

CRUTIAL: The Blueprint of a Reference Critical Information Infrastructure Architecture ^{*}

Paulo Veríssimo, Nuno Ferreira Neves, and Miguel Correia

University of Lisboa, Faculty of Sciences
Bloco C6, Campo Grande, 1749-016 Lisboa - Portugal
{pjuv,nuno,mpc}@di.fc.ul.pt <http://www.navigators.di.fc.ul.pt>

Abstract. The largely computerised nature of critical infrastructures and the pervasive interconnection of systems all over the world have generated the problem of achieving resilience of critical information infrastructures, namely against computer-borne attacks and severe faults. We believe the complexity of the problem prevents it from being solved using classical security methods.

The paper focuses on the computer systems behind electrical utility infrastructures, and proposes the blueprint of a distributed systems architecture that we believe may come to be useful as a reference for modern critical information infrastructures in general. The architecture is populated with a set of classes of techniques and algorithms, based on paradigms providing resilience to faults and attacks in an automatic way.

1 Introduction

The largely computerised nature of critical infrastructures on the one hand, and the pervasive interconnection of systems all over the world, on the other hand, have generated one of the most fascinating current problems of computer science and control engineering: *how to achieve resilience of critical information infrastructures*.

This problem is concerned with ensuring acceptable levels of service and in last resort, the integrity of systems themselves, when faced with threats of several kinds. In this paper we are concerned with threats centered in computers and control computers. These threats range from accidental events like natural faults or wrong manoeuvres [8], to attacks by hackers or terrorists [3, 7]. The problem affects systems with great socio-economic value, such as utility systems like electrical, gas or water, or telecommunication systems and computer networks like the Internet. In consequence, the high degree of interconnection is causing great concern, given the level of exposure of very high value systems and components to attacks that can be perpetrated in an anonymous and remote way.

Although we see an increase in the concern for using best practices of security in these systems, we believe that the problem is not completely understood, and

^{*} This work was mainly supported by the EC, through project IST-4-027513-STP (CRUTIAL), and also by the FCT, through LASIGE and projects POSI/EIA/61643/2004 (AJECT) and POSI/EIA/60334/2004 (RITAS) .

can not be solved with classical methods. Its complexity is mainly due to the hybrid composition of those infrastructures:

- The operational network, called generically SCADA (Supervisory Control and Data Acquisition)¹, composed of the computer systems that yield the operational ability to supervise, acquire data from, and control the physical processes. In fact, to the global computer system, SCADA computer systems (e.g., controllers) “are” the controlled processes (e.g., power generators), since by acting on the former, for example, through a network message, one changes the state of the latter.
- The standard corporate intranet, where usual departmental services and clients reside, and also the engineering and technical staff, who access the SCADA part through ad-hoc interconnections.
- The Internet, through which intranet users get to other intranets and/or the outside world, but to which, and often unwittingly, the SCADA network is sometimes connected to.

On the other hand, this mixture has given an unexpected inter-disciplinary nature to the problem: SCADA systems are real-time systems, with some reliability or even fault-tolerance concern, but they were classically not designed to be widely distributed or remotely accessed, let alone open to other more asynchronous and less trusted subsystems. Likewise, they were not designed with security in mind. In consequence, in scientific terms, our problem can be formulated as follows:

- The computer-related operation of a critical utility infrastructure is a distributed systems problem including interconnected SCADA/embedded networks, corporate intranets, and Internet/PSTN² access subsystems.
- This distributed systems problem is hard, for it simultaneously includes facets of real-time, fault tolerance, and security.

In our paper, we focus on the computer systems behind electrical utility infrastructures as an example, and we propose the blueprint of: a distributed systems architecture that we believe may come to be useful as a reference for modern critical information infrastructures; a set of classes of techniques and algorithms based on paradigms providing resilience to faults and attacks in an automatic way. This work is ongoing and is done in the context of the recently started CRUTIAL European project, CRITICAL UTILITY InfrastructurAL resilience [4], details of which are given in the end.

2 Rationale for the Model and Architecture

Before we proceed, let us bring some further insight on the problem of critical infrastructures:

¹ Or PCS (Process Control System).

² Public Switched Telephone Network

- Critical Information Infrastructures (CII) feature a lot of legacy subsystems and non-computer-standard components (controllers, sensors, actuators, etc.).
- Conventional security and protection techniques, when directly applied to CII controlling devices, sometimes stand in the way of their effective operation.

These two facts will not change at least for a long time, so they should be considered as additional research challenges. Despite security and dependability concerns with those individual components being a necessity, we believe that the crucial problem is with the forest, not the trees. That is, the problem of critical information infrastructure insecurity is mostly created by the informatics nature of many current infrastructures, and by the generic and non-structured network interconnection of CIIs, which bring several facets of exposure, from internal unprotected wireline or wireless links, to interconnections of SCADA and corporate intranets to the Internet and PSTN.

The problems that may result from this exposure to computer-borne threats range from wrong manoeuvring, to malicious actions coming from terminals located outside, somewhere in the Internet. The potential targets of these actions are computer control units, embedded components and systems, that is, devices connected to operational hardware (e.g., water pumps and filters, electrical power generators and power protections, dam gates, etc.) or to telecom hardware (core routers, base stations, etc.). The failure perspectives go from unavailability of services supposed to operate 24×7, to physical damage to infrastructures.

Whilst it seems non-controversial that such a status quo brings a certain level of threat, we know of no work that has tried to equate the problem by defining a reference model of “critical information infrastructure distributed systems architecture” providing the necessary global resilience against abnormal situations.

What can be done at architectural level to achieve resilient operation? Note that the crux of the problem lies with the fact that access to operational networks, such as remote SCADA manoeuvring, ended up entangled with access to corporate intranets and to public Internet, without there being computational and resilience models that *represent* this situation, unlike what exists in simpler, more homogeneous settings, e.g. classical web-based server infrastructures on Internet. Our point is that interference and threats start at the level of the macroscopic information flows between these subsystems, and can in consequence be stopped there³.

Now, given the simultaneous need for real-time, security and fault-tolerance, this problem is hard vis-a-vis existing paradigms. For example, many classical distributed systems paradigms handle each of those facets separately, and just solve part of the problem. A unifying approach has gained impressive momentum currently: *intrusion tolerance*. In short, instead of trying to prevent every single intrusion or fault, these are allowed, but tolerated: systems remain to some extent faulty and/or vulnerable, attacks on components can happen and some will be

³ But this should not prevent the study of techniques at the controller level.

successful, but the system has the means to trigger automatic mechanisms that prevent faults or intrusions from generating a system failure.

Our approach is thus equated along the following propositions:

PROPOSITION 1: Classical security and/or safety techniques alone will not solve the problem: they are largely based on prevention and ultimately disconnection.

There is a recent and positive trend to make SCADA systems and CIIs at large more secure [1, 7, 11]. However, classic engineering remedies place real-time and embedded (RTE) systems at most at the current level of commercial systems' security and dependability, which is known to be insufficient [3, 5, 12]: systems constantly suffer attacks, intrusions, some of them massive (worms); most defences are dedicated to generic, non targeted attacks; they degrade business, but only do virtual damage, unlike RTE systems, where there is a risk of physical damage. On the other hand, some current IT security techniques can negatively affect RTE system operation, w.r.t. availability and timeliness. For example, if the former are based on disconnection, significant performance degradation, or even defensive restrictions to access control.

PROPOSITION 2: Any solution, to be effective, has to involve automatic control of macroscopic command and information flows, occurring essentially between the physical or virtual LANs⁴ composing the critical information infrastructure architecture, by securing appropriate system-level properties.

We believe that a key to the solution lies with controlling the command and information flow at macroscopic level (organisation-level). We are talking about an architectural model, a set of architectural devices, and key algorithms, capable of achieving the above-mentioned control of the command and information flow. The devices and algorithms should be capable of securing a set of system-level properties characterising whatever is meant by correct and resilient behaviour.

PROPOSITION 3: We lack a reference architecture of “modern critical information infrastructure” considering different interconnection realms and different kinds of risk, throughout the physical and the information subsystems of a CII.

We must consider the physical or virtual LANs composing the operational SCADA/embedded networks, the corporate intranets, and the Internet/PSTN access networks, as different first order citizens of the architecture. Likewise, the notion that risk factors may vary and be difficult to perceive accurately, brings the need to reconcile uncertainty with predictability in architecture and algorithmics.

⁴ Local Area Networks.

3 CRUTIAL Architecture

The CRUTIAL architecture encompasses:

- Architectural configurations featuring trusted components in key places, which a priori induce prevention of some faults, and of certain attack and vulnerability combinations.
- Middleware devices that achieve runtime automatic tolerance of remaining faults and intrusions, supplying trusted services out of non-trustworthy components.
- Trustworthiness monitoring mechanisms detecting situations not predicted and/or beyond assumptions made, and adaptation mechanisms to survive those situations.
- Organisation-level security policies and access control models capable of securing information flows with different criticality within/in/out of a CII.




Fig. 1. CRUTIAL overall architecture (WAN of LANs connected by CIS, P processes live in the several nodes)

We build on results from the MAFTIA project⁵ in this field [15], but extend them significantly to attend the specific challenges of the critical information infrastructure problem, for example, timeliness, global access control, and above all non-stop operation and resilience.

Given the severity of threats expected, some key components are built using architectural hybridisation methods in order to achieve *trusted-trustworthy*

⁵ Malicious-and Accidental-Fault Tolerance for Internet Applications. The web site of the project is at www.maftia.org.

operation [15]: an architectural paradigm whereby components prevent the occurrence of some failure modes *by construction*, so that they can justifiably be trusted on their resistance to faults and hackers.

Intrusion tolerance mechanisms are selectively used in the CRUTIAL architecture, to build layers of progressively more trusted components and middleware subsystems, from baseline untrusted components (nodes, networks) [15]. This leads to an automation of the process of building trust: for example, at lower layers, basic intrusion tolerance mechanisms are used to construct a trustworthy communication subsystem, which can then be trusted by upper layers to securely communicate amongst participants without bothering about network intrusion threats.

One of the innovative aspects of this work, further to intrusion tolerance, is the resilience aspect, approached through two paradigms: *proactive-resilience* to achieve exhaustion-safety [9], to ensure perpetual, non-stop operation despite the continuous production of faults and intrusions; and *trustworthiness monitoring* to perform surveillance of the coverage stability of the system, that is, of whether it is still performing inside the assumed fault envelope or beyond assumptions made [2]. In the latter case, dependable adaptation mechanisms are triggered.

Finally, the desired control of the information flows is partly performed through protection mechanisms using an adaptation of *organisation-based access control* models [6] for implementing global-level security policies.

The mechanisms and algorithms in place achieve system-level properties of the following classes: trustworthiness, read as security and dependability, or resistance to faults and intrusions; timeliness, in the sense of meeting timing constraints raised by real world control and supervision; coverage stability, to ensure that variation or degradation of assumptions remains within a bounded envelope; dependable adaptability, to achieve predictability in uncertain conditions; resilience, read as correctness and continuity of service even beyond assumptions made.

3.1 Main architectural options

We view the system as a WAN-of-LANs, as introduced in [13]: there is a global interconnection network, the WAN, and packets are switched by it, through generic devices that we call *facility gateways*, which are the representative gateways of each LAN (the overall picture is shown in Figure 1). The WAN is a logical entity operated by the CII operator companies, which may or not use parts of public network as physical support. A LAN is a logical unit that may or not have physical reality (e.g., LAN segments vs. VLANs⁶). More than one LAN can be connected by the same facility gateway. All traffic originates from and goes to a LAN. As example LANs, the reader can envision: the administrative clients and the servers LANs; the operational (SCADA) clients and servers LANs; the engineering clients and servers LANs; the PSTN modem access LANs; the Internet and extranet access LANs, etc.

⁶ Virtual LANs

The facility gateways of a CRUTIAL critical information infrastructure are more than mere TCP/IP routers, collectively they act as a set of servers providing distributed services relevant to solving our problem: *achieving control of the command and information flow, and securing a set of necessary system-level properties*. CRUTIAL facility gateways are called *CRUTIAL Information Switches (CIS)*, which in a simplistic way could be seen as sophisticated circuit or application level firewalls combined with equally sophisticated intrusion detectors, connected by distributed protocols.

This set of servers must be intrusion-tolerant, prevent resource exhaustion providing perpetual operation, and be resilient against assumption coverage uncertainty, providing survivability. The services implemented on the servers must also secure the desired properties of flow control, in the presence of malicious traffic and commands, and in consequence be themselves intrusion-tolerant.

An assumed number of CIS can be corrupted. Therefore, a logical CIS may be implemented as a set of replicated physical units (CIS replicas) according to fault and intrusion tolerance needs. Likewise, CIS are interconnected with intrusion-tolerant protocols, in order to cooperate to implement the desired services.

3.2 CRUTIAL nodes

The structure of a CRUTIAL node or host closely follows the node structuring principles for intrusion-tolerant systems explained in [15]:

- The notion of *trusted* – versus *untrusted* – *hardware*. For example, most of the hardware of a CIS is considered to be untrusted, with small parts of it being considered trusted-trustworthy.
- The notion of *trusted support software*, trusted to execute a few critical functions correctly, the rest being subjected to malicious faults.
- The notion of *run-time environment*, offering trusted and untrusted software and operating system services in a homogeneous way.
- The notion of *trusted distributed components*, for example software functions implemented by collections of interacting CIS middleware.

A snapshot of the CRUTIAL architecture detailing nodes and their interconnection methods is depicted in three dimensions in Figure 2, where we can perceive the above-mentioned node structuring principles. In the context of this paper, we will focus on the CRUTIAL Information Switch (CIS) nodes. However, other specific nodes, for example, controllers needing to meet high trustworthiness standards, may be also built to a similar structure.

Firstly, there is the *hardware* dimension, which includes the node and networking devices that make up the physical distributed system. We assume that most of a node's operations run on untrusted hardware, e.g., the usual machinery of a computer, connected through the normal networking infrastructure, which we call the *payload channel*. However, some nodes– CIS, for example– may have pieces of hardware that are trusted, for example, that by construction intruders do not have direct access to the inside of those components. The type of trusted




Fig. 2. CRUTIAL node architecture

hardware featured in CIS is an *appliance board with processor*, which may or not have an *adapter to a control channel* (an alternative trusted network), as depicted in Figure 2. This appliance is plugged to the CIS's main hardware.

Secondly, services based on the trusted hardware are accessed through the *local support* services. The rationale behind our trusted components is the following: whilst we let a local node be compromised, we make sure that the trusted component operation is not undermined (crash failure assumption).

Thirdly, there is the *distributed software* provided by CRUTIAL: middleware layers on top of which distributed applications run, even in the presence of malicious faults (far right in Figure 2). In the context of this paper, we will discuss the layers of *middleware* running inside a CIS.

4 CRUTIAL Middleware

We now observe the part of the system made of the WAN and all the CIS (facility gateways) that interconnect all the internal LANs of the critical information infrastructure to the WAN.

We model this as a distributed system with N nodes (CIS). We use the weakest fault and synchrony models that allow to carry out the application tasks. So, we use the asynchronous/arbitrary model as a starting point, and strengthen it as needed. For example, by resorting to hybrid models [14] using wormholes, and assuming some form of partial synchrony.

We assume that the environment formed by the WAN and all the CIS is hostile (not trusted), and can thus be subjected to malicious (or arbitrary, or Byzantine) faults. On the other hand, LANs trust the services provided by the CIS, but are not necessarily trusted by the latter. That is, as we will see below, LANs have different degrees of trustworthiness, which the CIS distributed protocols have to take into account. CIS securely switch information flows as a service to edge LANs as clients.

We assume that faults (accidental, attacks, intrusions) continuously occur during the life-time of the system, and that a maximum number of f malicious

(or arbitrary, or Byzantine) faults can occur within a given interval. We assume that services running in the nodes (CIS) cooperate through distributed protocols in such an environment. In consequence, some of these nodes may be replicated for fault/intrusion tolerance.

Some of the services running in CIS may require some degree of timeliness, given that SCADA implies synchrony, and this is a hard problem with malicious faults. We also take into account that these systems should operate non-stop, a hard problem with resource exhaustion (the continued production of faults during the life-time of a perpetual execution system leads to the inevitable exhaustion of the quorum of nodes needed for correct operation [9]).

4.1 LAN-level services

A LAN is the top-level unit of the granularity of access control, regardless of possible finer controls. It is also and correspondingly, a unit of trust or mistrust thereof. In fact, we are not concerned with what happens inside a LAN, except that we may attribute it a different levels of trust.

Traffic (packets) originating from a LAN receive a label that reflects this level of trust, and contains access control information, amongst other useful things.

The trustworthiness of a label (that is, the degree in which it can or not be tampered with) can vary, depending on the criticality of the service. In the context of this paper, and without loss of generality, we assume it is an authenticated proof of a capacity.

4.2 WAN-level services

The collection of CIS implements a set of core services, aiming at achieving the objectives we placed as desirable for a reference model of *critical information infrastructure distributed systems architecture*:

- Byzantine-resilient information and command dissemination between CIS units, with authentication and cryptographic protection (broadcast, multicast, unicast).
- Pattern-sensitive information and command traffic analysis (behaviour and/or knowledge based intrusion detection) with Byzantine-resilient synchronisation and coordination between local Intrusion Detection Systems (IDSs).
- CIS egress/ingress access control based on LAN packet labels and/or additional mechanisms, implementing an instance of the global security policy.

The CIS middleware layers implement functionality at different levels of abstraction, as represented in Figure 3. As mentioned earlier, a middleware layer may overcome (through intrusion tolerance) the fault severity of lower layers and provide certain functions in a trustworthy way.

The lowest layer is the *Multipoint Network* module (MN), created over the physical infrastructure. Its main properties are the provision of multipoint addressing, basic secure channels, and management communications, hiding the

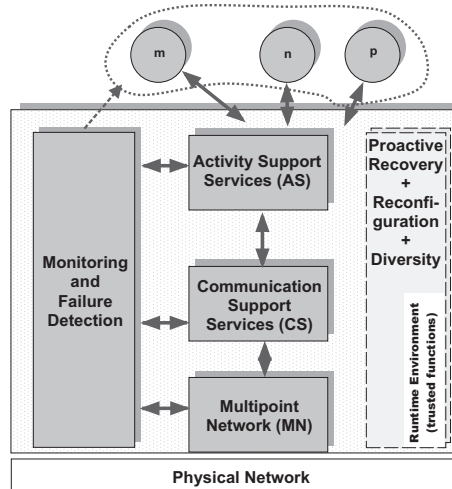


Fig. 3. CIS middleware

specificities of the underlying network. The *Communication Support Services* module (CS) implements basic cryptographic primitives, Byzantine agreement, consensus, group communication and other core services. The CS module depends on the MN module to access the network. The *Activity Support Services* module (AS) implements building blocks that assist participant activity, such as replication management (e.g., state machine, voting), IDS and firewall support functions, access control. It depends on the services provided by the CS module.

The block on the left of the figure generically implements Monitoring and Failure Detection. Failure detection assesses the connectivity and correctness of remote nodes, and the liveness of local processes. Trustworthiness monitoring and dependable adaptation mechanisms also reside in this module, and have interactions with all the modules on the right. Both the AS and CS modules depend on this information.

The block to the right represents the support services. Remember that these include the vanilla operating system's services, but also the trusted services supplied in support to the algorithms in the various modules. These are described and include: proactive recovery, reconfiguration, and diversity management.

5 CRUTIAL Information Switches (CIS)

Let us briefly discuss how CIS are made trusted-trustworthy components. CIS are built with a combination of untrusted and trusted hardware of varying degrees, depending on the needs and criticality of the traffic (sink or source) and the services they support. CIS individual resilience is enhanced by proactive resilience mechanisms, using a construct called Proactive Resilience Wormhole, described elsewhere [10], aiming at providing for perpetual execution of a given

set of CIS, despite continued intrusion and/or failure of an assumed simultaneous maximum number of CIS at an assumed maximum rate.

These notions can be recursively used to construct a logical CIS which is in fact a set of replicated physical CIS units, running some internal intrusion-tolerant protocols so that the whole appears to the protocol users as a single logical entity sinking/sourcing to/from a given LAN, but is in fact resilient to attacks on the CIS themselves. This is a powerful combination since the resilience of protocols running on such intrusion-tolerant CIS components is commensurate to arbitrary-failure counterparts.

CIS are in addition provided with trustworthiness monitoring subsystems, aiming at assessing the trustworthiness of the CIS itself: as a function of the evolution of the coverage of the assumptions underlying the whole FIT (fault and intrusion tolerant) design. As such, trustworthiness becomes a dynamic property, which provides further resilience to the CIS, through dependable adaptation: automatically reacting to environment uncertainty (changing fault and/or attack levels) and maintaining coverage stability, by changing operation parameters or modes automatically. Finally, for very high levels of resilience, CIS construction and or reconfiguration in the course of proactive recovery may be based on diversity techniques (ex. n-version programming, obfuscation, etc.).

6 Conclusion

The paper presents a blueprint of a distributed systems architecture for resilient critical information infrastructures, with respect to both accidental faults and malicious attacks and intrusions. The rationale for this work was based on three fundamental propositions: classical security and/or safety techniques alone will not be enough to solve the problem; any effective solution has to involve automatic control of macroscopic command and information flows between the LANs composing the CII; and, the unifying paradigm should be a reference architecture of “resilient critical information infrastructures” performing the integration of the different realms of a CII system.

The proposed solution encompasses a range of mechanisms of incremental effectiveness, to address from the lower to the highest criticality operations in a CII. Architectural configurations with trusted components in key places induce prevention of some attacks. Middleware software attains automatic tolerance of the remaining faults and intrusions. Trustworthiness enforcing and monitoring mechanisms allow unforeseen adaptation to extremely critical, not predicted situations, beyond the initial assumptions made.

Functionally, the information flow is controlled by basic mechanisms of the firewall and intrusion detection type, complemented and parameterised by organisation-level security policies and access control models, capable of securing information flows with different criticality within a CII and in/out of it.

Acknowledgements CRUTIAL is a project of the IST programme of the European Commission. Several institutions participate to the project: CESI RICERCA (Italy),

FCUL (Portugal), CNR-ISTI (Italy), LAAS-CNRS (France), K.U.Leuven-ELECTA (Belgium), CNIT (Italy). Details about the project and all the participants can be found in the CRUTIAL portal: <http://crutial.cesiricerca.it/>.

References

1. President's Critical Infrastructure Protection Board and Office of Energy Assurance U.S. Department of Energy. *21 Steps to Improve Cyber Security of SCADA Networks*. U.S. Department of Energy, 2002.
2. A. Bondavalli, S. Chiaradonna, D. Cotroneo, and L. Romano. Effective fault treatment for improving the dependability of COTS and legacy-based applications. *IEEE Transactions on Parallel and Distributed Systems*, 1(4):223–237, 2004.
3. J. Cieslewicz. Attacks and accidents: Policy to protect the power grid's critical computing and communication needs. Senior interdisciplinary honors thesis in international security studies, Stanford University, May 2004.
4. G. Dondossola, G. Deconinck, F. Di Giandomenico, S. Donatelli, M. Kaaniche, and P. Verissimo. Critical utility infrastructural resilience. In *International Workshop on Complex Network and Infrastructure Protection*, March 2006.
5. L. A. Gordon, M. P. Loeb, W. Lucyshyn, and R. Richardson. 2006 CSI/FBI computer crime and security survey. Computer Security Institute, 2006.
6. A. A. El Kalam, R. Elbaida, P. Balbiani, S. Benferhat, F. Cuppens, Y. Deswarte, A. Mige, C. Saurel, and G. Trouessin. Organization-based access control. In *IEEE 4th International Workshop on Policies for Distributed Systems and Networks*, pages 277–288, June 2003.
7. H. Li, G. W. Rosenwald, J. Jung, and C. Liu. Strategic power infrastructure defense. *Proceedings of the IEEE*, 93(5):918–933, May 2005.
8. V. Madani and D. Novosel. Getting a grip on the grid. *IEEE Spectrum*, 42(12):42–47, December 2005.
9. P. Sousa, N. F. Neves, and P. Verissimo. How resilient are distributed fault/intrusion-tolerant systems? In *Proceedings of the IEEE International Conference on Dependable Systems and Networks*, June 2005.
10. P. Sousa, N. F. Neves, and P. Verissimo. Resilient state machine replication. In *Proceedings of the 11th Pacific Rim International Symposium on Dependable Computing*, pages 305–309, December 2005.
11. J. Stamp, J. Dillinger, W. Young, and J. DePoy. Common vulnerabilities in critical infrastructure control systems. Technical report, Sandia National Laboratories, May 2003.
12. D. Turner, S. Entwisle, O. Friedrichs, D. Ahmad, J. Blackbird, M. Fossi, D. Hanson, S. Gordon, D. Cole, D. Cowlings, D. Morss, B. Bradley, P. Szor, E. Chien, J. Ward, J. Gough, and J. Talbot. Symantec Internet security threat report. Trends for January 05–June 05. Symantec, Volume VIII, September 2005.
13. P. Verissimo. Lessons learned with NavTech: a framework for reliable large-scale applications. DI/FCUL TR 02–17, Department of Informatics, University of Lisbon, December 2002.
14. P. Verissimo. Travelling through wormholes: a new look at distributed systems models. *SIGACTN: SIGACT News (ACM Special Interest Group on Automata and Computability Theory)*, 37(1):66–81, 2006.
15. P. Verissimo, N. F. Neves, C. Cachin, J. Poritz, D. Powell, Y. Deswarte, R. Stroud, and I. Welch. Intrusion-tolerant middleware: The road to automatic security. *IEEE Security & Privacy*, 4(4):54–62, Jul./Aug. 2006.