

Combining Theory and Systems Building Experiences and Challenges

S. Terzis

The Global and Pervasive Computing Group
Department of Computer and Information Sciences
University of Strathclyde, UK
`Sotirios.Terzis@cis.strath.co.uk`

Abstract. Having recently been through a successful project that followed a combined theory and systems building approach, in this position paper I reflect back on the experience in order to see whether there are any general lessons to be drawn. This reflective analysis leads to the conclusion that establishing a basic understanding and an appropriate level of abstraction is essential for systems builders and theoretician to collaborate successfully, with algorithmic work and the operationalisation of formal models playing a major role in this process. Within pervasive computing systems building research some prior consolidation work is required to make the combined research approach effective.

1 Motivation

The motivation in writing this paper stems from the experience in participating in the SECURE¹ project[2]. This project was a new and exciting experience as it involved research that set out from the start to follow a combined theory and systems building approach. What makes this project particularly interesting is the fact that it was considered a success, both by the members of the research team and external reviewers [6], leading to a number of publications in both theory and systems building forums, like the International Conference on Software Engineering and Formal Methods and the International Conference on Distributed Computing Systems, and the IEEE Pervasive Computing Magazine and the International Conference on Trust Management respectively. It also provided strong evidence on the potential of a combined theory and systems building research approach, and to some extent influenced the thinking of the second Global Computing IST FET research initiative that solicited proposals adopting such a combined approach[11].

More than six months after the official end of the project, and at a time when the need for a combined theory and systems building research approach in pervasive computing becomes pressing, it is a good time to reflect back on the experiences of the SECURE project in order to identify what were the challenges

¹ SECURE (Secure Collaboration among Ubiquitous Roaming Entities) FET funded project IST-2001-32486.

that the combined research approach presented and how they were addressed within the project, in an attempt to provide some useful insight to researchers and to further promote the combined research approach. However, readers should be warned that this reflective analysis is from the point of view of a system builder, and the views presented in this paper are the author's alone and in no way reflect the views of the rest of the SECURE project team.

2 A brief overview of the SECURE project

The aim of the SECURE project was to explore the human notion of trust as a basis for making access control decisions in a global computing environment. The intention of the SECURE project was not to accurately model or to fully replicate the way in which humans make trusting decisions, especially so as the process humans use is not fully understood, and is well-known to be easily defeated by adversaries as evidenced in the variety of reported trust scams. Instead, the aim was to focus on particular characteristics of this process that are desirable in a global computing context. These characteristics are the subjective nature and dynamic character of trust.

More specifically, in a global computing environment entities are expected to meet and collaborate with previously little known or even completely unknown entities without reliance to the availability of any security infrastructure [11]. The subjective nature of trust allows entities to operate autonomously in a decentralised manner, and to vary their trust assessment of similar situations. At the same time, the dynamic character of trust allows entities starting from complete lack of information about an entity to form and evolve an opinion about its trustworthiness, using evidence collected either through the experience of interacting with the entity in question or by propagation of other entities' opinions about its trustworthiness.

In this context, the SECURE project aimed to devise a decision making process based on three distinct but closely related models. A formally grounded trust model that allows entities to express opinions about the trustworthiness of others. A risk model that allows entities to assess the risks involved in an interaction with a particular entity. A collaboration model that combines the trust and risk models allowing entities to exploit the opinions about an entity's trustworthiness in assessing the risks an interaction with it entails and at the same time incorporates the dynamic aspects of trust enabling entities to form and evolve their opinions in accordance to collected evidence. In addition to the above models, a central objective of the project was the design and implementation of trust-based decision making engine, that incorporates the three models and is referred to as the SECURE Kernel [1]. Finally, in order to achieve these goals the project decided to follow an iterative approach based on application scenarios that in the early stages were used to drive the requirements elicitation process and in the latter stages to validate the results of the project.

2.1 Project outcomes

In this sections the aim is not to provide a comprehensive presentation of all the outcomes of the SECURE project. For this interested readers should refer to the series of project deliverables (see [12]). Instead, the aim is to provide a brief overview of those results most relevant to the current discussion.

Theoretical outcomes. One of the core outcomes of the SECURE project was a formally founded trust model[4]. According to the model *trust values* are members of an abstract domain with two orderings, an *information ordering* representing the amount of information or knowledge of a particular trust value, with a least element denoting the lack of any information or knowledge; and a *trust ordering* representing the degree of trustworthiness, with a least element denoting complete mistrust and a greatest element denoting absolute trust. The values of the domain form a complete partial order with a least element according to the information ordering and a complete lattice according to the trust ordering.

Moreover, the model also considers each entity as being equipped with a *trust box* that supports two operations. One that provides a trust value for each principal in accordance to *local trust policy* (π), and one that allows the entity to change its local trust policy. An important characteristic of the local trust policies is that they allow entities to express trust opinions in reference to the opinions of other entities. This characteristic in combination with the requirement that local trust policies are continuous functions, and the above properties of the trust value domain enables the definition of an entity’s overall trustworthiness in terms of fixed-point arithmetics, i.e. as the fixed point of a global trust function (II) induced by the collection of all local trust policies in the system.

Additionally, a more concrete instantiation of the above trust model based on *event structures* was also produced [9]. In this instance trust-mediated interactions were modelled in terms of *observable events* according to an event structure that defines *necessity* and *conflict* relations between them. This allows trust values to take the form of anti-tone functions from *event configurations*² to $(s, i, c) \in N^3$ triples denoting the amount of evidence, i.e. count of observed events, in support, inconclusive, and in conflict to a particular event configuration respectively. In this case, the two orderings of the trust value domain take the form of comparison operators for triples, allowing us to tell whether a triple represents more information, or more evidence in support of a configuration than another.

Besides the production of the formally founded trust model, another important outcome of the project was the production of a collection of techniques [8], and of a high-level programming language, namely a “calculus for trust management” (**ctm**) [3], that operationalise the model. The techniques enable entities to compute approximations of the trust values, i.e. fixed-points of the global trust

² Conflict-free and necessity closed subsets of events of a particular event structure.

function Π in several ways. More specifically, they allow the distributed computation of fixed-point approximations over a global, highly dynamic decentralised network, the checking of certain properties of the fixed-point value, i.e. whether it exceeds in terms of trustworthiness and/or information a certain threshold, and the exploitation of already computed intermediate results when local trust policies are updated during the fixed point computation. The `ctm`-language enables abstract description of entity behaviour allowing the expression of invariants that can be statically checked. At the same time, the fact that the development of the `ctm`-language was guided by process algebra, allows the exploitation of standard theory from process algebra to define various notions of equivalence between entities programmed in this language. This enables the application of another set of reasoning techniques for providing static guarantees.

Furthermore, another outcome of the project was the development of general reasoning techniques that can be applied to provide provable security-guarantees in the trust model [7]. More specifically, the exploitation of the concrete instantiation of the model described above, enabled the definition of a declarative language, a variant of pure-past temporal logic, for the specification of formal properties of past entity behaviour. This language combined with the definition of two alternative algorithms for the verification of such properties enables the definition of systems that provide a form of provable “security” properties of the form: “if principal p gains access to resource r at time t , then the past behaviour of p up until time t satisfies requirement ψ_r ”.

System building outcomes. In addition to the formally founded trust model, the SECURE project also developed a risk and a collaboration model. The risk model considers trust-mediated actions, each with a set of potential outcomes with an associated cost and benefit. It defines risk as the combination of the likelihood of an outcome occurring combined with its associated cost or benefit. On the other hand, the collaboration model combines the trust and risk models and defines three process, namely decision making, trust and risk evaluation [10]. The decision making process determines interaction outcome likelihood from the trust value of the interacting entity and uses a decision theoretic process to determine whether an interaction request should be accepted or not. The trust and risk evaluation processes enable entities to revise trust values and outcome costs respectively, to reflect collected evidence in the form of either direct observations of past interactions or indirect recommendations from other entities.

A core outcome of the SECURE project is a framework architecture for trust-based access control decision making that has been instantiated as a Java-based SECURE kernel [1]. The framework incorporates a series of components, including ones that provide generic implementations for of the trust and risk models. These components combined together are able to carry out the processes of the collaboration model. It should also be pointed out that the trust component of the kernel provides a generic implementation of the event structure based instantiation of the trust model outlined above. Moreover, the kernel has been

used in the implementation of a trust-based collaborative email spam filter that identifies spam email messages on the basis of the trustworthiness of the sender.

Finally, additional outcomes include the design and implementation of an entity recognition scheme and an evidence distribution framework.

3 Reflective analysis

The main challenge in any project attempting to follow a combined theory and system building research approach is that both fields over the years have built a significant volume of literature describing their respective research results. Consequently, it is not realistic from experts of either field to expect to be able to assimilate this amount of knowledge in order to identify the most appropriate theoretical and system models and to effectively combine them to address fundamental challenges. The SECURE project addressed this challenge through a research team that combined researchers with significant expertise in both fields. Not only that, but the work of these researchers was mainly focused to the field of their expertise. However, there were also some members of the team that their research traded on the boundaries between the two fields. This is not particularly unusual in the area of security systems research where the benefits of formal techniques in security policy and protocol verification are widely recognised.

Having a research team comprising of both theoreticians and systems builders within a project introduces problems of communication that present another significant challenge. These problems are further amplified in the context of highly technical discussions that are necessary to facilitate project progress. Within the SECURE project scenarios were extensively used to address the communication challenge. Scenarios can play a major role in any project to explore the strengths and weaknesses of developed models, both theoretical and systems ones. For this purpose a large number of quite rich scenarios is necessary. However, in the context of facilitating communication it is important that the project is limited to a quite small number of carefully selected scenarios. In selecting these scenarios it is important that they are both very simple in the interactions they describe in order to facilitate understanding of model concepts, and flexible enough to allow multiple perspectives to be explored ensuring coverage of all aspects of the model. In the SECURE project a scenario involving the use of an electronic purse for purchasing bus tickets with e-cash was used for this purpose. The main characteristic of the scenario is that it involves a number of trust relationships (client-bus company, bus company-bank, client-bank), but in each of them entity interactions can be largely simplified.

In addition to the use of scenarios, communication problems were also eased off by a clear understanding of core concepts of the problem domain. In this respect, the SECURE project had the advantage that it could build on previous work especially in the social sciences, which provided a comprehensive framework explaining various characteristics of trust and their interrelationships. Even though such previous work can be very useful in the beginning of the project, sustained effort is required throughout to ensure divergence does not creep in.

In this respect, maintaining and using a project glossary proved very useful for the SECURE project. Finally, the researchers with experience in both theory and systems building also played an important role in facilitating concept understanding.

The final challenge in combining theory and systems building is managing abstraction. In general, theoreticians are used to abstract away from implementation details and work with quite abstract constructs. In theoretical research abstraction is crucial as it allows researchers to focus on the essential properties of the problem domain. However, theoreticians need to also keep track of the fact that their abstract models are useful only in so far as they support the analysis and development of real systems. This requires particular attention to the operationalisation of the abstract models. In order to ensure this the SECURE project set out from the beginning as one of its objectives that the formal model of trust should also be operational. This was another area where the problem domain itself played an important role in providing a concrete target for the project. In the security domain the ability to formally verify that developed systems provide certain guarantees is key, as there is too much at risk. Not only that, but evidence based trust management has been previously criticised for failing to provide the traditionally expected strong guarantees. As a result, it was clear from quite early on that the project should also provide appropriate reasoning and verification techniques.

On the other hand, systems builders are used to work on much lower level of abstraction and quite often find it difficult to understand abstract theoretical models. So, quite often they tend to shy away from the theoretical models and just focus on their implementation work. This often results in a situation where the differences between accidental and real complexities of the problem domain are unclear, as systems builders fail to see that the same kind of functionality can be achieved in widely varied ways. In this respect, the SECURE project was lucky to include different research groups with systems building experience that worked independently and concurrently on the development of a number of application scenarios. This demonstrated from early on a variety of possible approaches to the problem, and helped in pointing out the accidental complexities.

However, identifying accidental complexities on its own is not enough, the systems building and theoretical point of view need to converge towards an appropriate level of abstraction. For the convergence to happen system builders and theoreticians need to engage with each other from the start and be prepared to listen to and adapt their work to accommodate the needs of the other side. In some sense keeping each other straight by not permitting researchers to shy away from the difficulties that the other field presents. In this process, the operationalisation of the formal models and algorithmic work can play a major role, as they are the point where the two aspects of the project are integrated. Within the SECURE project, this kind of process led to the development of the event based instantiation of the trust model and motivated the work on the distributed fixed-point computation algorithms.

More specifically, the abstract trust model, although it was clear that it captured the essential characteristics of the problem domain, it was also very clear that the fixed-point computations it required were very problematic in a distributed, highly dynamic and decentralised environment. As a result, this led to the investigation of the techniques for the distributed approximation of fixed-points mentioned above. At the same time, the fact that the abstract trust model left the exact format of the trust values unspecified made its implementation very application dependent. This made it very difficult for the SECURE kernel implementation to provide any integrity assurances to application developers. Moreover, the lack of an explicit notion of evidence within the abstract trust model made very difficult its integration with the risk and collaboration models where the notion was central. This meant that the details of how the three models would be integrated was also left to the application developers leading quite often to very ad hoc approaches. The event based instantiation addressed all these problems by introducing a particular format for the trust values tied to a notion of evidence, as event configuration histories. This process was facilitated by close collaboration between systems builders and theoreticians, including a number of exchange visits with clear objectives and a concrete target outcome.

4 Conclusions

In conclusion, a combined theory and systems building research approach requires the collaboration of systems builders and theoreticians. Such a collaboration presents some serious challenges, first in developing a common language and a basic understanding for scientific communication, and second in identifying an appropriate level of abstraction that both systems builders and theoreticians are able to work at. To overcome both of these challenges requires serious commitment and effort from all researchers involved in the project. For this reason, it is important that from the outset the project sets both systems building and theoretical objectives ensuring that it involves research of value for both fields. However, it is even more important that it identifies a clear product where both theoretical and system building work will be integrated, thus providing a clear target for the convergence of the work. At the same, it should not be underestimated how important it is for researchers to have from the start a clear picture of the expected benefits of the combined research approach. This is easier in some domains (e.g. security and safety critical systems engineering) than in others.

Taking the above into consideration, in the area of pervasive computing research there is still a lot of work to be done. Despite the fact that significant progress has been made over the years both in systems building and theoretical terms, work still remains largely disparate and disconnected, not only across both fields but also within them. It is certainly a recurring criticism of pervasive computing systems building research that very little effort has been put so far into the consolidation of results across research projects [5]. As a result, the community is still unclear on what really works. To make matters worse there does not exist any kind of consensus on the definition of the core concepts of

the domain, or as a matter of fact on what these concepts should be. However, such a consensus is crucial for the combined approach to succeed. So, maybe we (systems builders) should start by first putting our house in order.

References

1. Ciarn Bryce, Paul Couderc, Jean-Marc Seigneur, and Vinny Cahill. Implementation of the secure trust engine. In P. Herrmann, V. Issarny, and S. Shiu, editors, *Proceedings of the Third International Conference on Trust Management (iTrust 2005)*, volume 3477 of *LNCS*, pages 397–402. Springer, May 2005.
2. V. Cahill, E. Gray, J.-M. Seigneur, C. Jensen, Y. Chen, B. Shand, N. Dimmock, A. Twigg, J. Bacon, C. English, W. Wagealla, S. Terzis, P. Nixon, G. Serugendo, C. Bryce, M. Carbone, K. Krukow, and M. Nielsen. Using Trust for Secure Collaboration in Uncertain Environments. *Pervasive Computing Magazine*, 2(3):52–61, July-September 2003.
3. Marco Carbone. *Trust and Mobility*. PhD thesis, Department of Computer Science, University of Aarhus, 2005.
4. Marco Carbone, Mogens Nielsen, and Vladimiro Sassone. A formal model for trust in dynamic networks. In *Proceedings of the International Conference on Software Engineering and Formal Methods*, pages 54–63, Brisbane, Australia, September 2003.
5. Nigel Davies. Proof-of-concept demonstrators and other evils of application-led research. In *Proceedings of the Pervasive 2005 Workshop on What makes for good application-led research in ubiquitous computing?*, May 2005.
6. IST Results Feature. Trust me, i'm a machine, <http://istresults.cordis.lu/index.cfm/section/news/tpl/article/browsingtype/features/id/75057/highlights/secure>, 2005.
7. Karl Krukow, Mogens Nielsen, and Vladimiro Sassone. A framework for concrete reputation-systems with applications to history-based access control. In *Proceedings of the 12th ACM Conference on Computer and Communications Security (CCS'05)*, pages 260–269, Alexandria, VA, USA, November 2005.
8. Karl Krukow and Andrew Twigg. Distributed approximation of fixed-points in trust structures. In *Proceedings of the 25th IEEE International Conference on Distributed Computing Systems (ICDCS'05)*, pages 805–814, Columbus, OH, USA, June 2005.
9. Mogens Nielsen and Karl Krukow. On the formal modelling of trust in reputation-based systems. In J. Karhumki, H. Maurer, G. Paun, and G. Rozenberg, editors, *Theory Is Forever: Essays Dedicated to Arto Salomaa*, volume 3113 of *LNCS*, pages 192–204. Springer, 2004.
10. Sotirios Terzis, Waleed Wagealla, Colin English, and Paddy Nixon. Trust lifecycle management in a global computing environment. In C. Priami and P. Quaglia, editors, *Post-Proceedings of the Global Computing 2004 Workshop*, 2005.
11. Global Computing Initiative Website. <http://www.cordis.lu/ist/fet/gc.htm>, 2002.
12. SECURE Project Official Website. <http://secure.dsg.cs.tcd.ie>, 2002.