and retraining periods for anomaly detection based intrusion detection systems as well as the useful life of signatures for rule based systems. NetFlow is relatively easy to capture efficiently, and lends itself to compact storage, but its granularity lies somewhere between packet and session level data. It contains volume information, but no payload characterization; session start time and duration, but no pacing information. Aside from IP

protocol, source and destination addresses, and a few routing specific fields, port information for TCP and UDP, TCP flags, and message and code for ICMP, are the only application layer properties recorded. Within these limitations, this data can be used to characterize border flows, i.e. flows between enterprises and the internet core. With enough observation points, core flows can be estimated. If flows within the enterprise are also captured, internal traffic characteristics can be determined, as well.

Realistic traffic generation requires more information than can be obtained from net flow, but precisely 1 what is required depends on the way the data is to be used. for example, to evaluate an IDS that only uses packet header data, it is important to have a realistic mix of the relevant header parameters. Payloads should have realistic length distributions and proper checksums, etc. to create a realistic load on the network and protocol stack, but we need not be concerned with content otherwise. On the other hand, a system that analyzes application layer data in detail, using complex pattern matching requires that a great deal of attention be paid to the details of payload content so that the effort required to attempt matches on benign traffic is close to that required with “wild” data.

In addition, realistic traffic generation also requires the inclusion of a substantial noise component. Typical noise includes traffic that is directed to unoccupied addresses (systematic scans, misaddressed traffic, and responses to spoofed address traffic are among the sources), traffic that is not part of a proper conversation (DDoS, misconfigured or broken stacks, etc.) In terms of packet count, this traffic may account for over 90% of the volume on occasion and requires substantial processing resources from the system under evaluation. Systematic collection and analysis of network data from diverse sources and making it available to researchers, preferably under suitable non-disclosure agreements rather than in anonymized form, is a prerequisite for many research initiatives.

Preliminary work on applying the second approach to a rule based IDS, overly simplified, would proceed as follows:

1. For each IDS rule to be evaluated, configure and instrument an appropriate victim and an attacker that can launch the exploit to be detected.

2. Ascertain the ability of the IDS to detect the exploit under “ideal” circumstances. Forensic validation of the effect of the attack on the victim is also required. If the IDS is largely ineffective at this stage, further testing is unnecessary.

3. If possible, vary the attack to explore both the limits of the IDS rule and the effectiveness of the attack. The ideal result is an exact match. An overly broad rule allows false positives; one that is too narrow allows missed detections.

4. Stress test the IDS by increasing the load on the sensor. This can be done by using real traffic, possibly replayed at increasing rates or by crafting non-attack traffic designed to require extensive processing by the IDS(supra. In either case, the earlier attacks are embedded within the traffic stream. At this point, the IDS has been evaluated for detection ability and ability to function under load.

5. Deploy the IDS in the environment in which it will be used, along with an appropriate traffic capture device. The ideal device will capture the context for each IDS alert, a window before and after the traffic that caused the alert. Operate the system through several business cycles to obtain representative alerts. Several outcomes are possible:

(a) The number of alerts is sufficiently low so that they represent an acceptable false positive rate, even if they are all false.

(b) The number of alerts is small enough so that each one can be analyzed and classified as true or false, giving a rate that may or may not be acceptable. A larger sample may be required for significance.

(c) There are a large number of alerts. Careful sampling and analysis can estimate the false positive rate to any desired precision.

Note that deployment in a different environment is likely to require reevaluation of the false positive rate.

The overall evaluation result is a composite of the partial results. I believe that a similar approach may

be applicable to anomaly based systems, but the issues associated with realism for artificial data are even more critical here, especially for learning based systems where it is important that they not learn the right answers for the wrong reasons.

# Economic Solutions to Security Engineering – Aad van Moorsel

Contribution to Panel D: Evaluating the Dependability & Security of Networked Systems

**Aad van Moorsel**

Newcastle University, UK

Engineering security solutions in complex systems inherently involves trading off security against other system properties. Without good techniques, methodology and best practices, security engineering trivializes to binary argumentation about 'secure' or 'non-secure' systems. The promising technology we are after is that of 'economising security', which is based on a dual realisation. First, advanced engineering always boils down to valuing (monetary or otherwise) system properties. Hence, to include security in system engineering, valuation of security is a must. Vice versa, techniques from economy may be useful to establish trust. To illustrate the latter, truth telling about quality properties in service-oriented systems can be enforced by pricing and penalising suitably.

*Define 1 or 2 difficult problems (research challenges):*

- objectively engineer security, valuing security properties without trivialising it into an all or nothing proposition
- value security such that one can objectively trade off security with other system properties, such as performance

*For each problem, describe the impact of solving or failing to solve the problem and a description of specific challenges making the problem difficult:*

- without the ability to value security (based on either quantifiable metrics, heuristics or best practices), the debate about security trivialises to a binary argument of secure versus non-secure.
- without the ability to value security complex system cannot be expected to be engineered to be secure.

*Provide potential research directions for handling the problem, and metrics for measuring progress against the problem. Possibly, describe also the limitations of current approaches, promising approaches not yet fully explored and desirable approaches where relatively few ideas have been proposed:*

- it is known that the only way to deal with multi-objective optimisation is to map different metrics on a single scale, so as to create a (partial) ordering of design alternatives. Good methodology is required to facilitate such mapping.
- If a single metric is needed (per the above), one might as well choose business value (expressed in monetary value) as single metric.

- not yet fully explored is the fact that pricing and penalising can lead to truth telling, and hence to more justifiably trusted systems.
- quantification of security is desirable, but often not realistic. ‘Effort’ based models as well as threat based models suffer from the inability to parameterise the models, since they require information that belongs to the ‘not known and never to be known’ category.

# Using Risk As A Security Metric - Sami Saydjari

O. Sami Saydjari
Cyber Defense Agency, LLC
Wisconsin Rapids, WI 54494
**ssaydjari@CyberDefenseAgency.com**

Why measure security? To make good decisions about how to design security countermeasures, to choose between alternative security architectures, and to improve security during design and operations. So, in essence, measurement is a decision aid. What needs to be measured depends on the decision.

**Definitions**. Security is the freedom from the possibility of suffering damage or loss from malicious attack. A metric is a system of related measures enabling quantification of some characteristic. A measure is a dimension compared against a standard. *So a security metric is a system of related dimensions (compared against a standard) enabling quantification of the degree of freedom from possibility of suffering damage or loss from malicious attack.*

In the end, the decisions needed are to be made by the owners of systems. They must decide how much security they need just as they must decide how much they need of other quantities such as functionality, performance, and reliability. To make such decisions, often cost and benefits are weighed as two of the most important dimensions. To assess benefit with respect to security, one needs to look at how much expected loss can be avoided. Expected loss can be found by looking at the probabilities of the range of possible attacks succeeding multiplied by the consequences of the attacks. For a simple example, if there is a 10% chance of an attack succeeding over the next year and the consequences are $100 in lost profit, then the expect loss is $10 per year. In other words, the owner has a risk exposure of $10 per year. If the owner can invest $5 to drive the likelihood of that attack to near-zero, then it might be seen as a good investment.

Security metrics are needed to understand, compare, and guide security toward improvement both in operations and design. In this panel, we examine risk as a potential direct metric of security. One good property of risk as a security metric is that it is a dimension that directly addresses the definition of security by addressing both possibility and damages. Another good property is that it causes defenders to grapple with how adversaries really attack systems. Also, the metric's output is a statement of residual risk of a given system and can be used directly by a system owner to decide on acceptability of that risk.

One negative property is that the metric is not prescriptive on how to accomplish it. In fact, to the extent that methods have been developed, they are quite difficult to execute.

**Key Issues**. There are several issues in making risk a practical metric.

(1) One must have a complete understanding of all possible attacks. This is particularly difficult because attackers can be incredibly creative and innovative. Can attack trees help with encouraging breadth of coverage and reuse of attack knowledge?

(2) One must assess the probability of attackers attempting attacks and their probability of success. Are accurate models of adversary behavior available as well as good empirical data and a firm understanding of all the complex factors that drive the probabilities including hard-to-predict world events and the state of the economy?

(3) To assess the consequences for the attack, one needs a good model of the owner's mission, the owner's system, and how the system supports the mission. Such models are not yet well-developed. Also, non-monetary consequences to mission are hard to put units on. What can we do?

(4) It is very difficult to validate the accuracy of a risk metric because ground truth is difficult to establish. What are possible approaches to achieve validation?

**What are alternative metrics?** *Criteria-compliance* was one of the earliest security metrics. Such criteria prescribe design processes and independent checking of system properties by evaluators. Such metrics have the advantage of having a prescription included for all to inspect. Some criteria offer rank orderings on system quality (A1 versus B2, for example) so consumers can compare directly. On the downside, the metric is indirect with respect to the definition of security. One must work hard to convince oneself that following the design prescription results in a system that is truly well defended against the full range of attacks possible. Mappings between the processes and effectiveness against attack classes are almost never offered.

*Intrusion-Detection-based* metrics relate to the performance of intrusion detection systems of one sort or another. For example, one can measure the number of viruses detected per day, or the number that are missed by the detection system and go on to do damage. One can look at the number of external connections blocked by the firewall or the number of attacks blocked. The good thing about these metrics is that they are easily measured and readily available from the tools that are used by system owners. The bad thing is that it is not clear how they relate to the definition of security. If one system has more detections than another, does that mean it is better?

*Policy-Based* metrics similarly look at quantities like number of unauthorized login attempts, files accesses, and so on. These metrics have the same advantages and disadvantages as the intrusion-detection-based metrics. Unfortunately, such metrics may end up measuring the inadequacy of user training more than it measures actual system security.

*Incident-Based* metrics begin to get closer to directly measuring actual security. Such metrics look at the actual successful attacks that occur, the frequency and the real damages. This is how insurance companies often work in calculating how to insure against various events such as hurricanes, theft and fire. This approach is promising and, with time, can become a reliable and useful metric. Currently, there is insufficient data on attack incidence and damage assessments. Further, this empirical data is limited by our ability to see successful attacks (given relatively primitive detection capabilities) and by rare events with big consequences (such as a massive strategic attack against an entire nation) not being included in the observed data.

**Qualities of a Metric**. From the above analysis, we see that some qualities of a good metric include: (1) to measure the right thing (supports the decisions that need to be made, for example), (2) to be quantitatively measurable (damages in dollars of profit loss, for example), (3) the capability to be measured accurately and (4) to be validated against ground truth.

To these criteria, one could add the following desirable properties: (5) inexpensive (both in time and cost) to execute, (6) able to be refereed independently, and (7) repeatable so that the results are independent of the analyst performing the measuring, and (8) scalable from small single-computer systems to large nation-scale enterprise networks.

**Potential Research Directions**. A research agenda can be derived from the key issues discussed above. One must develop a framework for systematically identifying all of the relevant attack events that could happen to a system and of estimating the probabilities and the consequences of those events. Such estimations require the creation of good models and a means to validate those models. *In essence we require a valid theory of system attack.*

**Common Research Resources**. Once a theory of system attack is developed to create a combined estimate of risk, one needs to develop a system of collecting actual damages to real systems to compare against the forecast estimated damage based on the estimate of risk. Such surveys of actual systems could create an invaluable community dataset against which one could evaluate alternative candidate estimation models. Also, system test beds which have modeled missions (perhaps specialized to various domains such as SCADA, banking, and telecom) should be created so that experiments may be run using red and blue teams in simulated attack-defense situations to test models of how attack events manifest in terms of consequence to mission as measured in dollar damages, for example. Such test beds will require the creation of models of how missions map to system resources which support that mission.

# Evaluating the Dependability & Security of Networked Systems – modelling, simulation, predictive evaluation, assurance cases - Robin Bloomfield

Robin E Bloomfield, CSR City University and Adelard, reb@adelard.com

*Research challenges and impact*

The challenge is to assess the dependability (risks, resilience) of complex adaptive socio-technical systems of a wide range of criticalities, to know the limitations of these assessments and to able to communicate to a broad range of stakeholders.

As a subset of the challenge (and a measure of success)

- to be able to undertake risk assessments that lead to insurable systems and services and to provide risks that can be traded
- to stand still - so that there are no disasters. The demands of current systems, complexity of new systems, time to market, pressures, the structure of the supply chain, competence, the scale of reuse, evolving threats all make this non-trivial and uncertain.
- to allow a choice of doctrine between conservative approaches to technology adoption and an ability to make informed decisions to exploit the benefits of using computer based systems.

The ability to have confidence in systems and services is a fundamental requirement for the critical information infrastructures, the future Internet, ubiquitous and pervasive computing. Lack of appropriate trust can lead to unexpected losses (human, environmental, political, commercial). Lost trust is very hard to win back.

*State of practice, potential research directions, metrics, limitations*

The capability to develop assured systems (at what industry judges reasonable costs) is not a solved problem even for today's systems. Add the key words: adaptation, ubiquity, complexity and the problems are compounded and changed. The assessment that a system is fit for purpose, and continues to be fit for purpose as the environment, use and implementation change, is a complex socio-technical process. Historically there has been a continuing move away from a standards, compliance based approach, to one where goals of system behaviour are justified. So-called safety, assurance or dependability cases are becoming widespread.

Current best practice involves graphical summaries of the case, tracking of evidence status, propagation of changes through the case, automatic linking to requirements and management tools. However claim decomposition is normally very informal and argumentation is often not explicit. In fact the emphasis is on communication and knowledge management. The uptake of the approach is very strong in some sectors and there is a little empirical evaluation of the benefits of the "cases".

**One research direction is to develop the use of "assurance case" as an overarching framework. The "case" should be seen from a number of viewpoints**

1. As a boundary objective between the different stakeholders who have to agree (or not) the claims being made about the system.
2. As an argumentation framework that allows us to reason as formally as necessary about all the claims being made.
3. As part of a "signal processing system" – the licensing, compliance or certification process – that seeks to reject false claims and accept good ones.

Some more specific directions are:

*Claim semantics*. In terms of claims, the goals of systems, we are rich in metaphors (often anthropomorphic - trustworthy, resilience, robustness, intelligence, adaptable) but short of precise operational meanings that we can assess against. On approach is to develop an operational semantics (via models) of some instantiations of these metaphors (e.g. for specific critical information infrastructures) so we can understand the properties we are seeking to assess. This is an interdisciplinary challenge.

*Claim decomposition* Develop a rigorous approach to claim decomposition (and confidence propagation) where claims are structured on the basis of architecture, functionality and attributes. Separate probabilistic methods and deterministic approaches are available but the synthesis is challenging Need approaches that can scale.

*System of systems resilience assessment*. Many individual systems have mature methods for modelling and analysis. The interconnection and evolution of system and interdependencies within critical infrastructures are already the subject of separate EU and US research. There is a need to develop theories and methods that deepen our understanding and design abilities and understand their contribution in comparison with detailed specific simulation. Some promising approaches are developments of hazard analysis approaches and insights from complexity science, in particular to deal with uncertainty in structure. (Unexpected) Interdependencies are normally seen as detrimental but complex systems often appear more reliable than a naïve analysis would predict due to emergent redundancies and diversity. Need to model and understand this.

*Assumption doubt*. In well-engineered, critical, systems the limitations that we can convincingly claim about dependability are likely to come from assumption doubt. We need methods for dealing with this uncertainty (the unknown unknowns?) and new approaches to common mode failure assessment.

*Incorporate the socio-technical into risk assessment.* Risk communication, issues of trusted vs trustworthy, and the adaptation of people and organisation to evolving systems are important directions. Specifically explore benefits from responsibility modelling to provide a means to understand, identify and mitigate potential failures. Lastly need to explore and understand the interaction of risk assessment and markets: how uptake of technology is impacted; how insurance might work.

*Describe common research resources that the research community could create*

Develop common metaphors and mathematical models. Benchmark problems to drive the technical and scientific basis of cases. Small enough to be tractable (on a whiteboard?) but rich enough to expose the issues.

Jointly develop assurance cases. Appraise international differences in risk communication, risk appetitive and argumentation (e.g. probabilistic vs deterministic).

Develop assurance of components of critical infrastructure that are for the common good (e.g. DHS work). Develop specific tools and theories for subset properties within a common framework (e.g. timing, reliability, correctness). Assess resiliencies and vulnerabilities through hidden interdependencies. Possible experiment on simulation of large-scale systems. In some exercises synthetic but realistic countries are used. Consider benefits of such a research model – incorporate utilities, geographical information etc.

# Panel E

## Evaluating the Dependability & Security of Networked Systems – monitoring, operational assessment, auditing

# Research Issues in Quantitative Assessment of Trust- David M. Nicol

Dept. of Electrical and Computer Engineering, and Information Trust Institute

University of Illinois, Urbana-Champaign

The topic of metrics repeatedly appears on lists of pressing open issues in security research; this is a reflection of the difficulty of finding appropriate quantitative assessment of system trustworthiness. A promising means of making headway is to narrow the problem somewhat at first, by focusing research in this area on systems in some domain with sufficient structure to guide the development. The challenges I see include

- To find application domains (e.g., credentialing systems) that lend themselves to development of quantitative trust assessment.
- Within an application domain, identify key quantifiable measures that capture different dimensions of "trustworthiness" in data, and process.
- Develop models appropriate for estimating these measures,
- Develop off-line and on-line means of assessing trust through these measures.

The impact of success is significant. System developers and end-users want trustworthiness, and both want it inexpensively. Success at development of methodologies and tools for assessment at design time will reduce the cost of provisioning trustworthiness. The impact of failure is significant. Increasing larger and complex systems will be built without a firm grounding in assessment engineering---and confidence in trustworthiness will be based on ad-hoc assessments that leave room for gaps.

I see at least two significant challenges that makes quantitative trust assessment hard. One is time-scale: the designers of a well-engineered system will have given some consideration to trust, and the system is likely to work as intended most of the time. In such cases trust failures are rare, which creates the challenge of driving the assessment towards unusual corners of the operating envelope. Time-scale is problematic also, in that behaviors that initiate trust failure may occur at a much faster time-scale than the system-scale behavior needing trustworthiness, or may be otherwise decoupled from system-scale behavior by the need for a large volume of fast-scale events. An individual packet is unlikely to cause a computer service to fail, but a large number of them in close succession may. A related challenge is in either filtering or aggregating measures so as to reduce their volume (for analysis or monitoring) while retaining the essential features needed to capture the impact on trustworthiness.

I think that a promising avenue for collaborative research in this area is to partner with groups who are interested in development of trustworthy systems in particular domains, and are developing technologies to provide that trustworthiness. Such a team could develop design methodologies and tools where model-based trust assessment plays an important role, and test-beds where metrics identified in the design assessment process can be measured and monitored in experiments designed to both demonstrate trustworthiness, and to push the system into behavioral regimes where trustworthiness may fail.

# The case of networked industrial systems - Marcelo Masera

Marcelo Masera, JRC–EC

One of the major transformations of the last years has been the networking of industry: within their installations/plants, between the technical and the business sectors of companies, and with service providers (e.g. for maintenance purposes). This has taken place in parallel with a shift from proprietary technologies and standards, to COTS and typical commercially available ones (e.g. TCP/IP, Windows).

This transformation has been driven by the vast functional and operational advantages deriving from the more immediate access and processing of data related to the industrial processes. Unfortunately, it is rather evident the gap regarding the evaluation of the security implications of such changes. We know that being connected implies exposure to security threats, but technologies were deployed in industry without accompanying security assessments.

Why has this occurred? First of all, there was some lack of awareness of the implications, as the networking of industry was bringing forward for the first time the subject of cybersecurity.

Security was an issue in mission–critical systems (defence and aerospace), but not in the majority of industrial systems. With the increasing realisation that many systems can be the object of security incidents, and that they might cause not only local disturbances but major societal disturbances (so–called critical infrastructures), the need for proper evaluation means is evident.

In addition, several factors affect this lack of proper security evaluation, and some of them present important research challenges:

–Evaluation with analytic and simulation means of industrial process control and communication systems, with the characterization of the specific vulnerabilities and threats. This goes beyond the simple evaluation of components.

–A comprehensive security risk assessment method that can support the evaluation of a networked industrial environment, including their real time control systems, random and deliberate failures, and not the least system–of–systems.

–The issue of dynamically updating security evaluations with the acquisition of new information regarding vulnerabilities, threats, attack methods or the efficiency of countermeasures.

–Metrics of security and evaluation criteria: what is the relevance of Common Criteria, how to define assurance baselines, etc. This topic is also affected by the dynamic nature of security.

# Next Generation Attacks on the Internet - Evangelos Markatos

Evangelos Markatos
Foundation for Research
and Technology Hellas (FORTH-ICS)
Greece
markatos AT ics.forth.gr
Angelos Keromytis
Columbia University
USA
angelos@cs.columbia.edu
Distributed Computing Systems LAB - ICS-FORTH, Greece
Columbia University, NY USA

**Defending Against Next Generation Attacks Through Network/Endpoint Collaboration and Interaction**

*(Define 1 or 2 difficult problems (research challenges)*

Over the past few years we have seen the use of *Internet worms, i.e.,* malicious self-replicating programs, as a mechanism to rapidly invade and compromise large numbers of remote computers. Although the first worms released on the Internet were large-scale easy to- spot massive security incidents [MSB02, MPS+03, SM04, BCJ+05b], also known as *flash worms* [SMPW04], it is currently envisioned that future worms will be increasingly difficult to detect, and will be known as *stealth worms* [SPW02]. This is partly because the motives of the first worm developers were centered around the self gratification brought by the achievement of compromising large numbers of remote computers, the motives of recent worm and malware developers are centered around financial and political gains. Therefore, although recent attackers still want to be able to control a large number of compromised computers, they prefer to compromise these computers as quietly as possible, over a longer period of time, so as not to be detected by any security defenses. Thus, to achieve a stealthy behavior, these attackers have started using, or at least have the capacity to use a wide variety of mechanisms that will make their worms more difficult to detect. Such mechanisms might include:

• Encryption. Attackers may communicate with the potential victim using a secure (encrypted) connection, making it difficult for network-based Intrusion Detection Systems [Roe99, XCA+06] to spot their attempted attack.

• Metamorphism. The body of worms usually contains some initial code that will be executed when the worm invades the victim computer. Metamorphism obfuscates this code by adding various instructions to it, and/or by substituting blocks of instructions with equivalent blocks of other instructions [SF01]. In this way two "copies" of the worm would be completely different from each other confusing worm detection systems which depend their effectiveness on the fact the all copies of a worm are practically identical [SEVS04, KK04, AAM05].

• Polymorphism. Polymorphic approaches obfuscate the worm's body by encoding it and prepending a decoder. When propagating, the worm mutates its body so that two "copies" of the worm would look completely different from each other (modulo the body of the encoder) [Sz¨o05, DUMU03, K201]. Much like metamorphic approaches, polymorphic systems confuse worm detection systems which depend their effectiveness on the fact the all copies of a worm are practically identical [SEVS04, KK04, AAM05].

• Hit Lists. The first versions of recent worms selected their victims completely randomly, *i.e.*, by generating a random IP address in the range 0.0.0.0 to 255.255.255.255. It has been proposed however, that worms may be more effective if they first create a *hitlist* of all vulnerable computers and then attack only computers in the hitlist [SPW02, AAMA05]. This hithist may even be filtered to exclude honeypots [10]1. Armed with a hitlist, a worm is able to compromise a number of vulnerable computers generating the minimum amount of traffic possible, evading detection mechanisms based on visible traffic anomalies.

• Hybrid Worms. Traditional worms used to invade computers by exploiting vulnerabilities of applications listening for Internet connections. However, as more and more computers are hidden behind firewalls and do not listen for incoming Internet connections, they are theoretically protected from such types of attacks. However, to compromise computers protected behind firewalls, worm developers may exploit several different invasion paths including, infected email attachments, infected files shared through peer-to-peer (P2P) networks, and infected files accessed through locally shared disks [KE03]. In this way, an attack may enter an organization as an email attachment, may spread to individual departments through infected disk shares, and may jump from department to department through traditional remote procedure calls.

• Defense Mapping. Many of the proposed (and deployed) techniques for detecting and countering new attacks use honeypots as the early warning system [Spi03, DQG+04, YBP04, CBMM04, BCJ+05a, RMT05, MVS01]. However, recent work has shown that attackers can exploit certain features and aspects of a honeypot's behavior to identify and avoid such detectors [BFV05, SII05, RMT06]. Combined with hitlists, this can render worms (especially slow-spreading ones) and other automated attacks virtually undetectable.

• Client-side Attacks. In the past year (2005–2006), we have seen an increase in the use of zero-day attacks aimed at client software (especially browsers, but also various types of document viewers such as

[10] A honeypot is a computer waiting to be attacked. Once attacked, the honeypot records as much information as possible so that the security administrators will be able to characterize the attack and generate a signature for it.

Microsoft Word, Excel and PowerPoint, and Adobe Acrobat). Other than stand-alone, host-based intrusion detection/prevention mechanisms (such as virus scanners), very little has been done in hardening vulnerable client systems.

**Impact of failing to solve the problem**

*(For each problem, describe the impact of solving or failing to solve the problem and a description of specific challenges making the problem difficult)*

Compromised computers can be used to cause harm to third parties or even to cause harm to their traditional owners.

• Attacks to third parties. Recent worm writers organize compromised computers into botnets, *i.e.*, armies of hosts which are primarily used for malicious acts including

launching of DOS attacks, blackmailing, sending of SPAM mail, click fraud, theft of intellectual property, and even identity theft. One would envision that botnets in the future could be used for political purposes as well.

• Attacks to the owners of compromised computers. A compromised computer can be used to steal the private data and the identity of the owner of the computer. Once, however, ordinary users start to realize the dangers of a compromised computer, they will probably get increasingly less inclined to trust their computers for financial transactions as well as private communications. This will probably *impede the adoption of the information society* and may eventually reduce its overall spread and impact.

**Research Directions**

*(Provide potential research directions for handling the problem, and metrics for measuring progress against the problem. Possibly, describe also the limitations of current approaches, promising approaches not yet fully explored and desirable approaches where relatively few ideas have been proposed.)*

Over the last five years there has been significant research in the area of detection and containment of cyberattacks. Indeed, we believe that we have currently reached the point where it is possible to detect rapidly spreading and massively parallel flash worms. However, it is unclear whether we have the complete technical knowledge or the deployed mechanisms in order to detect and contain *stealth attacks*. Using a combination of the techniques described earlier, such attacks can become invisible (or at least very difficult to detect) to network-based defenses.

Our view is that such attacks can only be detected via large-scale collaboration among end-hosts: by exchanging and correlating relevant information, it is possible to identify slow and stealthy attacks, and to take appropriate measures to defend against them, or at least quarantine those nodes that appear to have been compromised. Specifically, we believe that it is increasingly important to include home and small business computers in the attack detection process. These computers are increasingly becoming the primary targets of most attackers. Therefore, including them in the worm detection process will increase the chances of worm detection. Exemplifying a large range of access patterns and a large range of

applications, these computers typically tend to have more representative configurations than the traditional honeypots currently being used in worm detection. Furthermore, ordinary computers being used by their regular owners are more difficult to be categorized as honeypots and to be avoided by future attacks. The inclusion, however, of home computers in the detection process, should (1) guarantee the safety of the end computer and (2) the minimum possible intrusion in the ordinary use of the computer.

On the other hand, we are not completely discounting network-based defenses: rather, we believe that such defenses must be integrated with end-host defenses. In the past, network and end-host security were viewed as two distinct areas that were meant to complement each other but kept separate. While this allowed for a clean separation between the respective security mechanisms, it also meant that the potential of both was stunted. Furthermore, by keeping them isolated, it was (and is) impossible to exploit scale for defensive purposes. Exploiting scale is something that attackers have learned to do well, as evidenced by such phenomena as distributed denial of service attacks, self-propagating worms and botnets.

The industry is beginning to follow such an approach, albeit in a fragmented, ad hoc fashion. For example, several enterprises exchange alert and IDS logs through sites such as DShield.org; anti-virus vendors with extensive presence on the desktop are correlating information about application behavior from thousands of hosts; network security and monitoring companies perform similar correlation using network traces and distributed blackholes (honeypots). To the extend that such approaches are being explored, they seem largely confined to the realm of information gathering. This also largely seems to be the situation with the DoD and the various agencies. For example, DARPA is currently funding the Application Communities effort, which seeks to leverage large software monocultures to distribute the task of attack monitoring — again, an approach confined to the end-host. Previous work (notably in the DARPA OASIS program) looked into the space of reactive security, but only considered small-scale environments. Arguably, we need to extend the reach of our collaboration-based mechanisms to counter such pervasive threats as DDoS and botnets.

### *APPLICATION-LEVEL REFLECTION ATTACKS*

Thus, we argue that it is important to transition into an network architecture design where networks and end-hosts, in various combinations, can elect to collaborate and coordinate their actions and reactions to better protect themselves (and, by implication, the network at large). There are several research issues arising in such an environment, including:

• what problems are best addressed through a collaborative approach

• new mechanisms at all levels of the network architecture (routers, protocols, end-hosts,

processes, hardware) that are “collaboration friendly”

• metrics that quantify the security of collaborative approaches over non-collaborative

approaches

• who to trust, and to what extend

• how to prevent attacks that exploit such mechanisms, including insider threats

• command-and-control *vs.* loose-coupling mechanism composition

Furthermore, in an era of distributed software services (what is fashionably called ”Web 2.0”), no single application, node, or network has enough information to detect and counterhigh-level semantic attacks, or

even some of the more conventional web-based malware (*e.g.*, cross-site scripting attacks). Large-scale distributed systems require large-scale distributed defenses. This is particularly so within specific application domains (such as healthcare and industrial SCADA control), where large-scale collaborative (but independent) defenses will allow better control to critical information and resources.

**Application-Level Reflection Attacks**

*(Define 1 or 2 difficult problems (research challenges)*

Over the last few years we have seen an increase in the use of botnets, that is, compromised computers being used by attackers for a wide variety of malicious purposes including sending spam email, blackmailing targeted victims, and preventing victim computers from functioning. Although we will continue to see the use of botnets in the near future, we also expect to see the increasing exploitation of *non-compromised computers* for malicious acts. In this scenario, attackers carefully trick non-compromised and possibly non-cooperative computers into acting on behalf of the attackers. For example, in PuppetNets [LAAA06] it was shown that attackers which control a web server may direct a large number of ordinary web clients towards repeatedly requesting web pages from a victim computer. Made possible by the intelligent use of Javascript, these repeated requests where completely transparent to the user in front of the browser, who, all things being equal, did not see any malicious attack being going on. Similarly, Athanasopoulos *et al.* showed that peers participating in the Gnutella file sharing peer-to-peer network can be easily tricked into believing that a victim computer serves a large number of popular files, which in turn, directed a large numbers of requests towards the victim computer, possibly depleting it of its resources [AAM06]. Overnet, another real world P2P system for file sharing, may also be abused in the same fashion [NR06]. Randal Vaughn and Gadi Evron, in a preliminary work, published some techniques to use the DNS [MD88] system as an amplification platform for Denial of Service attacks to third parties by sending malcrafted DNS requests [VE06]. We believe that in the near future we will see an increasing number of such vulnerabilities which will make possible the use of a large number of non-compromised and non-malicious computers into malicious activities.

We call the collective exploitation of these non-malicious clients *Application-Level Reflection Attacks*. As computer applications become increasingly diverse, such attacks may spill beyond the traditional world of computers towards other networks, including, for example, the telephone network. For example, by exploiting a number of Skype (or other Internet telephony) clients, attackers may jam the telephone numbers and/or faxes of victim organizations with bogus telephone calls. To make matters worse, jamming may also be directed to organizations providing vital information services as well, precluding these organizations from providing their service even at critical times. And, since this attack originates outside the traditional telephone network, it might be difficult to trace the attackers back using traditional telecom-style traceback mechanisms. To make matters worse, even when the attack is traced back, the last point in the trace may be a set of non-compromised computers tricked into making these bogus telephone calls. Thus, attackers who have the power to selectively clog a decent number of telephone lines, may use this power for blackmail, revenge, or even to terrorize a selected subset of the population.

**Impact of failing to solve the problem**

*(For each problem, describe the impact of solving or failing to solve the problem and a description of specific challenges making the problem difficult)*

As computers are interwoven within several other services in our society, computer security problems diffuse into the general fabric of the society. Thus, a problem which originally initiated in the cyberspace,

may easily transcend into other spaces as well. For example, although SPAM used to be a problem of people communicating via email, it is becoming increasingly easier for SPAMers to make SPAM telephone calls, to send SPAM voice messages, to send SPAM SMSes, to send SPAM MMSes, and in general to SPAM traditional lowbandwidth communication channels using the efficiency, precision, and speed of modern digital computers.

Such attacks can be used for several reasons including:

• Disrupting the operation of organizations who depend on reliable communication channels (*e.g.*, telephones, mail, Internet). Such organizations may range from airline reservation services to the friendly neighboring take out Chinese restaurant.

• Cyber-vandalism against neighborhoods or even towns by selectively clogging their communication channels in a time of need.

• Intimidating, or even terrorizing, large numbers of people by clogging their telephone access to local services such as hospitals, schools, etc.

**Research Directions**

*(Provide potential research directions for handling the problem, and metrics for measuring progress against the problem. Possibly, describe also the limitations of current approaches, promising approaches not yet fully explored and desirable approaches where relatively few ideas have been proposed.)*

The problem of Application-Level Reflection Attacks is relatively new and there exists very little research in this direction. However, as the problem is getting better understood, we envision research along the following lines:

• Document the extent of the problem. At the time of this writing it is not clear which systems can be inadvertently used for Application-Level Reflection Attacks.

• Document the impact of the problem. We need to develop scenarios which clearly show what is the impact of the problem and what is the financial, social, and political cost associated with it.

• Develop Detection Mechanisms. Since Application-Level Reflection Attacks consist of non-compromised computers which behave in a seemingly normal way, we need to develop sophisticated mechanisms for detecting them.

• Develop Defense Mechanisms. Once a malicious behavior is detected, we need to have in place defense mechanisms which will block, or at least restrain this type of attack. We envision defense mechanism both close to the host, as well as close to the victim(s).

In closing, we believe that Application-Level Reflection Attacks show that it is technically possible to attackers to perform major attacks against traditional services by manipulating the behavior of a large number of non-compromised computers. Since these attacks involve noncompromised computers, they may be more difficult to detect using traditional approaches, and may deliver a more effective blow, especially when they spread into critical infrastructures, such as the traditional telephone network.

**References**

[AAM05] P. Akritidis, K. Anagnostakis, and E. P. Markatos. Efficient content-based worm detection. In *Proceedings of the 40th IEEE International Conference on Communications (ICC '05)*, 2005.

[AAM06] Elias Athanasopoulos, Kostas G. Anagnostakis, and Evangelos P. Markatos. Misusing unstructured p2p systems to perform dos attacks: The network that never forgets. In *ACNS*, pages 130–145, 2006.

[AAMA05] S. Antonatos, P. Akritidis, E. P. Markatos, and K. G. Anagnostakis. Defending against hitlist worms using network address space randomization. In *WORM '05: Proceedings of the 2005 ACM workshop on Rapid malcode*, pages 30–40, New York, NY, USA, 2005. ACM Press.

[BCJ+05a] M. Bailey, E. Cooke, F. Jahanian, J. Nazario, and D. Watson. The Internet Motion Sensor: A Distributed Blackhole Monitoring System. In *Proceedings of the* 12*th ISOC Symposium on Network and Distributed Systems Security (SNDSS)*, pages 167–179, February 2005.

[BCJ+05b] Michael Bailey, Evan Cooke, Farnam Jahanian, David Watson, and Jose Nazario. The blaster worm: Then and now. *IEEE Security and Privacy*, 3(4):26–31, 2005.

[BFV05] J. Bethencourt, J. Franklin, and M. Vernon. Mapping Internet Sensors With Probe Response Attacks. In *Proceedings of the* 14*th USENIX Security Symposium*, pages 193–208, August 2005.

[CBMM04] E. Cooke, M. Bailey, Z. M. Mao, and D. McPherson. Toward Understanding Distributed Blackhole Placement. In *Proceedings of the ACM Workshop on Rapid Malcode (WORM)*, pages 54–64, October 2004.

[DQG+04] D. Dagon, X. Qin, G. Gu, W. Lee, J. Grizzard, J. Levine, and H. Owen. HoneyStat: Local Worm Detection Using Honepots. In *Proceedings of the* 7*th International Symposium on Recent Advances in Intrusion Detection (RAID)*, pages 39–58, October 2004.

[DUMU03] Theo Detristan, Tyll Ulenspiegel, Yann Malcom, and Mynheer Underduk. Polymorphic shellcode engine using spectrum analysis. *Phrack*, 11(61), August 2003.

[K201] K2. ADMmutate, 2001. http://www.ktwo.ca/ADMmutate-0.8.4.tar.gz.

[KE03] Darrell M. Kienzle and Matthew C. Elder. Recent worms: a survey and trends. In *WORM '03: Proceedings of the 2003 ACM workshop on Rapid malcode*, pages 1–10, 2003.

[KK04] Hyang-Ah Kim and Brad Karp. Autograph: Toward automated, distributed worm signature detection. In *Proceedings of the 13th USENIX Security Symposium*, pages 271–286, 2004.

[LAAA06] V. T. Lam, S. Antonatos, P. Akritidis, and K. G. Anagnostakis. Puppetnets: Misusing web browsers as a distributed attack infrastructure. In *Proceedings of the 13th ACM Conference on Computers and Communications Security*, 2006.

[MD88] P. Mockapetris and K. J. Dunlap. Development of the domain name system. In *SIGCOMM '88: Symposium proceedings on Communications architectures and protocols*, pages 123–133, New York, NY, USA, 1988. ACM Press.

[MPS+03] David Moore, Vern Paxson, Stefan Savage, Colleen Shannon, Stuart Staniford, and Nicholas Weaver. Inside the slammer worm. *IEEE Security and Privacy*, 1(4):33–39, 2003.

[MSB02] David Moore, Colleen Shannon, and Jeffery Brown. Code-red: a case study on the spread and victims of an internet worm. In *IMW '02: Proceedings of the 2nd ACM SIGCOMM Workshop on Internet measurment*, pages 273–284, 2002.

[MVS01] D. Moore, G. Voelker, and S. Savage. Inferring Internet Denial-of-Service Activity. In *Proceedings of the* 10*th USENIX Security Symposium*, pages 9–22, August 2001.

[NR06] Naoum Naoumov and Keith Ross. Exploiting p2p systems for ddos attacks. In *InfoScale '06: Proceedings of the 1st international conference on Scalable information systems*, page 47, New York, NY, USA, 2006. ACM Press.

[RMT05] M. A. Rajab, F. Monrose, and A. Terzis. On the Effectiveness of Distributed Worm Monitoring. In *Proceedings of the* 14*th USENIX Security Symposium*, pages 225–237, August 2005.

[RMT06] M. A. Rajab, F. Monrose, and A. Terzis. Fast and Evasive Attacks: Highlighting the Challenges Ahead. In *Proceedings of the* 9*th International Symposium on Recent Advances in Intrusion Detection (RAID)*, pages 206–225, September 2006.

[Roe99] Martin Roesch. Snort: Lightweight intrusion detection for networks. In *LISA*, pages 229–238, 1999.

[SEVS04] Sumeet Singh, Cristian Estan, George Varghese, and Stefan Savage. Automated worm fingerprinting. In *OSDI*, pages 45–60, 2004.

[SF01] P´eter Sz¨or and Peter Ferrie. Hunting for metamorphic. In *Proceedings of the Virus Bulletin Conference*, pages 123–144, September 2001.

[SII05] Y. Shinoda, K. Ikai, and M. Itoh. Vulnerabilities of Passive Internet Threat Monitors. In *Proceedings of the* 14*th USENIX Security Symposium*, pages 209– 224, August 2005.

[SM04] Colleen Shannon and David Moore. The spread of the witty worm. *IEEE Security and Privacy*, 2(4):46–50, 2004.

[SMPW04] Stuart Staniford, David Moore, Vern Paxson, and Nicholas Weaver. The top speed of flash worms. In *WORM '04: Proceedings of the 2004 ACM workshop on Rapid malcode*, pages 33–42, 2004.

[Spi03] L. Spitzner. *Honeypots: Tracking Hackers*. Addison-Wesley, 2003.

[SPW02] Stuart Staniford, Vern Paxson, and Nicholas Weaver. How to 0wn the Internet in your spare time. In *Proc. 11th USENIX Security Symposium*, San Francisco, CA, August 2002.

[Sz¨o05] P´eter Sz¨or. *The Art of Computer Virus Research and Defense*. Addison-Wesley Professional, February 2005.

[VE06] Randal Vaughn and Gadi Evron. DNS Amplification Attacks (Preliminary Release), 2006. http://www.isotf.org/news/DNS-Amplification-Attacks.pdf.

[XCA+06] Konstantinos Xinidis, Ioannis Charitakis, Spyros Antonatos, Kostas G. Anagnostakis, and Evangelos P. Markatos. An active splitter architecture for intrusion detection and prevention. *IEEE Trans. Dependable Sec. Comput.*, 3(1):31–44,2006.

[YBP04] V. Yegneswaran, P. Barford, and D. Plonka. On the Design and Use of Internet Sinks for Network Abuse Monitoring. In *Proceedings of the* 7*th International Symposium on Recent Advances in Intrusion Detection (RAID)*, pages 146–165, October 2004.

# Issues in Monitoring and Threat Information Sharing in Global Networks- Alfonso Valdes

Alfonso Valdes (valdes@csl.sri.com)

Philip Porras

SRI International

Enterprise system administrators routinely employ defenses such as firewalls, switched networks, Intrusion Detection and Prevention systems (IDS/IPS) in support of security policy. This is a significant improvement over the past, yet new attacks still succeed against enterprise systems. Even though IDS/IPS has become the norm, security administrators are still reluctant to share vulnerability and security event data (for example, IDS alerts) across enterprise boundaries. This is true for a number of reasons, including the following:

- Disclosing alerts may disclose details of an enterprise's cyber defenses that may be useful to an attacker. Indeed, a sufficiently sophisticated adversary can conduct "probe/response" attacks precisely to determine the classes and sources of alerts resulting from a particular set of exploits.
- Disclosing details of a successful attack may lead to financial loss, regulatory sanctions, or damage to reputation.

In spite of these obstacles, there are potential benefits to the enterprise and to society as a whole from security event sharing. Security event information that is sufficiently content rich may enable identification and timely mitigation of emerging threats. It would be possible to assess the scale of an attack, from a few isolated systems to a particular enterprise, to a sector, or to the network as a whole. Alerts from multiple administrative domains can be correlated with each other as well as with other global defensive information, such as honeynet results and results of monitors closer to the backbone. There would be potential to propagate defensive measures (such as blacklists and new detection signatures) in a way that is presently not possible.

Sharing among entities in a particular economic or infrastructure sector potentially provides detection of the leading edge of an attack, or possibly trial runs of an attack ultimately intended for more targets within a sector. In the case of infrastructure systems, such as Process Control and SCADA systems, such leading edge or trial run detection can prevent not merely financial and reputation loss but catastrophic physical consequences as well.

Clearly, these benefits point to a need to overcome the obstacles to this sort of sharing. Under the Cyber TA project, SRI is leading a consortium to implement a Threat Operations Center to which user sites can register and share (contribute to and submit queries to the repository) in a privacy preserving manner. This center recognizes that contributors do not trust each other, and that traffic between contributors and the repository may be intercepted. We are active in four research thrusts to enable the development of this important capability.

- **Data and traffic anonymity**, including user-configurable services to provide data field anonymization and traffic source anonymity through the use of the Tor onion routing network.
- **Encrypted Computation**, exploring advances in encrypted computation to enable large-scale threat correlation of high-sensitivity encrypted data, without requiring decryption prior to analysis.
- **Malware Analysis and Mitigation**, through, for example, high-volume repository level correlation.

- **Reference Implementation of Threat Ops Center**, a live privacy-enabled web portal of live attack information contributed to by remote sensors that introduce new collaborative protection services to help defend operational networks from new large-scale attacks.

Infrastructure systems present particular challenges because anonymity in a computer science sense may not be adequate to protect the contributor. This is because there are a comparatively small number of such systems (although each system may have a large number of nodes) and comparatively high similarity between systems. Knowledge of even general information about an attack or attempted attack may narrow the possible number of candidates to a very small number.

Not all obstacles to such information sharing are technological. The global nature of the Internet results in a situation where permitted information sharing in one country violates privacy laws in another, for example.

We envision a potential future in which critical network communities, such as SCADA operators, can participate in collaborative security frameworks that may one day enable a new era of fast-reaction Internet defenses, or at a minimum allow these operators to comprehend the commonalities in how their collective community is being targeted by external intruders. However, we also think this vision is likely without significant progress in new large-scale threat reconnaissance technology that is designed with privacy-preservation inherent in their design. SRI's Cyber-TA project is providing a significant step toward addressing the technical challenges of privacy-enabled secure collaboration.

# Evaluating the Dependability & Security of Networked Systems – monitoring, operational assessment, auditing - Fabio Martinelli

Fabio Martinelli

National Research Council of Italy (IIT-CNR)

Fabio.Martinelli@iit.cnr.it

*Define 1 or 2 difficult problems (research challenges)*

The increasingly distributed, autonomous, open and heterogeneous nature of the current and future ICT-based systems challenges usual security mechanisms and demands for a coherent set of methodologies, techniques and tools to identify, assess, monitor and enforce "correct" system behaviour. In particular, for a truly pervasive computing and communications infrastructure, it is absolutely essential to provide a rigorous framework for informed decisions on trust issues. Thus, a main research challenge is the development of architectures and frameworks for *trust management* covering several applications areas.

*For each problem, describe the impact of solving or failing to solve the problem and a description of specific challenges making the problem difficult.*

The scale of the emerging global infrastructure, combined with the need for fully autonomous operation, surpass the usefulness of existing security infrastructures such as authorization services and certificate issuance and validation services. Having a certified identity (maybe granted using sloppy/undocumented procedures) in a dynamic and open environment does not a priori guarantee an acceptable behavior and performance. In particular, it is not enough for informed decisions on access restrictions and controls, selection among potential candidates for interaction, and even less adequate for reasoning about the expected behavior and dependability of entities for which no prior knowledge is available. Entities need to be distinguished not only based on their static (certified) identities but also based on their (un)expected, dynamically varying qualities that are relevant to the specific interaction context. Decisions are often based on directly verifiable evidence, but in a highly open system could be also based on indirect evidence reported by other entities. Trust should be ideally based on directly verifiable evidence, although when this is no possible, weaker forms may be established by using approaches as recommendation and reputation mechanisms. Definitely, the reported information must be carefully validated and reporting behavior monitored.

As a matter of fact, the presence of many sources of uncertainty as well as on threats due to the heterogeneity of the entities involved and the lack of centralized control, makes this research challenge ambitious. Although several successful attempts have been already applied (e.g. simple reputation systems adopted in e-commerce application), there is a need to fully study trust management processes for dynamic open systems and employ them in several ways, both to achieve a deeper level of protection and also for achieving better system performances.

For instance, by having such a capability of establishing, monitoring and using trust relationships, one could improve situation awareness. Indeed, a globally distributed monitoring infrastructure could be created on-demand by adopting existing monitoring infrastructures developed for other purposes. The credibility and the integrity of the information acquired could be assessed. Similarly, trust and reputation management based on behaviour observation may be considered as a way to stimulate collaborative behaviours in dynamic and mobile virtual coalitions. Thus using trust and reputation as a social control mechanisms (as in the human world) could be a way to isolate maliciously acting entities in the digital world. The development of global trust management frameworks and languages would benefit an effort of standardization in order to fully exploit their potentialities. Also trust management aspects are necessary in usual Service Level Agreements (SLA).

*Provide potential research directions for handling the problem, and metrics for measuring progress against the problem. Possibly, describe also the limitations of current approaches, promising approaches not yet fully explored and desirable approaches where relatively few ideas have been proposed.*

As a whole, we need to develop a set *methodologies, tools and architectures* to define, represent, evaluate, apply, monitor and negotiate trust in mobile, pervasive and open systems. This involves at least:

- *Rigorous and computational models of trust*. In particular, research is necessary to study how to compose trust measures as well as to validate these measures. Concepts from social sciences, probability and game theory, risk and utility analysis, economics and law should be adopted. Interdisciplinary approach would be very useful, in particular for assessing the merits of each approach and facilitate cross fertilization.
- *Adaptive trust policy frameworks.* Policy languages and mechanisms should be developed for continuously managing trust relationships as well as contractual agreements. These policies should depend also the current context. Trust-based policies for access and in particular, continuous usage control should be developed.
- *Advanced real-time and policy-based monitoring frameworks*. Both *proactive and reactive* monitoring techniques can be adopted, e.g. in case of negative events the trust information could be spread to other users in a reliable way. Behavioural policy languages would be useful for precisely expressing the allowed and the malicious behaviour (also incorporating some quantitative measures as probabilities an time). Also policy-based monitoring engines should be applied to correctly monitor the entity behaviour from its correct specification as well as detecting and tracking malicious behaviours (attacks). This would help both on modelling correct and malicious one. This would be particularly useful on service oriented architectures were workflow engines elaborate languages for service provisioning.
- *Metrics, simulation and validation techniques and tools* for establishing the relative merits of each trust frameworks. The concept of trust management is definitely useful for many communication and computation paradigms involving dynamic and mobile coalitions as service oriented architectures, mobile ad hoc networks, GRID, P2P. While many concepts could be shared, the different communication and computation capabilities force to define and establish different trust management frameworks.

# A Universal Instrumentation for the Network - L. Todd Heberlein

## Security Audit

Traditionally assessing the security posture of a site has been performed through the use of network-based vulnerability scanners such as the Nessus Vulnerability Scanner [Ness06]. A major value proposition of such an approach is that hardware or software only needs to be installed at one location reducing the manual labor of installing software on large numbers of hosts frequently controlled by numerous departments and individuals. Perimeter firewalls and network-based intrusion detection/prevention solutions share a similar value proposition.

Unfortunately, this dominate model of centrally controlled security has been obsolete for some time. Many penetrations into organization are not made through vulnerable servers, which a network-based scanner might be able to detect, but through vulnerable clients, which are generally not detectable from network-based scanners. Encryption and switched networks largely render network-based detection systems useless. Encryption, client-based attacks, mobile devices, and internal threats limit the value provided by perimeter firewalls.

Surely, a security assessment, no matter how sophisticated, based on centrally generated information will be wrong; garbage in, garbage out. Even subtle dynamic behavior is critical to a correct assessment. For example, suppose a user on machine A uses ssh to login to machine B. Machine B may have no vulnerabilities, but a compromise of machine A can put machine B at risk. A malicious agent can Trojan machine A's ssh client in order to capture the password to host B, or, even if one-time passwords are used on machine B, a malicious agent can hijack a connection to host B once it is established. Either way, despite no vulnerabilities on host B, host B can be placed at risk by host A. An accurate security assessment needs to understand this dynamic: which systems regularly grant what type of access to which other systems.

## The Protection of Information in Networks

There is also the issue of measuring risk assessment with respect to unknown threats. While this may sound infeasible (how can you measure yourself against something you don't know about?), good estimates may be possible. For example, at the extreme, a networked environment should only allow actions that need to be taken – where the networked environment considers all possible control surfaces in the network (internal application controls, kernel mediated actions, personal firewalls, network infrastructure access control lists, etc.). In other words, the network implements the Saltzer and Schroeder design principles of complete mediation, least privilege, and fail-safe defaults [SaSc74]. The security assessment of an actual network can be measured by how closely it comes to meeting this theoretical configuration.

As an example, most network services have an asymmetric behavior – clients behave one way while servers behave another. Typically a web server does not need to make requests other web servers, and if network control surfaces were set by default to reflect this (e.g., preventing a web server from making web requests like a web browser), many of the most famous worms such as Code Red and Slammer would become non-issues. The question that must be answered then is:

*If the Saltzer and Shroeder principles have been known for over three decades, and if operating systems and network infrastructure have supported mechanisms to enforce much*

*of Saltzer and Schroeder throughout the network, why aren't we taking advantage of them to build more resistant and robust networks?*

I believe one answer may be in another of Saltzer and Schroeder's principles: the principle of psychological acceptability. The number of control surfaces in a network is large, each control surface typically comes with its own language and semantics, the location and means to modify these control surfaces are often poorly known, the semantics of the controls are often counterintuitive, the interaction between the various control surfaces is not well understood, and reporting of when and why a control surface interferes with a legitimate operation and identifying how to correct the problem is poorly documented.

As an example, ZDNet reported this year, "The firewall in Windows Vista will, by default, have half its protection turned off because *that is what enterprise customers have requested*, according to the software giant" (emphasis added) [Kota06]. The article goes on to quote Zone Labs general manager, Laura Yecies, "For consumers, [configuring Vista's firewall] is challenging at best." I remember at a National Computer Security Conference panel in the early 1990s, when asked why Sun shipped their operating system with a '+' in the hosts.equiv file, a Sun representative replied that is what their customers wanted.

## Research Challenges

The research challenge: **Design a universal instrumentation architecture that can effectively harvest the necessary information to perform a reasonably accurate security assessment and can make a "Saltzer and Schroeder class" network infrastructure a reality.**

Develop a comprehensive list of all information necessary to perform a comprehensive security audit of an organization. Collect this comprehensive information at a number of sites, and then perform relevant security assessment. Repeat the exercise, but each time "knock out" some combination of data (e.g., lack of knowledge of client vulnerabilities), and then compare knock out results with comprehensive analysis results.

Develop a comprehensive list of all control surfaces in a network. Develop the means to measure what activity can be mediated by each control surface. Measure actual activity, and calculate the "gap" between what each control surface needs to allow and what it does allow.

Instrument all control surfaces to collect appropriate audit information so that each observable activity (e.g., packet observed on a wire or a write to a file) can be mapped to (1) the user that instigated the activity, (2) the person(s) who installed the relevant software, and (3) the person(s) who wrote the relevant software. Furthermore, all failures of expected system performance should be traceable to the set of mediating elements in the network that contributed to the unexpected failure.

Develop a unified language that can capture all the syntactic and semantic requirements of the various network control surfaces, audit their uses, and perform a reasonable level of local computation.

## References

[Kota06] M. Kotadia, "Vista firewall shackled due to customer demand: Microsoft", ZDNet, 26 April 2006.

[Ness06] "Nessus", http://www.nessus.org/

[SaSc74] J.H. Saltzer, M.D. Schroeder, "Protection of Information in Computer Systems", Communications of the ACM 17,7, July 1974.

# Model-based validation, deployment, configuration synthesis and control, measurement of large scale systems - András Pataricza

**András Pataricza**

**Budapest University of Technology and Economics**

## 1. Problem formulation

### Configuration synthesis

The growing use of the MDA and COTS platforms offering high-level services results in a simultaneous increase in design productivity and the quality of the target implementation. Initial steps were undertaken to integrate validation and verification into the system design workflow promising a proof of correctness and compliance to the user requirements.

However, less attention has been paid to elevate the design of operational configurations to the same level of automation despite the availability of industrial solutions for system monitoring and supervision (IBM Tivoli, HP Openview). Both performance and the protection of the systems require adaptive policies coping with overloads, physical and software faults and intrusions, but their design workflow is heuristic and no design automation or proof of correctness methods have been elaborated.

### Exploitation of measurement results in the design flow

A tremendous amount of applications were and are continuously designed with little or no concern for dependability aspects as the widespread use of dependability concepts did not become to an integrate part of the design workflow (with the exception of security during the last decade).

The situation is even more critical in the case of existing applications which need a proper support in dependability consolidation; otherwise a large amount of existing IP will become obsolete or deliver poor services despite of their appropriateness in the pure functional sense. Recent platforms generate a tremendous amount of log data but little or no automated support is provided to identify dependability bottlenecks and targeted reinforcement of the systems.

Due to the lack of knowledge extraction from the available empirical data no best practice is fed back to the design workflow of applications and systems.

## 2. Specific challenges

### Configuration synthesis

Two main challenges dominate this problem: modeling difficulties, V&V complexity due to the large size and structural adaptivity of the systems. In the case of modeling the lack of advanced analysis methods for empirical data, automated model and parameter extraction, premature modeling concepts and tools and the lack of integration with application design are the main problems despite of the existence of candidate solutions in other, remote fields of science targeting complex systems. The huge variety of real-life faults moreover, evolving threads require a quite flexible approach but no over-abstractions leading to overly pessimistic results and an intolerable level of redundancy.

Complexity related problems necessitate even basic research as new modeling and analysis paradigm should evolve explicitly supporting the concept of structurally adaptive, content sensitive networked systems (this is a prerequisite for a guaranteed QoS in autonomic computing, SOA, semantic web, ambient intelligence and other mobile applications as well). Obviously, this kind of analysis and advanced model-based software technologies should be integrated.

**Exploitation of measurement results in the design flow**

Intelligent processing of large sets of measurement results is a traditional topic in control theory but no bridge is built between dependability and control communities. Recent experience is made it clear that failover processes should be more carefully designed and the lack of a proper model-based verification is a major source of hazards. At the same time the large number of components involved necessitates either a large computing in power or a novel abstraction mechanism. Similarly, the counter parts of dominant factor estimations, mapping of empirical parameters to models, preferably generated by automated model structure identification algorithms is necessitated as well. As all these methodologies to be elaborated need an extensive experimental work the entire approach should rely on already existing basic technologies adapted for the new purpose.

## 3. Potential research directions

**Configuration synthesis**

By the very nature of the scope of the proposal made above, it relies on the reformulation and adaptation of existing intelligent data processing algorithms, methods and tools borrowed for instance from the field of financial analysis. Early results using standard, commercial OLAP and data-mining tools did already indicate their potential for processing a large amount of operational and artificially generated dependability related log data. At the same time, these experiments highlighted the necessity of bridging the gap between automatically generated phenomenological models describing the main factors influencing the QoS and architectural design models.

**Exploitation of measurement results in the design flow**

In short the counterpart of system identification methodologies has to be created for large scale networked systems. On the one hand, faithful continuous abstractions are needed in which the abstraction mechanisms eliminate the details scaling up with the increasing numbers units represented. On the other hand, experimental methods are needed to validate the faithfulness of the methodology.

## 4. Common research resources

Both target topics share the common point that they can rely on existing technologies which are candidates for the integration of the results. Similarly, even commercial of the shelf tools exist for the different kinds of extraction from the experience. What is needed is on the one hand the creation of a rather interdisciplinary team composed of the experts representing different technical and scientific areas like dependable computing, control theory, intelligent data processing etc., on the other hand real-time logs and experimental test beds are require for the validation of the approaches.

# Position Paper - Takashi Nanya

Takashi Nanya

Research Center for Advanced Science and Technology

University of Tokyo

4-6-1 Komaba, Meguro-ku, Tokyo 153-8904 Japan

&

Center for Research and Development Strategy (CRDS)

Japan Science and Technology Agency (JST)

3, Nibancho, Chiyoda-ku, Tokyo 102-0084 Japan

## 1. Difficult problems

Problem 1: Define a simple set of metrics that can represent the degree of dependability for services provided by networked systems from user's point of view so that the value of dependability can be made visible and mapped to economic values.

Problem 2: Establish a methodology to verify the trustworthiness of, and evaluate the dependability of, summarized and prioritized data from a huge amount of information produced and recorded on globally networked systems.

## 2. Impact of the problem and specific challenges

The impact of solving Problem 1 will give a strong incentive for industries to work toward dependability-oriented design and development that never happened in the past when only performance and functionality can be priced in the market of information systems. Specific challenges that make the problem difficult includes 1) human-made-fault modeling, and 2) dependency of the required level and quality of service dependability on applications, environments and conditions with which information systems provide their services.

It is essential for dependable and secure information societies that Problem 2 is successfully solved. The impact of failing to solve Problem 2 may introduce serious threats against dependability and security of the information society confronting "information-explosion". A specific challenge that makes the problem difficult lies in defining criteria on what correct information is.

## 3. Potential research directions and approaches

Solving Problem 1 requires to take not only established research approaches for identification of the set on metrics, and for methods of measurement, calculation, validation, evaluation, but also interdisciplinary approaches where the evaluation results services are made visible and get mapped to economic values. The latter approach may well be taken with international, both industrial and academic, collaborations toward global standards.

Problem 2 may require at least two steps of approaches. The first step should be to explore a new theory and technologies to handle an enormous amount of information produced and recorded on global networks everyday in the world to extract useful information for the society, business and international security. The second step should be to explore a methodology to verify the trustworthiness and evaluate the dependability of the summarized and most likely prioritized information.

# Panel F

## Future Test beds

# Future Test Beds - Jim Clarke

Jim Clarke, Waterford Institute of Technology, Co-Chair Panel F. Future Tesbeds.

**Overall Position on the summit**

The phenomenal growth of the internet as a commercial and communications tool has fundamentally changed the way business and citizens conduct their business. The emergence of an associated global ecosystem environment (e-business, e-finance, e-health, e-government, infotainment, etc.) presents many challenges and opportunities for European citizens. Of particular importance is the creation of an appropriate trust, security and dependability infrastructure for the internet of the future, which will protect European organisations and citizens against attacks, frauds and other misuses and abuses. In addition, any security policies should reflect Europe's regulatory and cultural environment. However, since this is a global environment, Europe must co-operate with key trust, security and dependability experts from International countries such as the USA, Canada, Australia, Asia, etc. to ensure that any internationally agreed trust, security and dependability framework reflects the European needs.

The **EU/US Summit on Cyber Trust: System Dependability & Security,** will, thus, bring together leading trust, security and dependability influencers from the European Union, USA, Canada, Japan and Australia to collaborate on the design and development of a Global Trust, Security and Dependability framework and roadmap.

Therefore, in my opinion, the key objectives of the summit are to kick start the following very important activities:-

1. Global Trust, Security & dependability framework and roadmap development. This will be the first of a series of intensive workshop events with the participant countries to:
    a. Identify key trust, security and dependability priorities and potential rollout timetables;
    b. Develop technology roadmap and rollout strategies, which reflect a balance between the participants.
    c. Identify and propose solutions to any legal impediments, which may restrict the participant's ability to lead the design and implementation of internationally agreed security solutions for the future internet.
    d. Identify and build upon existing projects and testbeds that could be used to facilitate international large-scale testing and evaluation of new dependability and security architectures, technologies, protocols, privacy protection mechanisms, etc., together with support towards global standards.

2. Organisation of an international trust, security and dependability project to further develop and promote the framework and roadmap.

**Position paper on Panel F. Future Testbeds.**

The Terms of reference of Panel F: Future Testbeds is "Examining the key challenges related to the establishment of interconnected future test beds constituting international large-scale experimental facilities for supporting the testing and evaluation of new dependability and security architectures, technologies, protocols, privacy protection mechanisms, etc., together with support towards global standards". Some of the challenges the panel will address during the summit will include:-

- identification of the technologies and systems including software, services, applications that will require the need for future testbeds;
- identification and examination of current existing projects, facilities and test beds and necessary developments required for re-use; and/or whether there is a need for new developments, testbeds, facilities, etc.;
- the need for joint benchmarks, test scenarios, interconnection issues;
- establish points of contacts, protocols, procedures to follow in making contact;
- the need for inclusive participation of ALL the right stakeholders (technology providers, service providers, academia, citizens, policy makers, standards bodies, others TBD)

As shown in figure 1, Panel F will be transversal to the other panels and will draw heavily and build upon the presentations and discussions from the other five panels.

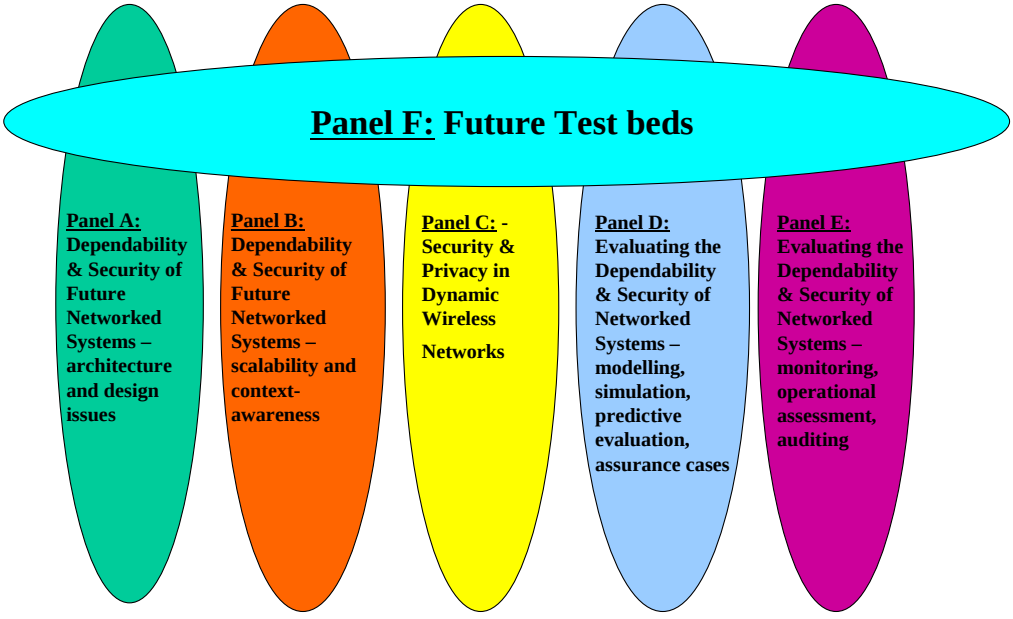


Figure 1. Positioning of Panel F w.r.t. other Panels A…E.

# Testbed Evolution – Zeta Dooly

Waterford Institute of Technology,

Ireland

zdooly@wit.ie

Test beds take many diverse forms and the spectrum of testing scenarios and requirements are diverse across US and Europe. As demand for services evolve, technology advancements and supporting infrastructures needs to meet this demand to provide adequate testing environments to deliver the desired level of quality services and products by all stakeholders including citizens, government, industry and academia.

Access to test data is often insufficient and difficult to obtain and is often only achieved to some extent by large organisations leaving the smaller businesses and innovators with little room for experimentation and structured testing exposure. There are a number of test bed initiatives emerging on a global scale, Planetlab and Geni for example but these are still at incubation stage where access to technical expertise is key to development and implementation and they are not yet at a stage where global collaboration and testing can be accessed and utilised easily by a wide community.

Test bed collaboration in the trust, security and dependability environment could expose a number of challenges specific to security service-oriented environments. Key challenges include:

- Confidentiality, trust of services that span multiple hosts.
- Privacy of data, protection mechanisms.
- Dependability and resilience of services that are subject to Fraud and attack.
- Stability of testing configurations and ease of re-configuration and simultaneous use and/or hand-over between different time-zones.
- Different regulatory environments.

The ability of a global test bed to facilitate diverse technologies will be key to the success of such an initiative and will include some of the following components:

- Virtual machines (such as Xen, SVISTA, Denali), virtual distributed environments.
- Web services (B2B services, eHealth, eGovernment).
- Rootkits, IDS, threat monitors, performance monitoring and management techniques, firewalls.
- Framework that could structure trust, security and dependability across regions and infrastructure.
- Mechanisms to support tool and technology evaluation processes throughout a cyclical, continuous testing process.
- Global testing standards.
- Techniques for trouble-shooting and fault tolerance management.

During this panel, challenges and opportunities in the Future Test Beds area will be explored and debated to identify the future requirements of a global testing environment that will complement existing work in this area.

# Enterprise Testbeds - Michael Bailey

Michael Bailey

University of Michigan

Ann Arbor, Michigan

As the needs of users and organizations change, network designers are constantly creating new communication methods, network protocols, and applications. Complete evaluation of these new techniques is often difficult in the context of existing production networks because they may rely on functionality that doesn't exist or, if they are deployable in current networks, may risk disruption of existing services. As a result, researchers have adopted a process of evaluation for these techniques that first involves analysis, then simulation or emulation, and finally small-scale experimentation in testbeds. More recently, new global-scale research testbeds have been proposed to build on the success of this process and address its limitations in order to allow researchers to re-conceive the Internet. While these testbed efforts provide a much-needed opportunity to reevaluate the way in which we bring together large numbers of individual networks through large-scale deployment and evaluation, the focus of these efforts is on the network core.

However, the edge of the network is already in the midst of a market driven re-evaluation. Over the past few years, enterprise networks undergone a rapid evolution fueled by huge increases in the number of hosts and applications as well as by the deployment of NATs, firewalls, and VPNs. Unlike the end-to-end reachability vision of yesterday's Internet, much of the Internet today is structured like a set of separate city states. Each state has its own internal organization, policies, and defended borders -- allowing little or no visibility into its operation from the outside. Enterprise networks no longer operate as miniature reflections of the broader Internet, and as a result, assumptions about topology, services, protocols, routing, and security that apply to one enterprise network do not apply to others. Understanding how these networks operate, and the impact of new techniques on their operation is critically important, as these networks are home to the sources and destinations that drive traffic across the Internet.

In this paper I advocate for a focus on testbeds to study new techniques for networking hosts within a single, very large enterprise networks. The different goals and properties of enterprise and backbone networks necessitate testbeds that account for these differences. For example, unlike the network backbone where the primary entity is a network device, enterprise testbeds consists of hundreds of thousands or even millions of individual users and hosts. Enterprise networks consist of a single administrative or policy organization, where backbones must rely on fragmented policy. The policies of the enterprise networks are primarily closed; in contrast to backbone networks where these policies are primarily open. In addition these enterprise networks consist of increasingly mobile users, while the backbone infrastructure remains primarily fixed.

The creation of new enterprise testbeds, or the expansion of existing testbeds to better account for the unique properties of enterprises, offer an excellent opportunity to explore the issues currently effecting enterprise networks. These include: Cache and proxy placement, security (e.g., malware propagation, attribution), telephony (e.g., VoIP, IM, email), policy enforcement and placement (e.g., virtual, ubiquitous firewalls), physical and virtual topologies (e.g., VLANs, MPLS), mobility (e.g., wireless), addressing (e.g., service based architectures, naming), critical infrastructure networks (e.g., financial networks, air traffic control) and their interfaces with the Internet, and measurement (e.g., DPI, host/network boundaries, heterogeneous data sources).

# Aiming for open and fair markets and balanced privacy - Pekka Nikander

Pekka Nikander
Helsinki Institute for Information Technology
Ericsson Research Nomadic Lab

**Abstract**

In this paper, we briefly discuss two major, high-level security-related challenges pertaining to global communication and distributed computing services. In a word, both challenges aim to create a sustainable networking infrastructure that would benefit the society. We aim towards a safe and ubiquitous communication framework that enables a competitive market for various kinds of services. A key issue is to ensure that the market will be free from negative network externalities.

The first challenge is architectural: We need to create a number of technical components that form a flexible and strong-enough technical framework for communications. The framework must allow open, fair, and sustainable *competition* between various players, where the players in different roles provide communication and other services. The main challenge here is on understanding the overall micro and macroeconomics of communication, where the communication services can be thought as being formed on digital telecom networks in a layered fashion. Another challenge is to make sure that such economic understanding is embedded in the design of the technical and regulatory means so that the various "layers" in the system are opened to real competition.

The second challenge is more technical but has broad architectural and system wide consequences, too. In order to prevent widespread and harmful price discrimination and other, Kafkaesque negative consequences of the loss of privacy, we need to provide the necessary technical, economic, and legislative means for balancing privacy and accountability in communications. There seems to be two main challenges here: finding the right *balance*, and finding a way to embed the necessary *technical means* to the future networks so that people's privacy cannot be illegally breached without their consent.

## 1. The challenge of creating open and fair markets

The two currently prevailing networking systems, i.e., the traditional telecom network and the Internet, have both been built without any conscious focus on creating an open and fair market place. The traditional telecom network was built with technical efficiency and widespread deployment in mind, consciously creating a number of regulated, regional monopolies. While the markets have been opened since, through legislative means, the structure of the technical system still shows its roots, leading to badly working markets in many cases. The Internet, on the other hand, grew from a co-operative research network, with little consideration of the incentives of the eventual commercial service providers and other stakeholders [CWSB02].

The next generation global network should be built with open, fair, and competitive markets in mind from the very beginning; relying on regulative and legislative means later on will not be enough. There are many reasons for this, and we spell out here but a few:

- As discussed amply and directly by Lawrence Lessig [Lessig99] and, from other points of view, e.g. by Thomas C. Shelling [Shelling78] and Geoffrey Moore [Moore99], any technology-based market with an economic networking effect becomes easily a closed one with one monopolistic major player. In other words, if customers prefer a closed proprietary solution for whatever

reasons, the economic networking effects are likely to direct the markets towards a monopolistic situation, where the technology owner has a considerable power over the markets due to excessively high switching costs. To counter such developments, the technology bases for the next generation global networks must be built from the ground up with the mindset of creating a competitive market place. With open, standardised technical interfaces that form well-thought, flexible business interfaces it becomes more likely that the resulting markets will have multiple players, reducing the probability of monopolies or oligopolies.

- The networking technology, including both hardware and software, deployed throughout the world forms a major asset whose upgrade and change is very expensive and hard due to the desire of compatibility. Changing any established networking technology is likely to take at least one technology generation (i.e. 5–8 years), or as observed so far, more likely two or more generations, i.e., in the order of 10–20 years. Consequently, if the markets get locked into a monopolistic or oligopolistic de facto technology standard, the society is likely to suffer from the negative consequences for a few decades.

- As shown by recent history in a number of crisis areas, efficient networking is critical to working civil society. Any repressive regime, being it corporate or governmental, seeks to limit people's ability to communicate with the external world. From this point of view, a network that is an amalgamation of mutually competing service providers is much harder to control and limit than a network that has one or just a few major players providing the vast majority of services. Consequently, by consciously creating an open network, where a number of service providers can compete at a fair market place, we indirectly support the very foundations of our civil society and democratic values.

**The difficulty of problems requiring multi-disciplinary approaches**

There are lots of difficulties even in trying to imagine how to proceed towards a new networking infrastructure based on the principles of openness, fairness, and sustainability. To begin with, we have little experience of truly open and fair communication markets. While one could claim that the World Wide Web is one such market place, the technical and organisational security foundations of WWW, i.e., the TLS protocol and underlying certification practises, are quite new and appear relatively weak when compared to, for example, the ones used in the existing banking system. Furthermore, even the WWW system contains a number of, strictly speaking, unnecessarily monopolistic technical structures, including the domain names and the system of TLS certificate authorities.

To take a broader view, the challenge of aiming towards a networking technology that forms the necessary technical and regulatory framework for a multitude of open and fair markets, each working at a different "layer" in the technology "protocol stack" and value chain, lies in the multi-dimensionality of the problem. Fundamentally, such a task requires technology experts that understand the details of communication and computing architectures, systems, and protocols. At the same time, economists are needed to understand, design, and model the potential effects of the technological choices that will form the foundations of the market place. Social and cognitive scientists are needed to anticipate the cognitive and emotional reactions of the general public to the new markets and services, including aspects such as the balance of powers, perceived safety and security, as well as trustworthiness and reputation. Finally, legal and regulatory understanding is needed to design the regulatory side of the markets.

**Understanding new networking primitives**

Recently, a number of prominent network researchers around the world have voiced, in various forms, an opinion that we need to seriously reconsider the communications architectures and even the basic communication primitives, as they are seen today [BLRSIW04] [vanJacobson06], [Shenker06]. From a technical point of view, it seems likely that neither the connectionless send-receive model, as used in the Internet, nor the connection-based model, as used in the traditional telecom network, may be inappropriate for a large number of existing and forthcoming network services. Instead, for example, we might consider a networking architecture based on the subscribe–publish model, where receivers subscribe to the data items or streams they want to receive [CKLRTVSS06].

If both the network architecture and primitives are to be revised, the resulting communications market place will be fundamentally different from the present ones. Considering, for example, the publish–subscribe primitives, an interesting question emerges: where is the value? Sometimes the value seems to be in being able to subscribe; e.g., subscribing to digital content. Sometimes, on the other hand, the value seems to be in the ability of publishing to a specific audience; e.g., publishing targeted product announcements. Hence, even at the level of the very primitives and primary communicating peers, it is an open question how the compensation structures should be modelled. If we further consider the other stakeholders, including access providers, backbone providers, caching and computing services, etc., and aim for open markets in all different service and provider categories, understanding the different potential value chains become tricky. As the providers will naturally aim for strong market positions, it becomes a major challenge to design the technical interfaces in such a way that they naturally support open and fair competition.

**A possible way forward**

We seem to be on a very early stage in addressing this architectural challenge. As stated above, a number of leading researchers seem to more-or-less agree that the current communication architectures and primitives are unlikely to be suitable for a large fraction of the future communication needs, but that seems to be roughly as far as we are. That is, a number of people are calling for a “clean slate” design, including a number of public bodies that have decided to fund such fundamental research. Consequently, the basic research in the area, including development, experimentation, and evaluation of new architectural concepts and new primitives, needs to be continued. However, we believe that it would be beneficial to enrich the current activities by creating a number of related activities that would focus on the economic, business, security, privacy, social, and cognitive aspects of the problem space.

## 2. The challenge of balancing privacy and accountability

With the foreseeable emergence of ubiquitous computing and networking, discussion around privacy has been fervent. A number of people have voiced their opinion that privacy is not needed and we should aim for a fully transparent society (e.g. [Brin99]). At the same time, it has been argued that loss of privacy is not only a social problem (e.g.[Solove04]) but also an economic one [Odlyzko03]. In general, it seems fair to claim that loss of privacy will be a genuine problem at least if it leads to wide spread, unacceptable forms of price discrimination, if it leads to a situation where the powerful can afford privacy but the average citizens do not, or if it leads to erosion of the trust of users to the networked services.

At the same time, wide spread unsolicited electronic mail (spam) and different kinds of electronic frauds (phishing, identity theft) have shown the perils of too strong anonymity, or rather, too weak reputation, authentication, and accountability mechanisms. Hence, while it seems reasonable to require stronger privacy due to a number of social and economic reasons, at the same time we need to require enhanced accountability. While this may sound like a contradiction, it looks like that existing cryptographic

methods could be combined in a way that provides simultaneously relatively strong privacy and the possibility of breaking that privacy, through explicit means, in the case that accountability requires it.

**Striking a balance or enhancing both**

As Odlyzko argues [Odlyzko03], wide spread aggressive price discrimination, made possible by loss of electronic privacy combined with moving of commerce onto the networks, is likely to undercut the foundations of our current economy, including such fundamental and well understood concepts as market price: “Prices that depend on the buyer would require a complete rethinking of [the notion of market price].” On the other hand, certain forms of price discrimination are usually considered to be good for the overall economy [Phlips83]. A third factor here is founded in human behaviour. People have a strong instinctive notion of economic and social fairness (consider e.g. [FS99], [HBBCFG04]), and overt price discrimination, in most forms, is considered very unfair. Consequently, considering the strong incentives businesses have for price discrimination, it seems likely that covert forms of price discrimination will exist where-ever technically possible, leading to strong pressures to reduce privacy even more.

A potential consequence of this and other prevailing trends is that only the rich and powerful can afford privacy, in terms of both electronic and physical presence. As argued, for example, by Solove, that in turn easily leads to a situation where the powerful become increasingly so due to the information asymmetry [Solove04]. (Indeed, the same argument is used by some openness proponents to argue for *complete* erosion of privacy. However, we believe that those who can afford would always be willing to pay for their privacy, making a completely transparent society an unreachable utopia.) Hence, we believe that in order to retain the relative balance of powers available in our current democratic societies, it is extremely important to make sure that effective and simple-to-use technical means for preserving personal information privacy will be available for all citizens.

While preservation of the prevailing level of physical privacy and simultaneous enhancement of electronic privacy seems important for sustainable economy and democracy, the current possibilities for accountability in the Internet are far from at a desirable level. Thus, almost paradoxically, we have both too little and too much privacy, at the same time, when working through the current network structures. The lack of accountability is amply shown by the current amounts of distributed denial-of-service attacks, spam, and phishing.

At the outset, there seems to be two different factors that contribute to the lack of accountability. Firstly, at the IP layer all Internet traffic is almost anonymous. That is, while the IP packets carry the so-called source address field, this field is no longer sufficient to identify the sender, due to e.g. ease of source address spoofing, proliferation of network address translation, and other active middle boxes. Secondly, the majority of unwanted IP packets are sent by zombies, or computers running programs unknown to their owners. These two factors, combined with the almost non-existing framework for tracking down the real life identity of the senders even in the cases where the IP address happens to be an accurate account of the source, make it very costly to find out the real source of any unwanted traffic. Finally, this situation must be contrasted with legitimate traffic, where the IP address does still function as a fairly reliable partial index of the sender’s identity, thereby allowing web sites and other services to track the sender’s probable identity even when higher layer data, such as HTTP cookies, are not available.

Hence, we believe that the future networking architectures must strengthen both privacy and accountability. In other words, we believe that in future communication systems those parties that know users’ real-life or other long-term identity should not be able to observe what the users do in the network, i.e., what services they use, and at the same time the service providers should not know the users’ real-life identity, location, or other attributes, unless they explicitly decide to reveal them.

**Existing and potential approaches**

In [ANN05], we presented one technical approach, based on pseudo-random sequences, for enhancing privacy through the existing Internet protocol stack. Basically, the same approach could be applied to any open communication system. If combined with techniques suitable for location privacy, such as indirection [SAZSS02], the methods results in an architectural "insulation layer" between inter-networking services and higher layer-services. That is, the network nodes responsible for actually passing the traffic cannot, by themselves, know who is the final destination, sender, or the content of the traffic, due to the anonymous and temporary addresses and encrypted content. At the same time, the peer nodes, by themselves, cannot know exactly where their peers are in the network due to indirection. However, by combining information from the access nodes, indirection infrastructure, and optionally other service nodes, the temporary identifiers can be strongly bound to long-term identifiers, which in turn could administratively be associated with real-life users.

While we believe that something like the combination outlined above might form a technical base for combined privacy and accountability for future networks, a lot remains to be done. Firstly, it is far from clear where the balance between privacy and accountability should lie. Hence, it might be desirable to create frameworks that allow the balance to be technically shifted, as demanded by regulation. At the same time, it must be remembered that encryption and indirection both are technologies that are extremely hard to forbid. Therefore, if the balance between privacy and accountability, as offered by the network, is generally felt unacceptable, a sizeable fraction of the users may seek for overlay technologies that allow them to shift the balance for their benefit. Secondly, even though we may have a glimpse of a technical solution, there remains challenges in integrating it to the kind of new network primitives discussed in Section 1, in figuring out if the solution is sufficient and generally acceptable, and there always lies the possibility of even better solutions.

## Summary

In this position paper, we have discussed the desirability of and the challenges involved in building a new communication framework that would form an open and fair playground for the various stakeholders. We have briefly envisioned a future where new networking primitives form a base for a new networking structure, with well-thought technical interfaces that form, together with new compensation instruments, a foundation upon which a multitude of open and fair market places could be based on. As a small but important part of that, we have discussed in a little bit more detail the necessity of strengthening both privacy and accountability, in order to avoid the negative effects of price discrimination and shifts of power resulting from increasing information asymmetries.

[ANN05] Jari Arkko, Pekka Nikander, and Mats Näslund, "Enhancing Privacy with Shared Pseudo Random Sequences," in Security Protocols, 13th International Workshop, Cambridge, 20-22 April, 2005.

[BLRSIW04] Hari Balakrishnan, Karthik Lakshaminarayanan, Sylvia Ratnasamy, Scott Shenker, Ion Stoica, Michael Walfish, "A layered naming architecture for the Internet", SIGCOMM 2004.

[Brin99] G. David Brin, "The Transparent Society: Will Technology Force Us to Choose Between Privacy and Freedom?", Perseus Books, 1999.

[CKLRTVSS06] Mohit Chawla, Teemu Koponen, Karthik Lakshminarayanan, Anirudh Ramachandran, Arsalan Tavakoli, Atul Vasu, Scott Shenker, and Ion Stoica, "Data-Oriented Network Architecture (DONA)", unpublished manuscript submitted to HotNets 2006.

[CWSB02] David D. Clark, John Wroclavski, Karen R. Sollins, and Robert Braden, "Tussle in Cyberspace: Defining Tomorrow's Internet", SIGCOMM 2002.

[FS99] Ernst Fehr and Klaus M. Schmidt, "A theory of fairness, competition, and cooperation," *the Quarterly Jornal of Economics*, 114, 817–868, 1999.

[HBBCFG04] Joseph Henrich, Robert Boyd, Samuel Bowles, Colin Camerer, Ernst Fehr, and Herbert Gintis (2004), *Foundations of Human Sociality: Economic Experiments and Ethnographic Evidence from Fifteen Small-Scale Societies*. Oxford University Press

[vanJacobson06] Van Jacobson, "If a Clean Slate is the solution, what was the problem?", Stanford University Clean Slate Seminar, 2006.

[Lessig99] Lawrence Lessig, *Code and Other Laws of Cyberspace*, Basic Books, 1999.

[Moore99] Geoffrey A. Moore, *Crossing the Chasm: Marketing and Selling High-Tech Products to Mainstream Customers*, Harper Business, 1999.

[Odlyzko03] Andrew Odlyzko, "Privacy, Economics, and Price Discrimination on the Internet," *ICEC2003: Fifth International Conference on Electronic Commerce*, N. Sadeh, ed., ACM Press, 2003.

[Phlips83] L. Phlips. The Economics of Price Discrimination. Cambridge Univ. Press, 1983.

[Schelling78] Thomas C. Schelling, *Micromotives and Macrobehavior*, W. W. Norton & Company, 1978

[Shenker06] Scott Shenker, "Rethinking the Internet Architecture", Stanford University Clean Slate Seminar, 2006.

[Solove04] Daniel J. Solove, "The Digital Person: Technology and Privacy in the Information Age", New York University Press, 2004.

[SAZSS02] Ion Stoica, Daniel Adkins, Shelley Zhuang, Scott Shenker, Sonesh Surana, "Internet Indirection Infrastructure," Proceedings of ACM SIGCOMM, August, 2002.

# Designing and Testing Networked Embedded Control Systems – Anthony Joseph

Anthony D. Joseph

Professor, EECS Department

University of California, Berkeley

The design, development, and testing of next-generation networked embedded controllers for SCADA/DCS systems is an expensive and complex multi-year process that introduces many pressing challenges for systems designers, developers, and vendors, and for the physical plant owners/operators. Many of these systems are at the heart of critical infrastructure components, such as power production and distribution systems, chemical processing plants, petroleum and gas refinery and distribution facilities, transportation systems, etc. One of the key differences between these systems and Information Technology systems is that damage to this infrastructure can ultimately lead to severe loss of life.

This cyber critical infrastructure is replaced or upgraded on three to twelve year cycles; as such, these systems represent significant capital and time investments, and, present significant opportunities for cyber vulnerability analysis and exploitation by malicious parties. The complexity of networked embedded controller-based systems is growing as these systems are increasingly being interconnected with enterprise networks, often for Total Quality Management purposes, and directly or indirectly with public networks, such as the Internet, for remote management purposes. Interconnection with enterprise networks introduces two threats: exposure to compromised desktop machines, and malic ious insiders. Desktop machines are vulnerable to compromise by e-mail- or removable media -borne worms and viruses; once compromised, they can be used to attack networked embedded controller systems. Similarly, the insider threat is also a growing problem. Paper design and analysis, along with small-scale emulation, are insufficient tools for addressing these threats, as they may not capture all of the interaction effects found in a large, real system. Failure to solve these systems design problems will allow the introduction of vulnerabilities that may be exploited by malicious parties, with serious consequences.

There are existing small-scale testbeds for evaluating networked embedded controllers; however we posit that incorporating extensive, realistic, large-scale emulation and testing of these systems into the design and development process would be an important and beneficial improvement. A key goal of this testing would be to anticipate types and scales of cyber attacks against systems under development, and to enable the accurate emulation of enterprise networks and embedded network controller systems. Another goal would be to identify best practices for the testing process.

The cyber DEfense Technology Experimental Research (DETER) testbed [ACM04, DETER06], supported by NSF and DHS, is a general-purpose experimental testbed for cybersecurity research and education in use since 2003 by academic, industrial, and government experimenters. As the largest open, free experimental facility dedicated to cyber security research, the testbed provides a unique, powerful tool for experiments with resource needs preventing them from being performed anywhere else. DETER is also a teaching platform for cybersecurity classes. DETER safely supports cybersecurity experiments, including experiments on “risky” code that cannot be performed in the Internet because of traffic volumes or the risk of escape. Examples for which DETER has already been useful include DDoS attack dynamics and defenses, virus and worm propagation and defenses, and attacks on network routing infrastructure. The DETER testbed, based upon Utah’s Emulab software, allows remote experimenters to allocate large numbers of nodes, link them with nearly-arbitrary topologies, load arbitrary code for routing, traffic

generation, defense mechanisms, and measurement tools, and execute their experiments.

We believe that the DETER testbed could be an excellent starting point for building a public testbed where all interested parties could design, model, and test their systems. In addition, the testbed would be an ideal place to store libraries of testing strategies and data. By using Field Programmable Gate Array technology, it should be possible to emulate the behaviors of actual networked embedded controllers, sensors, actuators, and the corresponding control/data traffic. The size of testbeds is often limited by many factors, such as cooling, power, space, and weight requirements. These limitations can be partially addressed by building multiple smaller testbeds, however the scope and scale of experiments is limited. To emulate particularly large systems (e.g., a large regional power grid) with tens of thousands of sensors and actuators, we propose the dynamic federation of multiple networked embedded controlle r testbeds into a single very large testbed (or for parallel development, multiple large testbeds). Such an approach offers administrative and flexibility, addresses the limitations associated with building very large testbeds, and it provides an on-demand ability to scale to large experiments. We are developing support for federation of multiple testbeds into a single logical testing environment. A testbed can also be used for other important applications – periodic operator training and system red-teaming. Both applications are challenging, but necessary requirements given continually changing systems and threat models. By using a testbed, owners/operators can explore different console screen application designs and layouts, and they can gauge operator reactions to various scenarios (representing both normal, abnormal, and attack situations). The DETER testbed has already been used for training and classroom exercises, and could easily be adapted for operator training and system red-teaming tasks.

We believe that a five to ten year research plan in this area should focus on two key areas:

· **Building multiple testbeds where developers and owners/operators can explore and test new devices, protocols, architectures, and applications.** We need to develop testbeds that incorporate actual networked embedded controllers, FPGA- and software-based controller emulators, and actual and emulated software plant control systems. The testbeds should also be capable of being federated into one or more large testbeds for large-scale experiments and tests.

· **Improving the training of the operators and security defenders responsible for these systems.** We need to collect best-practices for evaluating operators' behaviors and reactions, and deploy traffic/scenario generation utilities. In addition, we need to emulate real-world networked embedded controller installations and work with corporate and government agencies to develop comprehensive red-team training software and systems.

[DETER06] Terry Benzel, Bob Braden, Dongho Kim, Clifford Neuman, Anthony D. Joseph, and

Keith Sklower, Ron Ostrenga, and Stephen Schwab, *Experience with DETER: A Testbed*

*for Security Research.* Second IEEE Conference on Testbeds and Research

Infrastructures for the Development of Networks and Communities (TridentCom2006),

Barcelona, March 2006.

[ACM04] DETER/EMIST team, *Cyber Defense Technology Networking and Evaluation,*

Communications of the ACM, March 2004.

# Future Test Beds - Ollie Whitehouse

Ollie Whitehouse, Architect, Advanced Threat Research,

Symantec Security Response, Symantec Corporation

## Introduction

The following position paper deals with the subject of "*Future Test Beds: joint benchmarks, test scenarios and interconnected test beds for assessing new dependability and security architectures, technologies, protocols, privacy protection mechanisms, etc., together with support towards global standards.*". Currently there is still little to no standardization with regards to assessing the effectiveness of security architectures, technologies or the security of a protocol outside of national Governments[11]. The standards and testing criteria's which do exist are typically designed to assess the effectiveness of an overall system[12]. Alternatively these standards deal with a certain configuration or specific component of a system[13] to ensure it is fit for purpose.

When dealing with subject of test bed design to assess any system in terms of generic security or dependability there are many facets which need to be taken into consideration. These range from the standard interoperability testing which is normally conducted by either Special Internet Groups or formal vendor integration labs through to configuration assessments and more aggressive failure testing.

Future test beds will be required to provide a degree of testing not seen out of formal product assessments and on a scale never seen before. The services required from a test bed will include assessing configuration scenarios which were not envisaged or catered for during development. Assessment of implementation issues with the technology or protocol which require deep knowledge of target in question. Additionally the ability to assess the underlying technology related to components which make up part of the system but were not deemed appropriate to assess during the criteria evaluation.

The key to testing any system be it an architecture, technology or protocol is to look at it as a whole as well as in a number of different circumstances, configurations and environments. Only by taking this holistic approach can future test beds attempt to ensure that the results they produce are of value to widest audience while also providing the highest level of assurance that a majority if not all common vulnerabilities or failure scenarios have been both assessed for and detected.

**Research Challenges**

There are a number of unique research challenges when considering the design and implementation of future test beds. The goal of these test beds is obviously unique with a specific goal of understanding the security and dependability of a test subject.

- Developing appropriate benchmarks for the testing criteria which represent the appropriate impact of any failures in terms of security
- Developing a scalable methodology and supporting toolset for the assessment of architectures and systems components that can not be assessed from a network perspective

[11] http://www.cesg.gov.uk/site/iacs/index.cfm

[12] http://www.cesg.gov.uk/site/iacs/index.cfm?menuSelected=2&displayPage=2

[13] http://www.cesg.gov.uk/site/iacs/index.cfm?menuSelected=5&displayPage=5

- Developing an effective low cost means of implementing protocols. This is so they can be thoroughly tested in all states independent of vendor, implementation or standard in an economical way
- Development of a mechanism and supporting methodology for the testing of proprietary protocols which have no public documentation available and are not standards based
- Understanding the benefits of moving to "interconnected test-beds constituting international large-scale experimental facilities" versus smaller technology specific test beds

All of these present their own unique challenges. While a number of these have been addressed previously in the academic, public and private sectors they have never been extended to be truly technology and environment agnostic. For example the issue of network protocol testing has been approached in a number ways from the generic[14] to the specific[15] [16]. With the later of these being the most thorough in terms of coverage but also being the most resource intensive to develop and thus currently unsuitable for large scale test bed deployment. Oulu has performed a significant amount of research through their PROTOS[17] [18] program into eliminating software vulnerabilities in common network protocols. This has resulted in collaboration with national government critical national infrastructure organizations such as NISCC[19] in order to set priorities for further areas of research. However there are significant challenges in making this scale due to the amount of effort required to research, develop and implement the test cases.

**Benefits**
**Introduction**

Below is a position on a hybrid test bed that could integrate into existing test beds already deployed and actively used [20] [21] [22]. The purpose of which is to provide a new specific type of security test bed environment with unique features designed to make it ultimately both scalable and thorough. The areas discussed below are not designed to deal with all areas of dependability such as distributed enterprise security management testing, DDoS experimentation or experimenting with next generation internet protocols beyond IPv6. Instead it is designed to address a number of medium term challenges facing networks being deployed. That is the security of the devices and protocols which make up these networks.

**Benchmark Development**

Simply counting the number of times a system fails in testing does little to quantify the severity or breadth of user impact and help vendors prioritize which flaws should be fixed first. So to ensure that the benchmarks are meaningful in terms of security an appropriate set of benchmarks should developed

[14] http://www.immunitysec.com/downloads/advantages_of_block_based_analysis.pdf
[15] http://www.ee.oulu.fi/research/ouspg/protos/ PROTOS - Security Testing of Protocol Implementations
[16] http://imj.gatech.edu/papers/ISC-06.pdf.gz SNOOZE: toward a Stateful NetwOrk prOtocol fuzZEr
[17] http://www.ee.oulu.fi/research/ouspg/protos/sota/matine/method-thesis/di.pdf - A COLLABORATIVE METHOD FOR ASSESSING THE DEPENDENCIES OF CRITICAL INFORMATION INFRASTRUCTURES
[18] http://www.ee.oulu.fi/research/ouspg/protos/ PROTOS - Security Testing of Protocol Implementations
[19] http://www.niscc.gov.uk/ National Infrastructure Security Co-ordination Centre
[20] http://www.isi.edu/deter/ - DETER - A Laboratory for Security Research
[21] http://www.geant.net/ / http://www.geant2.net/ - GÉANT / GÉANT2
[22] http://www.geni.net/ - Global Environment for Network Innovations

however to do this a formal model must be used to categorize the observed failures. Two existing examples of models which already exist for application threat profiling are DREAD and STRIDE

The DREAD model:

- **D**amage potential
- **R**eproducibility
- **E**xploitability
- **A**ffected Users
- **D**iscoverability

The STRIDE model:

- **S**poofing
- **T**ampering
- **R**epudiation
- **I**nformation disclosure
- **D**enial of service (DoS)
- **E**levation of privilege

Using a similar approach will aid in the process of communicating the different aspects of the security testing conducted. This in turn allows the recipient of the data to understand where a product may have issues and then based on these areas either make appropriate risk management changes to address them or alternatively select a different architecture, technology or protocol. This in turn would be used to develop the benchmarks for the given technology or protocol.

Failure to a develop model by which the results can be categorized in order to benchmark will result in output that will require further validation. Or alternatively further investigation and thus offer little value to the end recipient. The benefit of developing a model would be a universally accepted system by which architectures, technologies and protocols can have their test bed results categorized. This in term will allow the development of appropriate benchmarks.

**Methodology for Non Network Assessable Components**

It should be accepted that there will be elements of any architecture or technology which have a significant bearing of the security of the overall system that can not be assessed from a network perspective. As a result there should be research performed in to the most pragmatic means of conducting a thorough assessment of these components in a scalable fashion. This will be one of the more complex research projects due to the vast array of different technologies which could come into the test bed for assessment.

Failure to assess the system in its entirety could result in low hanging fruit vulnerabilities which could significantly affect the overall benchmarks results of a test subject.

**Standard Based Stateful Protocol Testing Suite**

As previously mentioned this area has been and continues to be researched, yet there is still no universal mechanism for developing data driven protocol testing suits that also include state machines. This is extremely complex as it has to deal with not only with multiple layers of the OSI model but also components which may leverage encryption or encoding schemes.

Failure to develop such a system will make any protocol testing cost prohibitive to perform to any great depth due to the time and effort required to research and implement a working protocol stack before test cases can be developed.

**Proprietary Protocol Testing Suite**

During the course of an assessment there will be times where proprietary protocols or proprietary extensions to standards based protocols will be encountered. This instances may also include situations where the vendor either can not due to licensing restrictions or is not willing to release the protocol specification to the test bed. As with standards based protocols this area has been researched to some degree[23] [24] including some novel research into using Bioinformatics in protocol reverse engineering[25]. However if the proposed future test beds wish to ensure depth in their testing this research will need to be formalized and productized in a system which can scale beyond simplistic stateless protocols. This problem will also be significantly complicated by proprietary protocols which may leverage encryption technologies.

Failure to address this problem will result in the test beds in being ineffective in assessing components which provide exposure to a system. The fact that these protocols are proprietary will not hinder a motivated attack who wishes reverse engineer them in order to deploy in their own security test bed.

**Direction & Metrics**

As mentioned throughout this position paper there have been a number of historical approaches to perform assessments in a test bed environment of architectures, technologies and protocols. However none of these will currently scale to any significant size without requiring highly skilled individuals (the only exception being AutoDafe[26]) to develop the testing tools and supporting test cases, the result of which is there is a requirement for research to be performed into developing a easy to use and scalable data driven protocol testing tools for the use in test beds.

[23] http://research.microsoft.com/workshops/sysml/papers/sysml-Gopalratnam.pdf Automatically Extracting Fields from Unknown Network Protocols
[24] http://www.ub.utwente.nl/webdocs/ctit/1/000000ef.pdf Assessing Unknown Network Traffic
[25] http://www.4tphi.net/~awalters/PI/pi.pdf Network Protocol Analysis using Bioinformatics Algorithms
[26] http://autodafe.sourceforge.net AUTODAFE: an Act of Software Torture

# Future Test Beds - James E. Just

Define 1 or 2 difficult problems (research challenges)

Since I'm participating in the Future Test Beds session and my experience is in developing and testing defensive technologies, I will concentrate on near-real-time monitoring of systems and attacks against them in my briefing. Below are my thoughts on significant research challenges.

1. Predicting with reasonable accuracy the impacts of large scale cyber attacks and/or large scale cyber defenses based upon results from testing.

2. Protecting one or more networked COTS computers connected to the Internet) from all forms of remote attacks with acceptable user impacts.

3. Protecting COTS computers from Trojans with acceptable user impacts.

For each problem, describe the impact of solving or failing to solve the problem and a description of specific challenges making the problem difficult.

1. Without such a model or science, it is impossible to provide a meaningful answer to either "Is it worthwhile to spend $X of public or private funds to implement the set Y of defense technologies" or "Is defense A better than defense B given individual and unique test data from each" without massive testing efforts. As Tanenbaum states in his *Computer Networks*, "Unfortunately, understanding network performance is more of an art than a science. There is little underlying theory that is actually of any use in practice. The best we can do is give rules of thumb gained from hard experience and present examples taken from the real world."[27]

2. While there are great incentives to prevent or mitigate attacks at the network level, individual hosts (computers, cell phones, switches, etc.) are the units that are attacked. They present the greatest opportunity for identifying that an attack is underway, understanding details of the attack and stopping it early. It is important to think way beyond the Code Red and Slammer types of attacks. Also important are ways to mitigate damage from attacks and significantly speed up recovery from them.

3. Even if hosts can be protected from remote attacks, the problem of defending against malicious code that is already running on the host (e.g., a Trojan) is very large.

Provide potential research directions for handling the problem, and metrics for measuring progress against the problem. Possibly, describe also the limitations of current approaches, promising approaches not yet fully explored and desirable approaches where relatively few ideas have been proposed.

1. Fundamental research into what is it important to predict about the effectiveness of attacks and defenses and what can be predicted?

2. Detection of corrupted processes and threads; isolation of threads and processes.

3. Detection at source code, binary and running process levels.

[27] Source: A.S. Tanenbaum, *Computer Networks* (3rd Ed., Prentice Hall, 1996, p. 555,556)

# A Test Bed for Software and Services - Aad van Moorsel

Aad van Moorsel

Newcastle University, UK

The objective of a test bed for software and services is to allow experimentation at the application and services level.

Service-oriented software and utility-computing inspired service provisioning systems are being developed in many places. Realistically, the complexity and scale of such systems makes it difficult for academia and small/medium businesses to effectively try out their ideas. A test bed that allows realistic experiments to be run would open up opportunities for these parties to venture into service-oriented solutions.

Existing world-wide test beds Planetlab and Geni (Global Environment for Network Innovation), target, as the name explicitly indicates in the second case, network innovation. These platforms allow for almost unrestricted experimentation, but necessitates every user to install application level software from scratch. Reuse of application level solutions is difficult.

The envisaged test bed would allow for experiments such as the following:

- experimentation with security policies for computing service providers
- experimentation with bandwidth and computing reservation algorithms in fully virtualised environments (as in Violin)
- experimentation with fault tolerance solutions such as mediators using realistic web services
- experimentation with new services deployment solutions on top of virtualised environments (services as in GOLD and other projects)

The kind of technologies we envision being provided by a software and services test bed are, among others:

- virtual machines (such as Xen), virtual distributed environments (as for instance proposed in Violin)
- realistic Internet services (B2B services, medical applications, bioinformatics services, etc.), using various technologies (web services, REST)
- load balancers, firewalls, and other parts of realistic service provisioning systems
- service deployment software (as for instance in DynaSOAR or Smartfrog.org)

There are various technical challenges to be addressed before a services test bed can materialise—this proposal simply encourages debate about such software and service test bed.

# Beyond Test Beds – Henrique Madeira

Henrique Madeira
University of Coimbra - Portugal

henrique@dei.uc.pt

## The Problem

In spite of the effort put on the development of adequate test beds and the intensive research devoted to the mitigation of key problems such as experiment representativeness, intrusiveness and portability of tools, just to name a few, two important questions remain largely unanswered:

- How to analyze the usually large amount of raw data produced in dependability and security evaluation experiments (specific test beds and field data experiments), especially when the analysis is complex and has to take into account many aspects of the experimental setup (e.g., target systems, configurations, workload, etc)?
- How to compare results from different experiments or results from similar experiments across different systems if the tools, data formats, and the setup details are different and, often, incompatible?

Existing test beds store experiment data in log files in specific formats. In most of the cases, tools used in test beds (e.g., fault injection, robustness testing, benchmarks, security tools, etc) either provide rudimentary means to analyze data or, more frequently, just store the raw results in a spreadsheet format such as the Microsoft Excel®. Although this approach can be acceptable for specific analysis, it is clearly not enough when the amount of raw data is very large, the analysis required is complex, and particularly when heterogeneous data from different test beds has to be analyzed and cross-exploited.

### Proposed Solution: The Data Warehousing Approach

The proposed approach is to collect raw data (sanitized if needed) produced in different test beds and stores it in a multidimensional data structure (data warehouse). The data analysis is done through the use of widely available OLAP (On-Line Analytical Processing) tools and data mining tools such as the ones traditionally used in business decision support analysis. Existing test beds and associated tools are used as they are, and the only thing required is to export the data obtained in the experiments to the data warehouse, where all the analysis and cross-exploitation of results can be done in an efficient and general way.

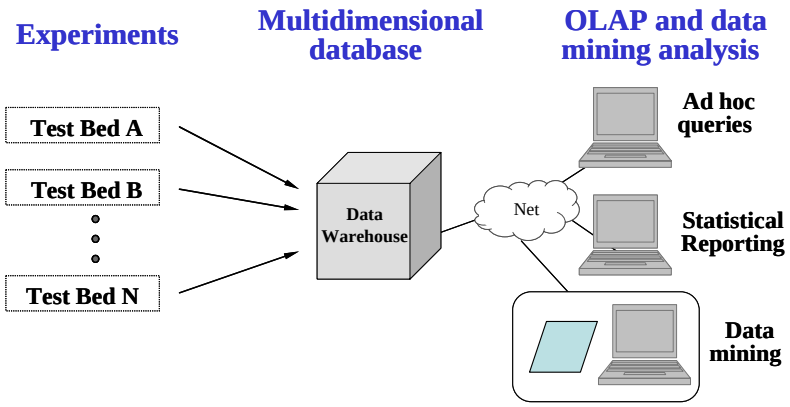

The multidimensional data structure (data warehouse) works as a common format to share and cross-exploitation data from different test beds. In fact, experimental evaluation of dependability and security is a multidimensional problem (actually, all experiments are multidimensional). The readouts collected in test beds during an experiment (called facts in the multidimensional jargon) are analyzed according to groups of experiment features called dimensions. For example, raw data representing things such as error detection efficiency or error recovery time are facts, while the dimensions represent the different target systems, configurations, workloads, faultloads, etc.

The OLAP tools allow easy computation of measures from the facts, filtered according to the attributes of the dimensions. By using different filters over dimensions attributes, the user can isolate the effect on the measures of specific aspects (dimension) such as different workloads, target systems, faultloads, etc. Furthermore, it is easy to accommodate sanitized data in a multidimensional data structure as the sanitization mainly affects the data stored in the dimensions (it is worth mentioning that the big bulk of data in a multidimensional model is in the facts, not in the dimensions).

In addition to using OLAP tools to analyze and cross-exploitation data from different test beds on demand, it is also possible to use the data warehouse to extract data into flat tables used by data mining tools. Data mining allows for the extraction of information (and knowledge) form datasets that generally cannot be achieved by user driven analysis, helping in the identification of key factors for the experiments done in each test bed. The main advantages of the proposed approach are the following:

- It is a general and ready to use approach:
  - As experiments are multidimensional problems, the use of multidimensional database and OLAP technologies for result analysis and sharing is applicable to all the experimental dependability and security evaluation scenarios.
  - Existing test beds can be used as they are; the only thing we need to know is the data format of the raw results produced during experiments, in order to read them into the data warehouse.
  - Data warehousing, OLAP, and data mining technologies are mature and readily available, and have proved (in the business decision support area) to be very effective in the integration of heterogeneous data from different systems in a common and sharable repository.
- It is easy to compare and cross-exploit raw results from different experiments, as all the raw data is stored in a common data warehouse (or in distributed data warehouses).
- It is easy to share raw results worldwide as the data stored in a data warehouse can be explored by web-enabled versions of OLAP tools. This way, it is possible to make available to the entire dependability community the raw data of dependability evaluation experiments (data is available for analysis through a web page). Furthermore, as the data warehouse stores raw data, it means that it is always possible to repeat the analysis or to perform a different analysis over the same raw data.

**Suggested Project**

Questions:

Are we using available test beds in the best way possible?

Are we extracting all the results and information form the test beds available?

Are we sharing the (raw) data we have got from our test beds?

Are we cross-exploiting results from different test beds?

Are we going to change things in future test beds?

Then, it is time to improve the way we share and cross-exploit available test beds and build a new kind of test bed: data warehouses that connect available test beds by the most important thing: data.

# 5 Appendix B - Workshop Agenda

## EU-US Summit Series on Cyber Trust: System Dependability & Security

### Proposed Working Schedule

*14th November*

*7:30pm* Registration

*8:30pm* Wine reception and buffet

*Finnstown Country House Hotel* http://www.finnstown-hotel.ie/

First Day: November 15th 2006

| | |
|---|---|
| **08:30– 08:40** | Welcome, Workshop objectives and format<br>*(EU representative - Willie Donnelly*<br>*US representative - Bill Sanders)* |
| **08:40 – 09:00**<br>**09:00 – 09:20**<br>**09:20 - 09:40**<br>**09:40 – 09:55** | Setting the scene for the Workshop<br>*Representative from the European Commission (Thomas Skordas)*<br>*Representative from NSF (Deborah Crawford)*<br>*Representative from DHS (Doug Maughan)*<br>*Questions/Discussion* |
| **09:55 – 10:00** | Coffee break |
| **10:00 – 12:00**<br><br>**Each presenter is limited to 10 minutes to facilitate open discussion** | **Panel A: Dependability & Security of Future Networked Systems – architecture and design issues**<br>Co-Chairs:<br>David Du (NSF, US) & Paulo Verissimo (Univ. Lisboa, Portugal)<br>Rapporteur: *Jim Clarke*<br><br>*Speakers:*<br>✓ Ravishankar Iyer (Univ. Of Illinois, US)<br>✓ Michel Riguidel (ENST, France)<br>✓ Felix Wu (UCDavis, US)<br>✓ Bart Preneel (Katholieke Univertsiteit Leuven, Belgium)<br>✓ Yair Amir (JHU, US)<br>✓ Neeraj Suri (TU Darmstadt, Germany) |

| | |
|---|---|
| **12:00-13:00** | **<u>Panel B:</u> Dependability & Security of Future Networked Systems – scalability and context-awareness**<br>Co-Chairs*:*<br>John Knight (Univ. Virginia, US) & Brian Randell (Univ. Newcastle upon Tyne, UK)<br>Rapporteur: *Ed Dawson*<br><br>*Speakers:*<br>✓ Jean-Claude Laprie (LAAS, France)<br>✓ Nick Weaver (Berkeley, US)<br>✓ Christof Fetzer (TU Dresden, Germany)<br>✓ George Kesidis (Penn. State, US)<br>✓ Gerard LeLann (INRIA, France)<br>✓ Ming-Yuh Huang (Boeing, US) |
| **13:00 - 13:45** | Lunch |
| **13:45 - 14:45** | **Continuation of Panel B (Discussions)** |
| **14:45– 15:00** | Coffee break |
| **15:00– 17:00** | **<u>Panel C:</u> - Security & Privacy in Dynamic Wireless Networks**<br>Co-Chairs:<br>Gene Tsudik (UC Irvine, US) & Roberto Baldoni, (U. Roma, Italy)<br>Rapporteur: *Michael Bailey*<br><br>*Speakers:*<br>✓ David Kotz (Dartmouth, US)<br>✓ Reijo Savola (VTT, Finland)<br>✓ Joe Evans (KU,US)<br>✓ Stephan Engberg (Priway, Denmark)<br>✓ Wenke Lee (Georgia Inst. Of Technology, Atlanta, US)<br>✓ Paddy Nixon (UCD, Ireland) |
| **17:00 – 17:20**<br><br>**17:20 - 17:40**<br><br>**17:40 – 17:50** | Australian perspective in securing future communication networks *(Ed Dawson)*<br><br>Japanese perspective in future networked dependable systems (*Takashi Nanya*)<br><br>*Questions/Discussion* |
| **17:50 – 18:00** | Concluding remarks on day 1 *(General Chairs of the event)*<br>Review of agenda for day 2 |

*7pm Meet in reception/ bar area for pre-dinner drink*

*7:30pm Conference Dinner Finnstown Country House Hotel*

Second Day: November 16th 2006

| | |
|---|---|
| **08:30 – 08:35** | Agenda and WS Format for today *(US representative - Karl Levitt)* |
| **08:35 – 10:15** | **Panel D: Evaluating the Dependability & Security of Networked Systems – modelling, simulation, predictive evaluation, assurance cases**<br>Co-Chairs:<br>William H. Sanders (Univ. of Illinois, US) & Dieter Gollmann (TU Hamburh-Harburg, Germany)<br>Rapporteur: *Stephan Engberg*<br>*Speakers:*<br>✓ John Rushby (SRI International, California, US)<br>✓ Bev Littlewood (City Unversity, UK)<br>✓ John McHugh (DAL,Canada)<br>✓ Aad van Moorsel (Newcastle, UK)<br>✓ O. Sami Saydjari (Cyber Defense Agency, LLC, US)<br>✓ Robin Bloomfield (City Univ + Adelard, UK) |
| **10:30 - 12:15** | **Panel E: Evaluating the Dependability & Security of Networked Systems – monitoring, operational assessment, auditing**<br>Co-Chairs:<br>David M. Nicol (Univ of Illinois, US) & Marcelo Masera (JRC, Italy)<br>Rapporteur: *Jim Just*<br>*Speakers:*<br>✓ Roy Maxion(CMU. US)<br>✓ Evangelos Markatos, (FORTH-Creta, Greece)<br>✓ Alfonso Valdes (SRI, US)<br>✓ Fabio Martinelli (IIT-CNR, Italy)<br>✓ Todd Heberlein (NetSQ, US)<br>✓ András Pataricza (Univ. of Budapest, Hungary) |
| **12:15 - 13:15** | Lunch |
| **13:15 - 15:15** | **Panel F: Future Test beds**<br>Co-Chairs:<br>Doug Maughan (DHS, US) & Jim Clarke (Waterford Institute of Technology, Ireland)<br>Rapporteur: *Zeta Dooly*<br>*Speakers:*<br>✓ Mike Bailey (Univ. of Michigan, US)<br>✓ Pekka Nikander (Ericsson NomadicLab, Finland)<br>✓ Anthony Joseph (Berkeley, US)<br>✓ Ollie Whitehouse (Symantec, UK)<br>✓ Jim Just (Global Info Tek, US)<br>✓ Henrique Madeira (Univ. of Coimbra, Portugal) |
| **15:30– 17:00** | Open discussion session with input from all working panels presented by US and EU Chairs. |
| **17:00 -17:15** | Concluding remarks *(US and EU representatives from funding Institutions)* |

# 6 Appendix C - Workshop Participant List

Michel Riguidel, ENST, France
Bart Preneel, Katholieke Univertsiteit Leuven, Belgium
Neeraj Suri, TU Darmstadt,, Germany
Paulo Verissimo , Univ. Lisboa, Portugal
Jean-Claude Laprie , LAAS, France
Christof Fetzer, TU Dresden, Germany
Gerard LeLann, Director of Research at INRIA ,
Brian Randell, Newcastle University, UK
Reijo Savola, VTT, Finland
Paddy Nixon, UCD, Ireland
Stephan Engberg, Priway, Denmark
Roberto Baldoni, U. Roma, Italy
Bev Littlewood, City Univ , UK
Robin Bloomfield, City Univ + Adelard, UK
Aad van Moorsel, Newcastle Uni, UK
Dieter Gollmann, TU Hamburg-Harburg, Germany
Evangelos Markatos, FORTH-Creta, Greece
Fabio Martinelli, IIT-CNR, Italy
András Pataricza, Univ. of Budapest, Hungary
Marcelo Masera, Joint Research Centre, EC, Italy
Henrique Madeira , Univ. of Coimbra, Portugal
Jim Clarke, Waterford Institute of Technology, Ireland
Ed Dawson, Queensland University, Australia
Sandro Etalle, Embedded sys area,
Yassine Lakhnech, Embedded sys area,
Wide Hogenhout, European Commission,
Jacques Bus, European Commission, Brussels
Thomas Skordas, European Commission, Brussels
Willie Donnelly, Waterford Institute of Technology, Ireland
Zeta Dooly, Waterford Institute of Technology, Ireland
Takashi Nanya, University of Tokyo, Tokyo, Japan
Ravishankar Iyer, Univ. of Illinois, Urbana-Champaign, US
Yair Amir, Johns Hopkins University, US
Felix Wu, Univ. of California, Davis, US
George Kesidis, Penn. State Univ., US
Nick Weaver, Univ. of California, Berkeley, US
Ming-Yuh Huang, Boeing Corporation, US
David Kotz, Dartmouth, US
Wenke Lee, Georgia Inst of Technology, US
Joe Evans, Univ. of Kansas, US
William Sanders, Univ. of Illinois, Urbana-Champaign, US
John Rushby, SRI Intl Computer Science Lab, US
John McHugh, Dalhousie Univ., Canada
O. Sami Saydjari, Cyber Defense Agency, US
David M. Nicol, Univ. of Illinois, Urbana-Champaign, US
Roy Maxion, Carnegie Mellon Univ., US
Alfonso Valdes, SRI Intl Computer Science Lab, US Todd Heberlein, Net Squared, Inc., US
Mike Bailey, Univ. of Michigan, US
Anthony Joseph, Univ. of California, Berkeley, US
Jim Just, Global InfoTek, Inc, US
Doug Maughan, Dept of Homeland Security, US
Gene Tsudik,Univ. of California, Irvine, US
John Knight,Univ. of Virginia, US
David Du, National Science Foundation, US
Karl Levitt, National Science Foundation, US