The Navigation Metaphor in Security Economics

Pieters, Wolter; Barendse, Jeroen; Ford, Margaret; Heath, Claude P. R.; Probst, Christian W.; Verbij, Ruud

Published in:
I E E Security & Privacy Magazine

Link to article, DOI:
[10.1109/MSP.2016.47](https://doi.org/10.1109/MSP.2016.47)

Publication date:
2016

Document Version
Peer reviewed version

[Link back to DTU Orbit](#)

Citation (APA):
Pieters, W., Barendse, J., Ford, M., Heath, C. P. R., Probst, C. W., & Verbij, R. (2016). The Navigation Metaphor in Security Economics. *I E E Security & Privacy Magazine*, 14(3), 14-21. <https://doi.org/10.1109/MSP.2016.47>

General rights

Copyright and moral rights for the publications made accessible in the public portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognise and abide by the legal requirements associated with these rights.

- Users may download and print one copy of any publication from the public portal for the purpose of private study or research.
- You may not further distribute the material or use it for any profit-making activity or commercial gain
- You may freely distribute the URL identifying the publication in the public portal

If you believe that this document breaches copyright please contact us providing details, and we will remove access to the work immediately and investigate your claim.

The navigation metaphor in security economics

Wolter Pieters

Jeroen Barendse

Margaret Ford

Claude P. R. Heath

Christian W. Probst

Ruud Verbij

Abstract

Security models and security economics have been separate developments for a long time. Models represent the organisation under scrutiny with possible attack paths, and security economics covers the effect and cost of attacks and counter-measures. This inhibits progress in decision support for security investment. The navigation metaphor merges these two concepts: navigation on security models can identify optimal attacker and defender decisions for multi-step attacks, based on “maps” of the system being studied. Routes on the map represent attacks on the system. Economic optimisation analyses can identify the most efficient routes for gaining access to certain targets from the point of view of an attacker; this insight is used to optimise the defences on these routes from the point of view of the defender. In this article, we discuss the achievements and the challenges of the navigation metaphor in cyber security.

Keywords: attack navigators, attacker profiles, navigation metaphor, security economics, visualisation

1 The navigation metaphor

In the physical world, navigation is a well-understood concept: if you want to get from A to B, you can use a navigation system to plan your route. Such systems can help you to optimise your behaviour by finding the most efficient route, and can even adapt dynamically when everybody takes the same route. However, navigation systems may also be used for other purposes. In order to prevent you from reaching a particular destination, I could identify your most likely routes and sabotage them, thereby significantly affecting your travel time, or even the likelihood of your arriving at all.

This is precisely the idea of the “navigation metaphor” for cyber security, which is being developed in the TRE_SPASS project. By identifying the most efficient routes

for gaining access to certain targets from the point of view of an attacker, the defences on these routes can be optimised for the benefit of the defender. This combines models of security architectures (maps) with economic optimisation strategies (route planning), integrating these different strands of security research. The navigation metaphor, in conjunction with insights from risk management and visual design, provides a very powerful security decision support tool. Specifically, the navigation metaphor supports the economic analysis and visualisation of multi-step attacks: just as successive roads lead to a single physical location, so ordered combinations of actions are required to reach the attacker's goal. By using navigation as an analogy, it becomes easier to motivate and explain security investment decisions to a wide audience.

In this article, we introduce the navigation metaphor in cyber security and discuss its strengths and limitations. We do this by comparison with different views and steps in real-world maps and navigation:

- a satellite view which offers realistic aerial pictures, but not the underlying infrastructure;
- the map view, which shows the underlying infrastructure and enables route calculation, but does not show the actual look of your surroundings;
- route calculation, which optimises travel to a specific goal, as well as strategies for blocking such travel.

We illustrate our techniques with a case study on online services via interactive TV, as a running example.

2 Satellite Images

The satellite view in navigation systems provides an overview of the natural surroundings. This direct mapping between the local environment and the display presented to the driver, promotes navigation and orientation, supporting the widespread adoption of navigation systems. Obtaining satellite images is relatively easy; creating a map from them is not, especially when travel is on foot or by bike. While the big, public roads are well known and easy to detect on the image, smaller, private tracks and paths may be hidden by vegetation, buildings, or even clouds, or may have been forgotten and are only found by chance.

In navigation systems for cyber security, the satellite image corresponds to the external view of the organisation, its divisions, computers, infrastructure. As with its geographical counterpart, some of these components are visible, others may be private, others again may exist but have been forgotten. As in geographical navigation systems, the navigation metaphor for cyber security depends on the collection of contextual information. Acquiring such information is the first step in analysing the security of a system with the navigation metaphor.

As a running example, we will look into a TRE_SPASS project case study exploring the delivery of home-based services via a TV interface. Our case study partner, despite its small footprint, achieves significant social impact by delivering services through a range of partners in different sectors. The development of this case study illustrates the value of using attack navigator maps and visualisations to support service planning. This process includes mapping multiple partner relationships, with their complex inter-dependencies and diverse security implications.

In our case study, the enterprise architecture modelling language ArchiMate [5] was first used to map the digital services, in collaboration with our case study partners. The resulting diagrams (“maps”) were easily understood by stakeholders, but needed greater focus on social, organisational and partnership features. The partners felt that the resulting map lacked essential elements of the system.

In the next session, we aimed to establish the satellite view of the proposed service instead. We conducted a pre-study briefing with senior managers, exploring the organisation’s goals, culture, business model, and projects past, present and future. From these discussions, a set of security concerns emerged relating to untrusted behaviours by both users and outsiders. Building on text-based target graphs developed from the results of the pre-study, we asked each participant to consider a typical service that they work on and to construct a drawing of how they would undertake a typical data management task.

Subsequently, we introduced *LEGO* bricks to the case study service design participants. This led them to co-construct a rich multi-perspectival picture of data-sharing as a part of the service [6]. We asked participants to model (using the colours and language of ArchiMate) the central actors (yellow bricks), infrastructure (green bricks), data (blue bricks), and locations (pink tiles) (Fig. 1). The group agreed on a narrative associated with the modelling process, and the weighting and positioning (and repositioning) of elements of the model.

The satellite view of the proposed service showed regions of trust between actors in the modelled environment, and data flows crossing the boundaries of these regions. Actors were modelled as *LEGO* avatars to represent the central actors and the control strengths of selected points along data-paths. Participatory techniques included mapping (the domain target), sequencing (order of linked events), listing, placing (the relevance of previously established values), comparing (the viewpoints of different characters), and linking (assessing the implications of actions upon different actors).

Figure 1: Digital collage of the results from two *LEGO* participatory mapping sessions. Case study participants designing an IPTV home-banking service used colour-coded *LEGO* bricks and figures to map the service infrastructure and the role of business actors. Two sessions are summarised here in a digital collage, showing the relationships that emerged. In the central loop, the service is carried forward, clockwise, starting with the client and moving to the provider and their business partners. Below, in a different loop, the banking platform assists cloud-based transactions made with a card. Above, the client receives income.

The satellite views for our navigation metaphor can be refined over time; in a second *LEGO* session, the group reflected upon and remodelled the weaker parts of the service design, adding bricks where necessary. The nature of actors was represented by avatars to show their business role, world view and degree of

influence, some being diminutive and others overbearing (Fig. 1). Lines of defence were added to show relationships and areas of vulnerability.

A particular benefit of three-dimensional physical modelling is the ease of incorporating annotations, while the group keeps track of any working assumptions. Apart from the use of avatars, the group developed colour-coding for different types of relationships and data-flows, sometimes increasing the height of defences for greater control strength. A key feature of the resulting model is the ability to view it from every side and track all the risk implications for each of the participating entities from their own perspective. This is vitally important in smaller, more flexible organisations which rely heavily on working in partnership in what have been called relational services [3], where human relationships effect the continuation and extension of the service.

3 Maps

The satellite view of an organisation developed in the first phase forms the basis for the navigator map. The goal of the first step was to obtain a precise satellite image, revealing as many elements as possible. The goal of the second step is to translate this image into a formalised map for cyber security navigation.

There is quite a history of map-style network models in security. Trees may be used to model the infrastructure by representing containment, in the sense that a computer is located in a room in a building in the world. In many domains, and also in security, infrastructure tree models are not sufficiently expressive to deal with the real world. Information is not contained within one clear boundary or perimeter, like a safe or an offline machine, but can be accessed via many possible routes. A network model (graph or map) is therefore more suitable in such cases, although it comes at the cost of more complex analysis (e.g. [1]).

However, such network models have generally been limited to the technical parts of an organisation's infrastructure, typically representing computer networks and hops of a hacker from one node to another. Focusing only on the computer network limits the analysis possibilities, just as navigating only along highways limits navigation possibilities. Many attacks contain some form of social engineering and/or physical access, and it is therefore vital to include humans *and* physical locations in the map, with their role in obtaining access. The navigation metaphor introduced in this paper uses such "socio-technical" network models (Fig. 2) as maps, which form the basis for navigating to the goal of an attack [8]. These maps are essentially graphs with nodes and connections, enabling the analysis of routes. In the Fig. 2, one can recognise formalised elements from the Internet service case study, such as locations (bank, home), digital infrastructure (computers and connections) and actors with possessions (Alice, Charlie).

Figure 2: Example network model (map) for socio-technical security. Entities are represented as boxes, and data as circles. Dashed boxes represent access control policies.

Developing a map from the satellite view is largely based on domain knowledge of members of the organisation, supported by software tools that assist in the translation according to the rules of the map formalism. Although some information will be lost in the translation from satellite image to map, the mathematical structure of a map enables quantitative analyses and optimisations.

Network models allow users to model entities and relations within their access control space, independently of attacker goals. For example, rather than thinking about vulnerabilities on a particular database, the user would first map the infrastructure around this database, including digital, physical, and social access. This decoupling of attack opportunities and modelling is crucial for ensuring that no relevant attack opportunities are overlooked. As geographical maps represent different entities and their connections, cyber security maps represent different assets and possible activities. This information forms the basis for maps that can represent “attack navigation” in socio-technical systems – identifying possible attacks by picking a starting state and a target node on the map.

4 Routes

With geographical maps, navigation amounts to identifying and optimising the means to reach a goal using different modes of transportation and routes; in the attack navigator, this amounts to identifying and optimising ways in which an attacker can reach a particular asset [8]. This analysis is based on assumptions of economic rationality: short and cheap routes are better for the attacker. (Long and costly ones are obviously better from a defender point of view.)

Some kind of route planning can already start from the satellite view. Once the most important aspects, such as the main actors, locations, and assets have been identified and agreed upon by all participants, the environment can then be explored from many different perspectives. In particular, physical models enable the stakeholders to move actors around physically. This enables stakeholders to evaluate the attacker perspective and strategy in relation to future users of the service, exploring the strategic and economic implications of different relationships. These include the potential for financial abuse of users of the system, economic risks posed by rivals and even partner organisations, as well as more standard technical and communications security risks.

More formal analysis is possible based on the map, in particular answering the question of optimal routes for the attacker. This is equivalent to planning a journey by identifying the destination, and generating a sequence of actions to perform to reach that destination. In comparison to road navigation, attack navigation deals not only with where the attacker can go, but also with what he can take along, for example credentials. Therefore, an optimal route for the attacker may involve visiting different places or people, obtaining credentials from those, and then reaching the information that can be disclosed with those credentials. In this sense, “has-access-to” is a fundamental relation in the navigation system: the attacker has certain access in the beginning, and may gain access to new places, people, and credentials by his actions.

However, such access is not cost-free. Travelling on a navigator map, like in the real world, costs something. This “distance” can be represented as cost, time, or likelihood of success/failure, associated with connections and access control policies on the map. All of these constitute *difficulty metrics* that say something about the expected effort an adversary has to expend in order to gain access. For example, opening a door by force may take 3 minutes, but with a key it may only take 10 seconds. As well as time, money (costs) may be used as annotation, but also for example likelihood of detection.

4.1 Vulnerability functions

The route proposed by a navigation system depends on settings such as car speed and efficiency preferences: if you have a faster car, you may want to take a different route, depending on infrastructure constraints such as road condition and speed limits. An obvious example is a shorter route that is only accessible to 4WD vehicles. The infrastructure properties *plus* the properties of your car determine how much time a particular part of the route will take you (with infinite for impossible), and thereby also determine the optimal overall route(s).

In security analysis, this corresponds to the relationship between the system properties represented in the map and attacker properties, which together determine properties of attack steps, such as required time and likelihood of success. Such a

basic separation between attacker properties and system properties has been proposed in the Factor Analysis of Information Risk taxonomy [14], whereby the likelihood of success of an attack step is determined by *threat capability* and *control strength*. The attack navigator uses vulnerability functions describing a relation between the threat capability and the likelihood of success to represent “resistance” of system components to attacks. In essence, this is a calculation of *action properties* from *agent properties*, similar to the way in which the outcome in many games is determined by the “skill levels” of the characters plus a source of randomness (dice). From a more scientific point of view, item response theory and Elo ratings (e.g., of chess players) achieve similar conceptual benefits [13].

4.2 Attacker profiles

The separation of attacker properties from system properties implies that attacker profiles are required in addition to navigator maps. Attacker profiles [2] are equivalent to individual car details entered into a navigation system; they specify attacker attributes such as skill and budget, and can be used to identify feasible attacks and their properties.

Within the TRE_SPASS project, we have developed several different strategies for specifying attacker profiles. For example, we can assume that attackers will be unable to execute attack steps where the difficulty level exceeds their skill level [9]. In a more quantitative setting, we can estimate the likelihood of success based on the difference between difficulty and skill [13]. Additional constraints may be imposed by the time and budget available to the attacker. In contrast to the skill constraints, these constraints are “additive”, in the sense that the attacker cannot execute attacks for which the *sum* of the costs of the steps exceeds the available budget. In a visual sense, skill or budget levels can be represented by the length of a bar, a number of blocks/items, or simply a number. When the attacker takes an action, his skill has to be higher than the corresponding difficulty, and his budget will be decreased by the cost of the action.

Although time, budget, and skill tell us something about the adversary, they do not provide any information about the attacker motivation or strategy. Not all attackers with the same time, budget, and skill will aim at the same attack vectors. Therefore, the question of the difficulty of attack paths should be complemented by the interest of the attacker in such attack paths, typically expressed in terms of the expected utility, which in turn depends on attacker motivation. Finally, attackers may not even choose the paths with the highest expected utility, for example because they have limited information, which forces us to make adversary strategy (or lack thereof) explicit [12].

4.3 Routing and weakest links

Using attacker profiles and the map, the attack navigator computes possible routes for an attacker to reach the goal of an attack. Sets of possible attack paths can be represented as trees. Attack trees [10] identify the different options available for an attacker to achieve a goal, and the properties of such attack paths, for example likelihood of success, cost, and time. Using extensions to this framework, defences can be added as well (attack-defence trees). Unlike existing attack tree frameworks, navigator maps can be used to generate an attack tree for each combination of goal and attacker, making the analysis more flexible. Whereas traditional attack trees do

not keep track of the system components involved in attack steps, attack navigators do exactly this.

This also provides a definition for what is called “weakest link”: the weakness of a link (system component) is determined by how much the utility for adversaries decreases when you remove the link (component) from the system, calculated over different possible attacker goals. In other words, if you remove an element or link from the system, how much more difficult or costly would reaching a goal become for the attacker, in comparison to the expected gain upon reaching the goal? Rather than setting a predefined goal that an attacker would be interested in, this becomes a multi-asset question, basically evaluating which assets the attacker can gain access to with positive expected utility, and how [11]. This is an economic analysis of contributions of system components to attackers’ expected utility.

From the perspective of an attacker, the generated attack trees also provide information that can be used to determine the optimal attack strategy, which in turn, from the defender point of view, may be used to determine which attack vectors are more likely, and where to direct investments. Based on the most efficient routes for gaining access to certain targets from the point of view of an attacker, we can then optimise the defences on these routes from the point of view of the defender.

4.4 Routes in the Internet Service

In contrast to navigation systems, we can use both the satellite view and the map of an organisation for identifying routes, that is possible attacks. While for the satellite view and the map this can be a manual process, it can also be (partially) automated by use of dedicated tools on navigator maps [8]. For the Internet service case study we developed an initial attack tree, annotated with values relating to the different attack steps. As with standard navigation systems, the attack navigator can visualise the attacks as routes through the model. However, the project has also developed new techniques for visualising vulnerability in such attack trees based on parameters such as difficulty, time, cost and probability per attack step, thereby visualising economic aspects of the attacks. This approach also allows for navigating specific threats, by highlighting important paths, zooming in, and reordering the tree (Fig. 3), each visualisation showing a different perspective on the same scenario. Visualisations of routes thus communicate the attacker perspective to the defenders. Based on such visualisations, stakeholders can identify the weak links in their system, which may suggest possible improvements by revising the architecture.

Figure 3: Radial attack tree for the Internet service case study. Colour, transparency and thickness of each edge is based on difficulty, required time and likelihood of success of the corresponding attack step.

Visualisations are an essential part of navigation systems, and also in the attack navigator they provide a greater level of understanding, as well as a clear means of communicating with the organisation. It also fostered a greater sense of ownership within the partner organisation, who expressed a desire to retain and update the emerging *LEGO* model even between workshop sessions. Their enthusiasm is typified by this quote: “I am in absolute awe at how you have managed to visualise and portray our sessions. It is very exciting to see our organisation all down on paper and at the same time very challenging in terms of next steps.” They also highlighted how “mapping out where we are, where we want to go, how we’ll get there clarifies all sorts of things.” In this respect, the mapping and navigation process provides a bridge between the detailed data gathering and analysis processes required to support sound security decision-making, and the commercial imperative to present key issues in a clear and concise form when dealing with senior decision-makers.

5 Optimisation

Once the attacker routes are known, economic analysis can answer different questions, in particular optimal strategies for both attacker (routes) and defender (“roadblocks”, i.e. controls to make routes more difficult). In particular, the effects of changes in the map (added controls) on optimal attacker strategies can be investigated, thereby providing a metric for the effectiveness of controls. This applies to digital architectures (such as cloud infrastructure) as well as socio-technical systems (such as the Internet service above).

We have applied the navigation metaphor in several additional case studies. For the Estonian Internet voting system, we estimated costs of individual attack steps, for example based on black market prices of infected machines. The optimal route depends on what an attacker is actually trying to achieve: changing a single vote, or changing the complete election result. Thus, we were able to map the required number of votes against the optimal strategy, and the cost of this strategy for that number of votes [15]. Using the navigation metaphor, one can compare such results against alternative architectures, such as an equivalent paper voting system. Also, one could specify requirements in terms of the minimum cost for an adversary to hijack one seat in parliament (or ten seats).

Another case study concerns fraud opportunities in telecommunication services, where the maps consist of value models of service architectures, and the routes of combinations of services that lead to monetary gain for an attacker [7]. For example, rogue telecommunication operators abroad may arrange a high volume of calls to their numbers, in order to obtain interconnection fees from honest operators. The analyses show possible adversary strategies, in terms of e.g. the required number of calls for making a profit. Again, changes to the architecture can be investigated in order to reduce the utility of adversaries. These different implementation variants of the metaphor illustrate its power in economic analyses of security.

6 The TRE_sPASS vision

Based on the methods and cases above, we believe the navigation metaphor offers great potential for further integration of economic and system models in cyber security, as well as visual designs. Navigation systems perform economic optimisation in terms of shortening travel times for travellers. Attack navigators do the same for potential cyber attackers, but now it becomes relevant how defenders can change the map. This combination of security architecture models and security economics enables new research directions as well as practical applications.

The navigation metaphor makes it easy to explain economic considerations of cyber security decisions to stakeholders, as shown in the Internet service case. If I change something on the map, how would the situation change for the attacker, from an economic point of view? This is basically a minimax optimisation [4] supported by the map. The combination of a map (model of the infrastructure) and an attacker profile also enables “adversary course of action” type reasoning for different types of attackers, providing flexibility under changing threat environments. If the threat environment changes, one simply re-runs the analysis with a different attacker profile. These are the key innovations compared to the state of the art.

It is notoriously hard to compare risk assessment approaches, because their effectiveness is ultimately determined by the losses due to real attacks. Even if one

could measure this, this would beg the question whether the attacks that actually occurred are representative for the threat environment of the system. For example, if method A predicts an annual loss expectancy of € 1 million, and method B predicts € 5 million, then what is the best method if the actual average annual loss after 5 years is € 2 million? Acknowledging these comparison difficulties, these are the main advantages we see of the navigation metaphor with respect to existing approaches:

1. The navigation metaphor enables a better conceptual understanding of weak links in a system, thereby supporting investment decisions;
2. The navigation metaphor stimulates consideration of different attacker profiles for the same system, acknowledging different economic strategies / utility functions of attackers;
3. The navigation metaphor invites the user to take the perspective of an attacker looking for opportunities, enabling new forms of learning;
4. The navigation metaphor stimulates innovation in visualisation of security economics and risk.

Obviously the navigation metaphor does not solve all problems of security risk management. Firstly, there are limits to the level of detail that can be represented on a map. Technical details of attacks, such as manipulations of code or program flow, need different types of analyses and visualisations. Maps can still show high-level access paths for gaining the required access, e.g. via different servers, via infected USB drives, or via social engineering. Secondly, the navigation metaphor does not solve the problem of data availability. Although our work helps stakeholders map their systems, there may still be uncertainties or errors in their input, which can propagate to the results. Links with security economics are important in terms of identifying costs of actions and impact of successful attacks. Thirdly, analysis of the maps, like many economic frameworks, depends on models of adversary choices, which can be complicated to construct and validate. We are currently investigating probabilistic models for this purpose.

The navigation metaphor is most useful in identifying attack opportunities in complex socio-technical systems, where attacks consist of multiple steps of which the relation is not immediately obvious. The strength of the metaphor lies in representing heterogeneous elements in a single formalism and analysis. By contrast, finding attack opportunities for individual system components, be they human or technical, should be identified by different types of tools.

On a more strategic level, we have been asked whether attack navigator tools could give advantages to attackers as well, in particular when methods are published and tools are open source. Ultimately, this relates to security by obscurity: can we deflect attackers by withholding information about system design? Even if the answer were positive, it would probably be more effective to protect the system design data, rather than tools for analysis of the system for optimal attack paths, based on the navigation metaphor.

We see the following key challenges for the coming years:

1. Further integration of different types of difficulty metrics for use in navigation analysis, such as CVSS (common vulnerability scoring system) or CWSS (common weakness scoring system) values, but also the results of social engineering experiments;

2. More advanced methods for taking attacker motivation and strategy into account in the analysis, in addition to “resource” properties such as skill, budget, and available time;
3. “App-style” interfaces for map development and maintenance, as well as “first-person” visualisation of routes.

We hope to be able to contribute to these challenges in the final stages of the TRE_SPASS project. In doing so, we aim at further improving stakeholder understanding of economic reasoning and investment decisions in cyber security, based on a well-understood metaphor. Whereas geographic navigation is about optimisation of reaching a goal on earth, cyber defence is about making such optimisation harder for the attacker in cyberspace. For our stakeholders, the navigation metaphor enables grasping and rethinking this fundamental economic relation between attackers and defenders.

Acknowledgements

The authors wish to thank Lizzie-Coles Kemp and Peter A. Hall for their input in developing the *LEGO* methods. The research leading to these results has received funding from the European Union’s Seventh Framework Programme (FP7/2007-2013) under grant agreement ICT-318003 (TRE_SPASS). This publication reflects only the authors’ views and the Union is not liable for any use that may be made of the information contained herein.

References

- [1] P. Ammann, D. Wijesekera, and S. Kaushik. Scalable, graph-based network vulnerability analysis. In *Proceedings of the 9th ACM Conference on Computer and Communications Security*, pages 217–224, New York, 2002. ACM.
- [2] T. Casey, P. Koeberl, and C. Vishik. Threat agents: A necessary component of threat analysis. In *Proceedings of the Sixth Annual Workshop on Cyber Security and Information Intelligence Research*, CSIIRW ’10, pages 56:1–56:4, New York, NY, USA, 2010. ACM.
- [3] C. Cipolla and E. Manzini. Relational services. *Knowledge, Technology & Policy*, 22(1):45–50, 2009.
- [4] L. A. Cox Jr. Game theory and risk analysis. *Risk Analysis*, 29(8):1062–1068, 2009.
- [5] E. Grandry, C. Feltus, and E. Dubois. Conceptual integration of enterprise architecture management and security risk management. In *Enterprise Distributed Object Computing Conference Workshops (EDOCW), 2013 17th IEEE International*, pages 114–123, Sept 2013.
- [6] C. Heath, L. Coles-Kemp, and P. Hall. Logical Lego? Co-constructed perspectives on service design. In *Proceedings of NordDesign 2014*, pages 416–425. Aalto Design Factory, 2014.
- [7] D. Ionita, S. K. Koenen, and R. J. Wieringa. Modelling telecom fraud with e3value. Technical Report TR-CTIT-14-11, Centre for Telematics and Information Technology, University of Twente, Enschede, October 2014.

- [8] F. Kammuller and C. W. Probst. Invalidating policies using structural information. In *Security and Privacy Workshops (SPW)*, 2013 IEEE, pages 76–81, May 2013.
- [9] A. Lenin, J. Willemson, and D. P. Sari. Attacker profiling in quantitative security assessment based on attack trees. In Karin Bernsmed and Simone Fischer-Hübner, editors, *Secure IT Systems*, volume 8788 of *Lecture Notes in Computer Science*, pages 199–212. Springer International Publishing, 2014.
- [10] S. Mauw and M. Oostdijk. Foundations of attack trees. In D. Won and S. Kim, editors, *Proc. 8th Annual International Conference on Information Security and Cryptology, ICISC'05*, volume 3935 of *Lecture Notes in Computer Science*, pages 186–198. Springer, 2006.
- [11] W. Pieters. Defining “the weakest link”: Comparative security in complex systems of systems. In *Cloud Computing Technology and Science (CloudCom)*, 2013 IEEE 5th International Conference on, volume 2, pages 39–44, Dec 2013.
- [12] W. Pieters and M. Davarynejad. Calculating adversarial risk from attack trees: Control strength and probabilistic attackers. In *3rd International Workshop on Quantitative Aspects in Security Assurance (QASA)*, Lecture Notes in Computer Science. Springer, 2014.
- [13] W. Pieters, S. H. G. Van der Ven, and C. W. Probst. A move in the security measurement stalemate: Elo-style ratings to quantify vulnerability. In *Proceedings of the 2012 New Security Paradigms Workshop*, NSPW '12, pages 1–14. ACM, 2012.
- [14] The Open Group. Risk taxonomy. Technical Report C081, The Open Group, 2009.
- [15] R. Verbij. Dutch e-voting opportunities: Risk assessment framework based on attacker resources. Master’s thesis, University of Twente, 2014.

Author Biographies

Wolter Pieters is the technical leader of the TRE_SPASS project at the University of Twente and an assistant professor of cyber risk at Delft University of Technology. His research interests include electronic voting, verification of security properties, and philosophy and ethics of cybersecurity. Pieters received a PhD in information security from Radboud University Nijmegen, The Netherlands. Contact him at w.pieters@tudelft.nl.

Jeroen Barendse (1973, NL) is partner and Director of LUST and LUSTlab, a multi-disciplinary design studio and research-based media & technology laboratory. He received his bachelor Graphic Design at the Academy for the Arts in Arnhem and continued at post-graduate design program Werkplaats Typografie. From 2001, Jeroen has taught Graphic Design, Typography and Interactive Design at the Academy of Arts, Arnhem and he was guest professor at the post graduate department of ERBA-Valence and Leeds Metropolitan University. Contact him at jeroen@lust.nl.

Margaret Ford is a specialist in digital identity, with extensive experience in the design, delivery and management of enterprise scale network security systems. Author of numerous research reports into state of the art Electronic Identity, with a special interest in social aspects of identity and risk management. Contact her at margaret.ford@chyp.com.

Claude Heath is a widely exhibited visual artist specialising in drawing, based at Royal Holloway, London, with a first degree in Philosophy from King's College, London. His PhD at Queen Mary, London, devised exploratory drawing methods for the study of human interaction and social practices, within the Cognitive Science Group of the Department of Electronic Engineering and Computer Science. Contact him at Claude.Heath@rhul.ac.uk.

Christian W Probst is Associate Professor in the Department of Applied Mathematics and Computer Science at the Technical University of Denmark. Christian is technical co-lead of the TRE_SPASS project. His work addresses safety and security properties of systems and organisations. Christian's group has created ExASyM, a system model that supports the identification of insider threats in organisations. Contact him at cwpr@dtu.dk.

Ruud Verbij is information security advisor at KPMG Netherlands. He works within the Information Protection Services team on a broad range of security related engagements, e.g. from performing technical penetration tests to PKI and IT audits. He has a Master's degree in information security from the University of Twente, Netherlands. Contact him at Verbij.Ruud@kpmg.nl.