



The Navigation Metaphor in Security Economics

Wolter Pieters | Delft University of Technology

Jeroen Barendse | LUST

Margaret Ford | Consult Hyperion

Claude P.R. Heath | Royal Holloway, University of London

Christian W. Probst | Technical University of Denmark

Ruud Verbij | KPMG Netherlands

By combining security architecture models (maps) with economic optimization strategies (route planning), the navigation metaphor for cybersecurity encourages strategic security investment decisions.

In the physical world, navigation is a well-understood concept: if you want to get from point A to point B, you can use a navigation system to plan your route. Such systems help optimize your behavior by finding the most efficient route and even adapt dynamically, for instance, when a particular route becomes congested. However, navigation systems might also be used for other purposes. For example, to prevent you from reaching a particular destination, an attacker could identify and then sabotage your most likely routes, significantly affecting your travel time or the likelihood of your arriving at all.

This is precisely the idea behind the *navigation metaphor* for cybersecurity being developed by the TRE₃PASS (Technology-Supported Risk Estimation by Predictive Assessment of Socio-Technical Security) project. By identifying—from an attacker’s viewpoint—the most efficient routes for gaining access to certain targets, we can optimize these routes’ defenses for the defender’s benefit. This project combines security architecture models (maps) with economic optimization strategies (route planning). The navigation metaphor,

in conjunction with risk management and visual-design insights, provides a powerful security decision support tool. Specifically, the navigation metaphor supports the economic analysis and visualization of multistep attacks: just as successive roads lead to a single physical location, ordered combinations of actions are required to reach an attacker’s goal. Using navigation as an analogy makes it easier to motivate and explain security investment decisions to a wide audience.

In this article, we introduce the navigation metaphor for cybersecurity, discussing its strengths and limitations by comparing it to the following real-world map and navigation views and steps:

- *satellite view*, which offers realistic aerial pictures but not the underlying infrastructure;
- *map view*, which shows the underlying infrastructure and enables route calculation but doesn’t show the surroundings’ actual appearance; and
- *route calculation*, which optimizes travel to a specific location as well as strategies for blocking such travel.

To illustrate our techniques, we use a case study on online services via an interactive TV interface as a running example.

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 and has supported widespread adoption of navigation systems. Whereas obtaining satellite images is relatively easy, creating a map from them is not, especially when travel is on foot or by bike. Big public roads are well-known and easy to detect in the image, but smaller, private tracks and paths might be hidden by vegetation, buildings, or clouds, or might have been forgotten and, thus, found only by chance.

In navigation systems for cybersecurity, the satellite image corresponds to the external view of the organization and its divisions, computers, and infrastructure. Some of these components are visible, others are private, and, again, others might exist but have been forgotten. The navigation metaphor for cybersecurity depends on contextual information. Acquiring such information is the first step in analyzing a system's security.

Our TRE_GPASS case study explored the delivery of home-based banking services via a TV interface. Despite its small footprint, the organization we worked with achieves significant social impact by delivering services through a range of partners in different sectors. By mapping these multiple partnerships—with their complex interdependencies and diverse security implications—we aimed to illustrate the value of using attack navigator maps and visualizations to support service planning.

We used the enterprise architecture modeling language ArchiMate to map our partner's digital services.¹ Although the stakeholders easily understood the resulting diagrams (maps), they felt that essential social, organizational, and partnership features of the system were lacking.

We next established a satellite view of the service. We conducted a prestudy briefing with senior managers, exploring the organization's goals; culture; business model; and past, present, and future projects. From these discussions a set of security concerns emerged relating to untrusted behaviors by both users and outsiders. Building on text-based target graphs developed from the prestudy results, we asked participants to consider a typical service that they work on and then to construct a representation of how they would undertake a typical data management task.

Using LEGO bricks, the participants co-constructed a rich multiperspective picture of data sharing as a part



Figure 1. Digital collage of two LEGO mapping sessions. Case study participants designing an Internet Protocol television home-banking service used color-coded LEGO bricks and figures to map the service infrastructure and the role of business actors. Yellow bricks represented central actors; green bricks, infrastructure; blue bricks, data; and pink tiles, locations. In the central loop, the service is carried forward, clockwise, starting with the client and moving to the provider and its business partners. Below, in a different loop, the banking platform supports cloud-based transactions made with a card. Above, the client receives income.

of the service.² Specifically, we asked participants to use the colors and language of ArchiMate to model the central actors (yellow bricks), infrastructure (green bricks), data (blue bricks), and locations (pink tiles; see Figure 1). The group agreed on a narrative associated with the modeling process, and the weighting and positioning (and repositioning) of the model elements.

The satellite view of the service showed regions of trust between actors, as well as dataflows crossing these regions' boundaries. LEGO avatars represented the central actors and the control strengths of selected points along data paths. Participatory techniques included mapping (the domain target), sequencing (the order of linked events), listing, placing (the relevance of previously established values), comparing (the different characters' viewpoints), and linking (the implications of actions for different actors). In a second

LEGO modeling session, the participants reflected on and remodeled the weaker parts of the service design, adding bricks where necessary. Avatars represented the nature of actors to demonstrate their business role, world view, and degree of influence, with some being diminutive and others overbearing (see Figure 1). Lines of defense were added to show relationships and areas of vulnerability.

A particular benefit of 3D physical modeling is that it's easy to incorporate annotations so that the group can keep track of any working assumptions. Apart from the use of avatars, the group developed color coding for different types of relationships and dataflows, sometimes increasing the height of defenses for greater control strength. They could view the resulting model from every angle and track each participating entity's risk implications from its own perspective. This is vitally important in smaller, more flexible organizations that rely heavily on partnerships with relational services,³ where human relationships affect the services' continuation and extension.

Maps

In the first phase, we developed the organization's satellite view as the basis for the navigator map. The goal was to obtain a precise satellite image, revealing as many elements as possible. In the second phase, we sought to translate this image into a formalized map for cybersecurity navigation.

There's a history of map-style network models in security. Trees have been 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. However, infrastructure tree models aren't always sufficiently expressive to deal with the real world. Information isn't contained within one clear boundary or perimeter, like a safe or an offline machine, but rather can be accessed via many possible routes. In such cases, a network model (a graph or map) is more suitable, although it requires more complex analysis (see, for example, the work of Paul Ammann and his colleagues⁴).

However, such network models have generally been limited to the technical parts of an organization's infrastructure, typically representing computer networks and hops of a hacker from one node to another. Focusing only on the computer network limits analysis possibilities, just as navigating only along highways limits navigation possibilities. Many attacks contain some form of social engineering or physical access; therefore, it's vital to include humans and physical locations in the map, along with their roles in obtaining access. The navigation metaphor uses such sociotechnical network models as maps (see Figure 2); these form the basis

for navigating to the goal of an attack.⁵ These maps are essentially graphs with nodes and connections, enabling route analysis. In Figure 2, you'll recognize the formalized elements from the case study, including locations (bank and home), digital infrastructure (computers and connections), and actors with possessions (Alice and Charlie).

Developing a map from the satellite view is largely based on the domain knowledge of organization members, supported by software tools that help translate this knowledge according to the rules of the map formalism. Although some information is lost in the translation from satellite image to map, the map's mathematical structure enables quantitative analyses and optimizations.

Network models allow users to model entities and relations within their access control space, independent of attacker goals. For example, rather than thinking about a particular database's vulnerabilities, users would first map the infrastructure around this database, including digital, physical, and social access relations. This decoupling of attack opportunities and modeling is crucial for ensuring that no relevant attack opportunities are overlooked. Just as geographical maps represent different entities and their connections, cybersecurity maps represent different assets and possible activities. This information forms the basis for maps that represent attack navigation in sociotechnical systems—identifying possible attacks by picking a starting state and a target node on the map.

Routes

With geographic maps, navigation entails identifying and optimizing the means to reaching a goal using different modes of transportation and routes; in the attack navigator, this amounts to identifying and optimizing ways in which attackers can reach a particular asset.⁵ This analysis is based on assumptions of economic rationality: short and cheap routes are better for attackers (long and costly ones are obviously better for defenders).

Some kind of route planning can begin with the satellite view. Once the participants have identified and agreed on the most important aspects, such as the main actors, locations, and assets, the environment can be explored from many different perspectives. In physical models, the actors can be moved around, letting stakeholders evaluate the attacker perspective and strategy in relation to future users of the service and explore different relationships' strategic and economic implications. These include the potential for financial abuse of system users and the economic risks posed by rivals and partner organizations as well as more standard technical and communications security risks.

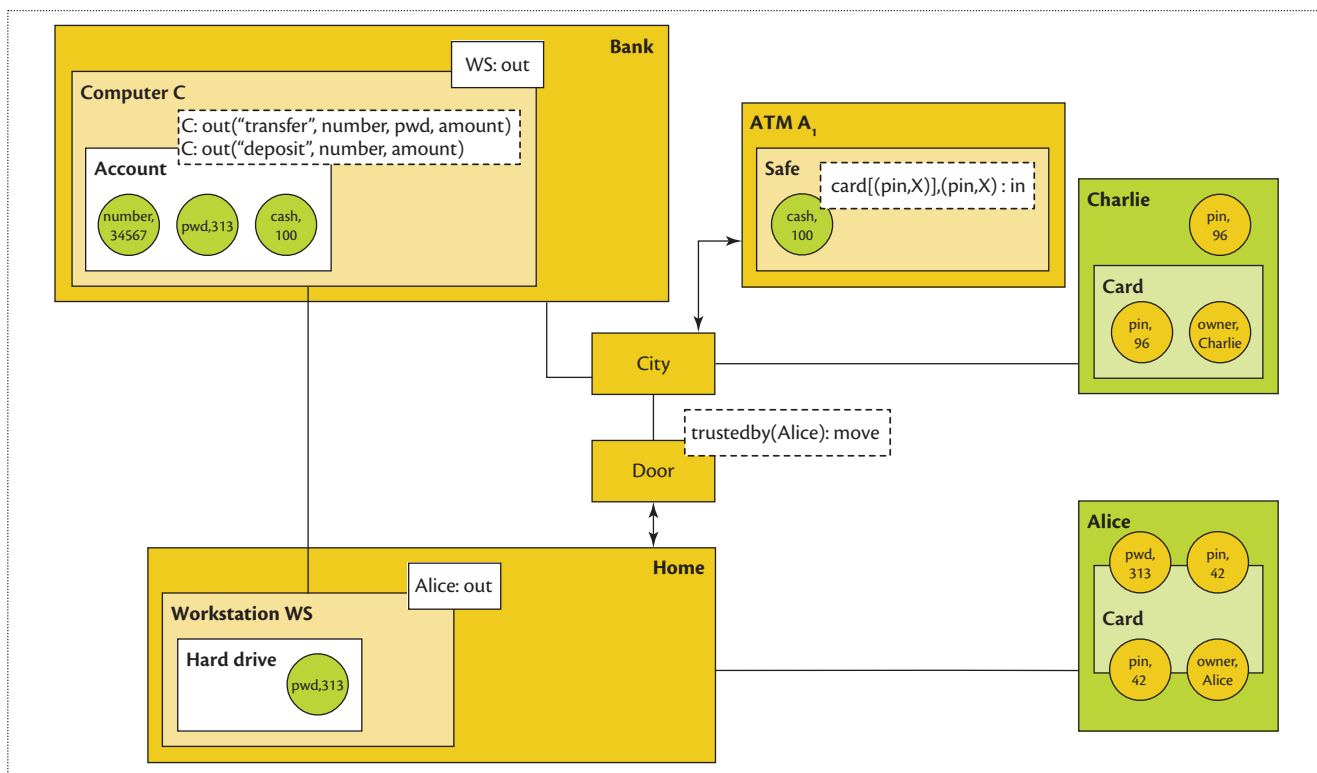


Figure 2. Example network model, or map, for sociotechnical security. Entities are represented as boxes, and data as circles. Dashed boxes represent access control policies.

More formal analysis is possible based on the map, in particular, answering the question of optimal attacker routes. This is equivalent to planning a journey by identifying the destination and generating a sequence of actions for reaching that destination. Compared to road navigation, attack navigation deals with not only where attackers can go but also what they can take along, such as credentials. Therefore, an attacker's optimal route might involve visiting different places or people to obtain their credentials 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: attackers have certain access in the beginning and might gain access to new places, people, and credentials.

However, such access comes at a price. Distance can be represented as the cost, time, or likelihood of success or failure, associated with connections and access control policies on the map. All of these constitute difficulty metrics that say something about the expected effort adversaries have to expend to gain access. For example, opening a door by force might take 3 minutes, but using a key might take only 10 seconds. In addition to time, potential options for annotation of attack steps include costs and the likelihood of detection when executing the step.

Vulnerability Functions

The route proposed by a navigation system depends on settings such as car speed and efficiency preferences as well as on infrastructure constraints such as road conditions and speed limits. An obvious example is the system recommending a shorter route that's accessible only to four-wheel drive vehicles. The infrastructure properties plus your car's properties determine how much time a particular part of the route will take you (up to "infinite" for impossible routes) and thereby also determine the optimal overall route or routes.

In security analysis, this concept corresponds to the system properties represented in the map and the attacker properties, which together determine the properties of the attack steps, such as required time and likelihood of success. This basic separation between attacker properties and system properties has been proposed in the Factor Analysis of Information Risk (FAIR) taxonomy, whereby the likelihood of an attack step's success is determined by threat capability and control strength.⁶ The attack navigator uses vulnerability functions describing a relation between the threat capability and the likelihood of success to represent the system components' resistance to attacks. In essence, this representation of resistance enables a calculation of action properties from agent properties,

similar to the way in which the outcome in many games is determined by the players' skill level plus a source of randomness (such as dice). From a more scientific point of view, item response theory and Elo ratings (for example, of chess players) achieve similar conceptual benefits.⁷

Attacker Profiles

The separation of attacker properties from system properties implies that attacker profiles are required in addition to navigator maps. Attacker profiles 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.⁸

In the TRE_SPASS project, we've 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.⁷ Additional constraints might be imposed by the attacker's available time and budget. In contrast to the skill constraints, these constraints are additive in the sense that attackers can't execute attacks for which the sum of the costs of the steps exceeds their available budget. Skill or budget levels can be represented visually by the length of a bar, a number of blocks or items, or simply a number. When attackers act, their skill must be higher than the corresponding difficulty, and their budget will decrease by the action's cost.

Although time, budget, and skill tell us something about the adversary, they don't provide information about attacker motivation or strategy. Not all attackers with the same time, budget, and skill will aim for the same attack vectors. Therefore, questions about the difficulty of attack paths should be complemented by attacker interest in such paths, typically expressed in terms of the expected utility, which in turn depends on attacker motivation. Finally, attackers might not even choose the paths with the highest expected utility, because, for example, they have limited information; this forces us to make adversary strategy (or lack thereof) explicit.¹⁰

Routing and Weakest Links

Using attacker profiles and the map, the attack navigator computes the possible routes attackers might take to reach their goal. Sets of possible attack paths can be represented as trees. Attack trees identify the different options available for attackers to achieve a goal and the properties of such attack paths, for example, likelihood of success, cost, and time.¹¹ Using extensions to

this framework, defenses can be added as well (attack-defense trees). Unlike existing attack tree frameworks, navigator maps can generate an attack tree for each combination of goal and attacker, making the analysis more flexible. Whereas traditional attack trees don't track the system components involved in attack steps, attack navigators do exactly this.

This also defines the weakest link (system component), which is determined by how much the utility for adversaries decreases when you remove the link 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 will reaching a goal become for attackers, compared to their expected gain upon reaching the goal? Rather than setting a predefined goal that attackers would be interested in, this becomes a question of evaluating which assets attackers can access with positive expected utility, and how.¹²

From the attacker's perspective, the generated attack trees also provide information that can be used to determine the optimal attack strategy, which the defender might use to determine which attack vectors are more likely and where to direct investments. By determining the most efficient routes for gaining access to certain targets from an attacker's viewpoint, we can then optimize the defenses on these routes from the defender's viewpoint.

Routes in the Internet Service

In contrast to physical navigation systems, we can use both the satellite view and map of the organization to identify possible attack routes. This process can be partially automated by using dedicated tools on navigator maps.⁵ For our 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 visualize the attacks as routes through the model. However, our project has also developed techniques for visualizing vulnerability in such attack trees based on such economic parameters as difficulty, time, cost, and probability per attack step. This approach allows navigation of specific threats by highlighting important paths, zooming in, and reordering the tree (see Figure 3), with each visualization offering a different perspective on the same scenario. Route visualizations thus communicate the attacker perspective to the defenders. On the basis of such visualizations, stakeholders can identify their system's weak links and consider making improvements by revising the architecture.

Visualizations are an essential part of navigation systems, providing a greater understanding as well as a clear means of communicating with the organization. In our

case study, they fostered a greater sense of ownership in the partner organization, which expressed a desire to retain and update the emerging LEGO model between workshop sessions. The organization's enthusiasm is typified by this quote: "I am in absolute awe at how you have managed to visualize and portray our sessions. It is very exciting to see our organization all down on paper and at the same time very challenging in terms of next steps." Participants 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 required to support sound security decision making and the commercial imperative to present key issues in a clear, concise form when dealing with senior decision makers.

Optimization

Once the attack routes are known, economic analysis can determine the attacker's optimal routes as well as the defender's roadblocks—controls to make these routes more difficult. The effects of changes, or added controls, to the map on optimal attacker strategies can be investigated, thereby providing a metric for the controls' effectiveness. This applies to digital architectures such as cloud infrastructure as well as to sociotechnical systems such as the Internet service described.

We've applied the navigation metaphor in other case studies as well. For the Estonian Internet voting system, we estimated costs of individual attack steps based on, for example, the black-market prices of infected machines. The optimal route depended on whether attackers were trying to change a single vote or the final election result. Thus, we mapped the required number of votes against the optimal strategy and the cost of this strategy for that number of votes.¹³ The navigation metaphor could also have been used to compare such results against alternative architectures, such as an equivalent paper voting system, or to specify requirements for the minimum cost for attackers to hijack one seat (or 10 seats) in the parliament.

Another case study concerned fraud in telecommunication services. Here, the maps consisted of value models of service architectures and the routes of service combinations leading to monetary gain for attackers.¹⁴ For example, rogue telecommunication operators abroad might arrange a high volume of calls to their numbers to obtain interconnection fees from honest operators. The analyses showed possible adversary strategies in terms of, for example, the number of calls required to make a profit. Again, changes to the architecture could've been investigated to reduce the utility of attackers. The various ways in which the navigation metaphor can be implemented illustrate its power in economic security analyses.

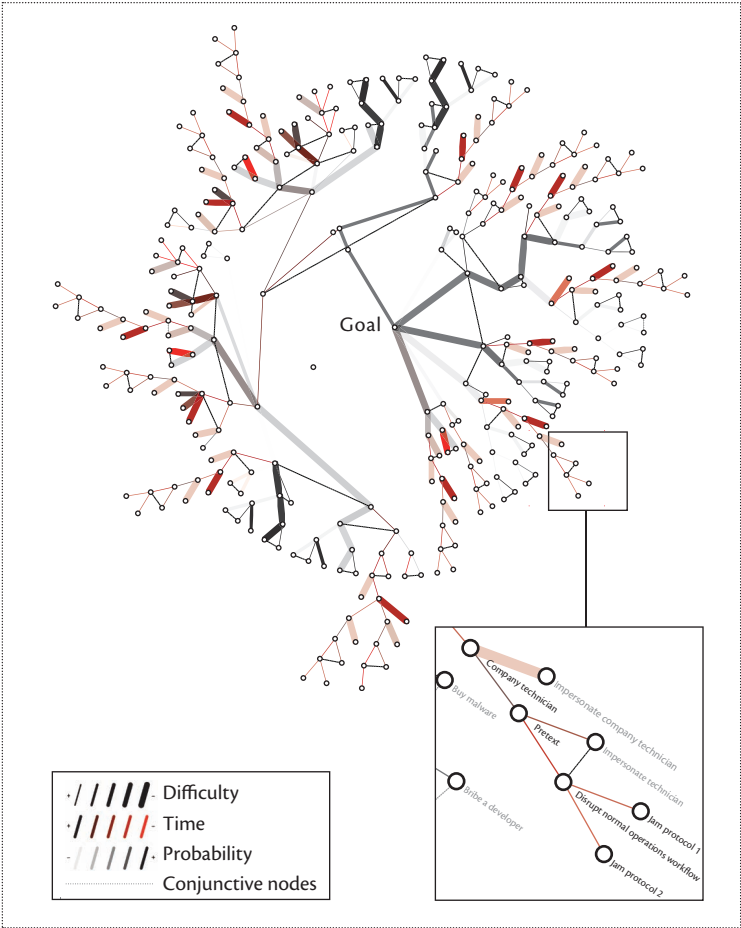


Figure 3. Radial attack tree for the Internet service case study. Each edge's color, transparency, and thickness are based on difficulty, required time, and likelihood of the corresponding attack step's success.

The TRE_{PASS} Vision

The navigation metaphor offers great potential for further integration of economic and system models in cybersecurity and visual designs. Navigation systems achieve economic optimization by shortening travelers' travel times. Attack navigators do the same for potential cyberattackers, 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 the economics of cybersecurity decisions to stakeholders. They can visualize how changing something on the map changes the situation for the attacker from an economic viewpoint. This is basically a minimax optimization supported by the map.¹⁵ The combined map (infrastructure model) and attacker profile also enables "adversary course of action"-type reasoning for different attacker types, providing flexibility under changing threat environments. If the threat environment changes,

stakeholders can simply rerun the analysis with a different attacker profile. These are the key innovations compared to the state of the art.

It's notoriously hard to compare risk assessment approaches, because their effectiveness is ultimately determined by losses due to real attacks. Even if we could measure this, it would beg the question of whether the attacks that occurred were actually representative of the system's threat environment. For example, if method A predicts an annual loss expectancy of €1 million, and method B predicts €5 million, then what's the best method if the actual average annual loss after five years is €2 million? Acknowledging these comparison difficulties, we propose that the navigation metaphor has the following main advantages over existing approaches:

- It enables a better conceptual understanding of a system's weak links, thereby supporting investment decisions.
- It encourages consideration of different attacker profiles for the same system, acknowledging different attacker economic strategies and utility functions.
- It introduces new forms of learning by inviting the user to take the perspective of an attacker looking for opportunities.
- It stimulates innovation in visualization of security economics and risk.

Obviously, the navigation metaphor doesn't solve all security risk management problems. First, there are limits to the level of detail that can be represented on a map. An attack's technical details, such as code or program flow manipulations, need different types of analyses and visualizations. Maps can still show high-level access paths for gaining the required access, for example, via different servers, infected USB drives, or social engineering. Second, the navigation metaphor doesn't solve the problem of data availability. Although our work helps stakeholders map their systems, their input might contain uncertainties or errors, which can propagate to the results. Links with security economics are important for identifying costs of actions and impact of successful attacks. Third, analysis of the maps, like many economic frameworks, depends on models of adversary choices, which can be complicated to construct and validate. We're currently investigating probabilistic models for this purpose.

The navigation metaphor is most useful in identifying attack opportunities in complex sociotechnical systems, where attacks consist of multiple steps whose connections might not be immediately obvious. The metaphor's strength lies in its representation of 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.

We've been asked whether attack navigator tools could also give attackers an advantage, particularly when the methods are published and the 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 the tools that analyze optimal attack paths.

We foresee several key challenges for the coming years. The first is to further integrate not only different types of difficulty metrics for navigation analysis, such as CVSS (Common Vulnerability Scoring System) or CWSS (Common Weakness Scoring System) values, but also the results of social engineering experiments. The second is to develop more advanced methods for considering attacker motivation and strategy during analysis, in addition to resource properties such as skill, budget, and available time. Finally, there will be a need for app-style interfaces for map development and maintenance as well as first-person visualization of routes.

We hope to help tackle these challenges in the final stages of the TRE_SPASS project and, in doing so, help stakeholders use the well-understood metaphor of navigation to support better economic reasoning and investment decision making in cybersecurity. Whereas geographic navigation is about optimizing goal attainment on earth, cyberdefense is about making such optimization harder for attackers in cyberspace. The navigation metaphor helps stakeholders grasp and rethink this fundamental economic relation between attackers and defenders. ■

Acknowledgments

The authors thank Lizzie Coles-Kemp and Peter A. Hall for their input in developing the LEGO methods. The research leading to these results was funded by 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 European Union isn't liable for any use that might be made of the information contained herein.

References

1. E. Grandry, C. Feltus, and E. Dubois, "Conceptual Integration of Enterprise Architecture Management and Security Risk Management," *Proc. 17th IEEE Int'l Enterprise Distributed Object Computing Conference Workshops* (EDOCW 13), 2013, pp. 114–123.

2. C. Heath, L. Coles-Kemp, and P. Hall, "Logical Lego? Co-constructed Perspectives on Service Design," *Proc. 10th Biannual Conf. Design and Development* (Nord-Design 14), 2014, pp. 416–425.
3. C. Cipolla and E. Manzini, "Relational Services," *Knowledge, Technology & Policy*, vol. 22, no. 1, 2009, pp. 45–50.
4. P. Ammann, D. Wijesekera, and S. Kaushik, "Scalable, Graph-Based Network Vulnerability Analysis," *Proc. 9th ACM Conf. Computer and Communications Security* (CCS 02), 2002, pp. 217–224.
5. F. Kammüller and C.W. Probst, "Invalidating Policies Using Structural Information," *Proc. IEEE Security and Privacy Workshops* (SPW 13), 2013, pp. 76–81.
6. *Risk Taxonomy*, tech. report C081, The Open Group, 2009.
7. 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," *Proc. New Security Paradigms Workshop* (NSPW 12), 2012, pp. 1–14.
8. T. Casey, P. Koeberl, and C. Vishik, "Threat Agents: A Necessary Component of Threat Analysis," *Proc. 6th Ann. Workshop Cyber Security and Information Intelligence Research* (CSIIRW 10), 2010, pp. 56:1–56:4.
9. A. Lenin, J. Willemson, and D.P. Sari, "Attacker Profiling in Quantitative Security Assessment Based on Attack Trees," *Secure IT Systems*, LNCS 8788, K. Bernsmed and S. Fischer-Hübner, eds., Springer, 2014, pp. 199–212.
10. W. Pieters and M. Davarynejad, "Calculating Adversarial Risk from Attack Trees: Control Strength and Probabilistic Attackers," *Proc. 9th Int'l Workshop Data Privacy Management, Autonomous Spontaneous Security, and Security Assurance* (DPM 14), LNCS 8872, Springer, 2014, pp. 201–215.
11. S. Mauw and M. Oostdijk, "Foundations of Attack Trees," *Proc. 8th Ann. Int'l Conf. Information Security and Cryptology* (ICISC 05), LNCS 3935, Springer, 2006, pp. 186–198.
12. W. Pieters, "Defining 'the Weakest Link': Comparative Security in Complex Systems of Systems," *Proc. IEEE 5th Int'l Conf. Cloud Computing Technology and Science* (CloudCom 13), 2013, pp. 39–44.
13. R. Verbij, "Dutch E-voting Opportunities: Risk Assessment Framework based on Attacker Resources," master's thesis, Dept. Computer Science, Univ. of Twente, 2014.
14. D. Ionita, S.K. Koenen, and R.J. Wieringa, *Modelling Telecom Fraud with e3value*, tech. report TR-CTIT-14-11, Centre for Telematics and Information Technology, Univ. of Twente, 2014.
15. L.A. Cox Jr., "Game Theory and Risk Analysis," *Risk Analysis*, vol. 29, no. 8, 2009, pp. 1062–1068.

Wolter Pieters is the former 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 social engineering, security metrics and testing, and philosophy and ethics of cybersecurity. Pieters received a PhD in information security from Radboud University Nijmegen. Contact him at w.pieters@tudelft.nl.

Jeroen Barendse is partner and director of LUST and LUSTlab. His research interests include data visualization, interaction, and interfaces. Barendse received a BA in graphic design from the Academy for the Arts. Contact him at jeroen@lust.nl.

Margaret Ford is a specialist in digital identity at Consult Hyperion. Her research interests include state-of-the-art electronic identity, particularly social aspects of identity and risk management. Ford received a postgraduate diploma in information systems from Thames Valley University. Contact her at margaret.ford@chyp.com.

Claude P.R. Heath is a Royal Holloway, University of London-based researcher and visual artist specializing in drawing. His research interests include exploratory drawing methods for the study of human interaction and social practices. Heath received a PhD in cognitive science from Queen Mary. Contact him at Claude.Heath@rhul.ac.uk.

Christian W. Probst is an associate professor in the Department of Applied Mathematics and Computer Science at the Technical University of Denmark and is technical colead of the TRE_SPASS project. His research interests include safety and security properties of systems and organizations. Probst received a PhD in engineering from Saarland University. Contact him at cwpr@dtu.dk.

Ruud Verbij is an information security advisor at KPMG Netherlands. He works on a broad range of security-related engagements, including technical penetration tests to public-key infrastructure and IT audits. Verbij received an MSc in information security from the University of Twente. Contact him at Verbij.Ruud@kpmg.nl.



Selected CS articles and columns are also available for free at <http://ComputingNow.computer.org>.