Thou Shalt Not Trust non-Trustworthy Systems

Paulo Esteves Veríssimo
Univ. of Lisboa, Faculty of Sciences – Portugal
http://www.di.fc.ul.pt/~pjv

E-mail: pjv@di.fc.ul.pt

1. Introduction

Computer systems and ICT at large (information and communication technologies) are on the verge of a strange era: on the one hand, everyday we require more from applications as seen by users (response, determinism, robustness, security); on the other hand, improvements in infrastructure technology peer with asymmetry and instability (access networks, mobility, de-regulation, and so forth). This evolution of distributed computing and applications has put new challenges on models, architectures and systems. In essence, we should look for paradigms that help us reconcile uncertainty with predictability.

Grand challenges require drastic changes, and they are happening: in the hybrid, dynamic and decentralised way we start looking at system design, once quite homogeneous, static, centralised, and in the cross-fertilising way we now look at previously disjoint scientific fields. Two issues are central to modern design of dependable and secure dynamic distributed systems:

- the confluence between classical dependability and security, met essentially but not only by the concept of common 'accidental fault and malicious intrusion tolerance'.
- and the necessary but often forgotten link between trust (dependence or belief on some system's properties) and trustworthiness (the merit of that system to be trusted, the degree to which it meets those properties, or its dependability).

The uncertainty described above, together with the vast amount of exposure to wrong-doing endured by current systems, forms an explosive combination. In order to handle it and obtain assurance on the correct operation of systems, all efforts are not too much. The tolerance perspective on security, currently termed intrusion tolerance, shed new light on a road darkened by the growing difficulty of preventing every intrusion on large and complex and uncertain installations. The combination of fault and intrusion tolerance closes the final gap, by allowing the design of systems that become simultaneously secure and dependable through the same class of mechanisms.

This is not enough though, we need system design principles that ensure a global and accurate view of the relation between trust and trustworthiness. This goes well beyond technological factors: 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 will nevertheless be deployed due to market pressure: will be perceived with suspicion by users; will be managed by a restricted group of "experts", increasing info-exclusion; may very well be mismanaged, yielding cyber-crime, e-frauds, cyber terrorism and sabotage.

2. Core challenges

If we were to enumerate core challenges, two come immediately to mind: adaptability; availability. There are many ways to achieve adaptation, to remain available. However, one can and should ask: how many will survive in the new order of uncertainty versus predictability?

2.1. Trustworthy adaptability

As the complexity of systems is growing out of control, amplified by the advent of ambient-intelligent pervasive and ubiquitous computing, and by the growing deployment of "always-on" complex systems, it becomes practically impossible to predict a system's behaviour at design time. In consequence, this requires adaptation.

However, if any trust is to be put on the operation of these systems, it had better be up to the extent of their trustworthiness. Adaptation should not undermine trustworthiness, adaptation should be dependable. Today, when adaptation (e.g., QoS) is largely driven by heuristics, the former objective is a research challenge, requiring innovative approaches such as proactive-reactive design under uncertainty, adapt-

ing functional and non-functional properties while providing guarantees on adaptation result, autonomous and decentralised system algorithmics, trustworthy monitoring and update for continuously-on systems.

2.2. Rethinking availability

Classical dependability being based on aprioristic and all-or-nothing criteria, it is not surprising that availability has been designed-in (by redundancy) according to forecasted fault modes, and predicted/ contracted to the user upon deployment. A system is either available, or unavailable. Whilst availability will remain a crucial attribute of systems, the future will bring new variables that will invalidate the above-mentioned status-quo: dynamics, uncertainty, evolvability, mobility, energy, maliciousness.

In consequence, this requires some rethinking of the concept of availability, an approach where availability becomes itself an evolvable and survivable attribute, in essence conditioned by: evolution of the environment; oscillation in QoS of the infrastructures. However, guarantees must still be met in some form, and this is another hard research problem, encompassing technology and societal factors (such as managing user expectations). Essentially, a totally new perspective on the 24x7 problem equation in a world of threats, what we might describe as 'acceptable availability'.

2.3. Hybrid and modular system models

Consider a component or sub-system for which a given benign behaviour or failure mode is assumed. Let us call it a trusted component. That is, we believe the component possesses a set of "good" properties. How often is it the case in current system designs when the merit of that system to be trusted— the degree to which it meets those properties, or trustworthiness— were not analysed, to assess whether it was commensurate to the trust been put on the former? Furthermore, how can we enforce trustworthiness of the component vis-a-vis the assumed behaviour, that is, coverage of such assumptions, given the unpredictability of attacks and the elusiveness of vulnerabilities?

One sensible way is to cast the notion of "better" components in the system model from the very start. Firstly, resorting to modular and hybrid distributed systems models, representing the fact that the presence and severity of vulnerabilities, attacks and intrusions varies from component to component. Secondly, enforcing these assumptions by construction, using architectural hybridisation, so that failure assumptions are in fact enforced by the architecture and the construction of the system components, and thus substantiated. That is, a trusted component is made trustworthy enough to match the implied trust.

3. Critical Infrastructures, the perfect example

The problem of resilience of critical infrastructures is not completely understood, mainly to the hybrid composition of these infrastructures: SCADA (supervisory, control and data acquisition) systems which yield the operational ability to supervise, acquire data and control; interconnections to the standard corporate intranets and often unwittingly to the Internet; advent of distributed generation.

There is in consequence an increase in the risk associated with using technologies in ways not adequate to the roles they were planned for. This risk, on the other hand, is not well mastered, and current configurations probably risk far more damaging failure scenarios than probably anticipated. Since it is economically infeasible to globally change technologies, one has to make legacy systems work better whilst allowing the system to live in this interconnected future. Suddenly, the potential to cause physical damage is accrued to the known Internet-borne failure scenarios

In fact, to a researcher, they present a set of significant challenges very much related to the previous discussion: enormous, increasing, unstoppable dependence on the above for all societal activity (e.g., Lisbon Agenda); uncertainty, generated by openness, dynamics and interdependence of the "Internet" to which they are connected; predictability required of services provided by CI's.

These problems are complex and must be tackled with the right weapons. Not surprisingly:

- Simultaneously under a security and a dependability viewpoint, a trustworthiness perspective;
- Achieving predictability in uncertain conditions, a dependable adaptability perspective;
- Encompassing correctness and continuity of service under a holistic viewpoint, a resilience perspective.

References

[1] A. C. P. V. Pedro Martins, Paulo Sousa. A new programming model for dependable adaptive real-time applications.

[2] P. Veríssimo. Travelling through wormholes: a new look at distributed systems models.

[3] P. Veríssimo. Uncertainty and predictability: Can they be reconciled? In *Future Directions in Distributed Computing*, volume 2584, pages 108–113. 2003.

[4] P. E. Veríssimo, N. F. Neves, and M. P. Correia. Intrusion-tolerant architectures: Concepts and design. In R. Lemos, C. Gacek, and A. Romanovsky, editors, *Architecting Dependable Systems*, volume 2677. 2003.