Dan Ionita

Model-driven information security risk assessment of socio-technical systems

Model-Driven Information Security Risk Assessment of Socio-Technical Systems

Dan Ioniţă

MODEL-DRIVEN INFORMATION SECURITY RISK ASSESSMENT OF SOCIO-TECHNICAL SYSTEMS

DISSERTATION

to obtain
the degree of doctor at the University of Twente,
on the authority of the rector magnificus,
prof.dr. T.T.M. Palstra,
on account of the decision of the graduation committee
to be publicly defended
on Thursday, 8th of March 2018 at 12:45PM

by

Dan Ioniţă

born on the 16th of April 1988
in Bucharest, Romania.

This dissertation is approved by:
Supervisor prof. dr. R.J. Wieringa

IDS Ph.D. Thesis Series No. 18-456
Institute on Digital Society
P.O. Box 217, 7500 AE Enschede, The Netherlands

SIKS Dissertation Series No. 2018-06
The research reported in this thesis has been carried out under the auspices of SIKS, the Dutch Research School for Information and Knowledge Systems.

This research was funded through the European Union Seventh Framework Programme (FP7/2007-2013) under grant agreement ICT-318003 (TREsPASS).

ISBN: 978-90-365-4483-2
ISSN: 1381-3617
DOI number: 10.3990/1.9789036544832
https://doi.org/10.3990/1.9789036544832

Typeset with LaTeX.
Cover and print: AIO proefschrift

Graduation committee

Chairman and Secretary:	prof. dr.	J.N. Kok	University of Twente
Supervisor:	prof. dr.	R.J. Wieringa	University of Twente
Committee members:	prof. dr.	P.H. Hartel	University of Twente
			Delft University of Technology
	prof. dr.	M.U. Reichert	Ulm University Germany
	prof. dr.	M. Junger	University of Twente
	prof. dr.	J. Zdravkovic	Stockholm University
	dr.	J. Gordijn	Vrije Universiteit Amsterdam

Question everything.

English summary

This dissertation explores the role of conceptual models in assessing the risks pertaining to the development and operation of socio-technical systems. Specifically, it introduces a variety of risk assessment techniques built around different types of conceptual models not traditionally used in risk management. They range from coordination process models to argumentation models and from tangible models to value models. The dissertation does not, however, aim at to produce an exhaustive list. Instead, it is meant to shed light on how existing conceptual modelling paradigms can support the risk assessment processes, as well as discuss the applicability of different modelling approaches to the identification or analysis of different kinds of risks.

I start by introducing a distinction between models serving as input to a risk assessment and models which are produced as a result of a risk assessment. I give examples of ontologies from the fields of enterprise modelling and argumentation which have the potential to empower analysts to better understand the system being assessed, to streamline the assessment process, to quantify risks, or to communicate results. In the remainder of the thesis, I propose several model-driven modelling and analysis approaches which can be used stand-alone but can also augment existing risk management processes. The approaches are centered around three modelling paradigms:

- Tangible modelling - i.e. "physical" modeling using graspable three-dimensional tokens - and its benefits on the collaborative effort required to construct correct and complete models of socio-technical systems. I conclude that tangible modelling can reduce the modelling effort - especially when modelling is done as a group - and that it has beneficial effects on the quality of the resulting models when the modellers have a technical background. These effects are significant if there is some relationship between the appearance of the tangible tokens and their meaning. But they are heavily mitigated by the profile of the modellers: people with a technical background produce tangible models which closely adhere to the prescribed syntax of the language while people with a background in social sciences tend to produce rich pictures.

- Argumentation modelling - i.e. recording the rationale behind claims - and how it can support the security decision making process. Results show that structuring the risk assessment as a set of arguments forces risk assessors to make their assumptions explicit and that maintaining a mapping between risks and countermeasures increases the defensibility of the resulting security requirements. Simple, informal argumentation structures provide a basis for making risk assessment more transparent, but also more collaborative.

- Value modelling - i.e. understanding the value transfers which underpin any commercial information system - and how they can be used to quantify risks, identify vulnerabilities

to fraud, and rationalize processes. I find that value models, and in particular the *e³value* modelling ontology, provide the ability to quantify risks in terms of their business impact. I show how the ontology - with a small extension - can be used to automatically generate and rank fraud scenarios. Finally, I propose an approach for extracting value models from process models which opens the door to rationalizing business processes in terms of their financial sustainability.

The three approaches are in principle complementary, as they each address different aspects of risk assessment or different types of risk.

Overall, I find that conceptual models, especially ones with a usable graphical representation, increase justifiability by making the inner workings of the risk assessment easier to understand for both the assessors and external stakeholders. Justifiability is important because risk assessment of socio-technical systems (1) often involves experts from different domains, (2) needs to inform the broader *Governance, Risk and Compliance* capabilities, and (3) should be both defensible and re-visitable.

Nederlandse samenvatting –Summary in Dutch–

Dit proefschrift onderzoekt de rol van conceptuele modellen bij het bepalen van risico's met betrekking tot de ontwikkeling en het gebruik van sociotechnische systemen. In het bijzonder introduceert het verschillende technieken voor risicoanalyse gebaseerd op verschillende conceptuele modellen die van oudsher niet in de risicoanalyse werden gebruikt. Deze modellen variëren van coördinatieproces-modellen tot argumentmodellen en van tastbare modellen tot waardemodellen. Het doel is niet om een volledige lijst met modellen te geven. In plaats daarvan is het doel om inzicht te geven over hoe bestaande paradigma's voor conceptuele modellering kunnen bijdragen aan het maken van een risicoanalyse, en om inzicht te geven in de de toepasbaarheid van de verschillende modelleertechnieken voor de identificatie of analyse van verschillende soorten risico's.

Ik begin met het vaststellen van een verschil tussen modellen die dienen als invoer voor een risicoanalyse en modellen die juist een resultaat van een risicoanalyse zijn. Ik geef voorbeelden van ontologieën uit argumentatietheorie en *enterprise modeling* die analisten de mogelijkheid kunnen geven om een beter beeld van het te analyseren systeem te krijgen, om de analyse te kunnen stroomlijnen, om risico's te kwantificeren of om resultaten te communiceren. In de rest van het proefschrift stel ik diverse model-gedreven modellerings- en analysemethodes voor, die zelfstandig gebruikt kunnen worden, maar ook als aanvulling op bestaande methoden voor risicoanalyse kunnen dienen. De methoden richten zich op drie modelleringsparadigma's:

- Tastbare modellering – materiële modellering door het gebruik maken van grijpbare, driedimensionale objecten – en de voordelen van de noodzakelijke samenwerking om juiste en complete modellen van sociotechnische systemen te maken. Ik concludeer dat tastbare modellering het modelleringsproces kan vereenvoudigen – in het bijzonder wanneer het modelleren door een groep wordt gedaan – en dat, wanneer de betrokkenen een technische achtergrond hebben, het een positieve uitwerking op de kwaliteit van het resulterende model heeft. De effecten zijn relevant als er een relatie is tussen het uiterlijk en de betekenis van de tastbare objecten. Echter worden deze effecten sterk tenietgedaan door het type persoon dat modelleert: personen met een technische achtergrond produceren tastbare modellen die nauw aansluiten bij de voorgeschreven syntaxis van de taal, terwijl personen met een achtergrond in de sociale wetenschappen vaker *rich pictures* maken die meer op cartoons lijken..

- Argumentmodellering – het vastleggen van de rationale achter beweringen – en hoe dit de besluitvorming over beveiliging kan ondersteunen. Resultaten tonen aan dat door de risicoanalyse als een verzameling argumenten te structureren, risicobeoordelaars gedwongen worden hun aannames expliciet te maken en dat het bijhouden van de relatie

tussen risico's en tegenmaatregelen, de weerbaarheid van de resulterende beveiligingsmaatregelen vergroten. Simpele, informele argumentaties maken de risicoanalyse niet alleen transparanter, maar nodigen ook uit tot meer samenwerking tijdens het proces.

- Aaardemodellering – het begrijpen van de waardeoverdracht van ieder commercieel informatie systeem – en hoe dit kan dienen om risico's te kwantificeren, fraudegevoelige scenario's kan blootleggen en processen kan rationaliseren. Ik ben van mening dat waardemodellen, in het bijzonder de *e*^{3}*value* modelleringsontologie, de mogelijkheid bieden om risico's te kwantificeren op basis van hun uitwerking op de bedrijfsvoering. Ik laat zien hoe de ontologie – met een kleine uitbreiding – kan worden gebruikt om automatisch frauduleuze scenario's te genereren en te rangschikken. Tot slot stel ik een methode voor om waardemodellen uit procesmodellen te kunnen afleiden, wat de deur opent om de financiële houdbaarheid van bedrijfsprocessen te verbeteren.

De drie methoden zijn in principe complementair, aangezien elk zich op een ander aspect van risicoanalyse richt of zich richt op een ander type risico.

Samenvattend, de resultaten laten zien dat conceptuele modellen, in het bijzonder modellen met een bruikbare grafische weergave, de rechtvaardiging van risico-analyses kunnen verbeteren door de interne structuur van die analyses zichtbaar te maken voor de verschillende partijen die bij risico-analyse betrokken zijn. Het kunnen verantwoorden is van belang aangezien de risicoanalyse van sociotechnische systemen (1) vaak de betrokkenheid van experts van verschillende vakgebieden vereisen, (2) de *Governance, Risk and Compliance* capabilities dienen te informeren en (3) zowel verdedigbaar als herzienbaar moeten zijn.

Aknowledgements

I remember leaving Romania with two suitcases and a lot of enthusiasm. It felt like going on a business trip. Now, more than six years later, I realize this trip was in actuality a truly transformative experience: I now need a truck to move all my belongings. But everything in that truck would pale in comparison to the experiences I've lived and the knowledge I've gathered throughout my stay in The Netherlands. And I owe each and every one of them to the amazing people I've met throughout this journey.

First and foremost, I'd like to express my gratitude towards Roel, who went from being my teacher, to my thesis supervisor, to my doctoral promoter, and now my business partner. Thank you for being my mentor. Thank you for guiding me throughout my journey into academia. Working with you has been an honor and a privilege. I'm also grateful towards my graduation committee for taking the time to review my work and provide valuable feedback.

The research leading up to this dissertation was not conducted in isolation. It's the result of hundreds of hours of brainstorming, performing experiments, and sending comments back and forth with some of the smartest people I've had the privilege of meeting: Alexandr, Jaap, Julia, Lorena, Margaret, and Wolter.

Even with all their help, I probably wouldn't have made it past my first week without the endless amount of support I received from Suse, Bertine, and Gert-Jan. As far as I am concerned, you guys really are the beating heart of the group.

I'm very grateful for having some of the best collegues I could ask for. Ali, Alexandr, Andreas, Bence, Chris, Erik, Ines, Elmer, Elefteria, Herson, Jan-Willem, Marco, Prince, Riccardo, Robson, Roeland, Steven, Susanne, Thijs, Tim, and Yuxi, thank you for making lunches and coffee-breaks the highlight of my day.

The reason this adventure was even possible is also the reason I regret every day of it: my family. Mom, dad, I know I don't say this nearly enough but I love you! No amount of acknowledgments could come close to expressing my appreciation for making me who I am.

Laura... nothing I say here would do justice to the spectacular, intimate, and crazy moments we spent together. Thanks for being for being both dedicated and understanding. Both fun and beautiful. Both loving and lovable. Trr.

Alex, Antonia, Aykan, Christos, Cristi, Dirk, Dimitris, Gaby, Ivana, Kiril, Kostacos, Mircea, Moustafa, Razvan, Robert, Vassilis, and Vincy, thank you for putting up with me. The countless awesome moments we spent together helped me stay sane.

And finally, special thanks to Alex, Costin, Emy, George, Narcis, Nicu, Paco, Puiu, and Vlad for making sure trips back home were both frequent and fun. Thanks for keeping me close even when I was far away.

Rotterdam, March 2018
Dan Ioniţă

Table of Contents

List of Figures

List of Tables

List of Acronyms

A

ACE	Adaptive Communication Environment
ADSL	Asymmetric Digital Subscriber Line
AI	Artificial Intelligence
ASAP	As Soon As Possible

D

BPEL	Business Process Execution Language
BPMN	Business Process Modelling Notation

C

CAE	Claims, Arguments and Evidence
COBIT	Control Objectives for Information and Related Technology
CORAS	Control Objectives for Information and Related Technology
CV	Coefficient of Variation

D

DARPRA	Defense Advanced Research Projects Agency
DoS	Denial of Service
DDoS	Distributed Denial of Service

E

EA	Enterprise Architecture
EM	Enterprise Modelling
EU	European Union

F

FAIR	Factor Analysis of Information Risk

G

GDPR	General Data Protection Regulation
GRC	Governance, Risk, and Compliance
GSN	Goal Structuring Notation

I

IFIP	International Federation for Information Processing
IP	Internet Protocol
IS	Information Security
ISO	International Standards Association
ISP	Internet Service Provider
ISRM	Information Security Risk Management
IT	Information Technology
IEEE	Institute of Electrical and Electronics Engineers

N

NFC	Near Field Communication

P

PBX	Private Branch Exchange
PISA	Personal Information Security Assistant

Q

QoS	Quality of Service

R

RA	Risk Assessment
REA	Resource Event Agent
RISA	RIsk assessment in Security Argumentation
RM	Risk Management
ROI	Return on Investment
ROSI	Return on Security Investment
RSF	Revenue Sharing Fraud

S

SLA	Service Level Agreement
SME	Small or Medium Enterprise
SRA	Structured Risk Analysis

T

TREsPASS	Technology-supported Risk Estimation by Predictive Assessment of Socio-technical Security
TSP	Telecommunication Services Provider

U

UML	Unified Modelling Language

V

VM	Virtual Machine

Part I

Introduction and Background

1 Introduction

As more aspects of life transition to the digital domain, computer systems become increasingly complex but also more social. This opens up plenty of opportunities, but also brings about new risks. In the face of major leaks and well-publicized hacking incidents, companies are facing increasing pressure to improve their security. But assessing a socio-technical system is no trivial task: it often requires intimate knowledge of the system, awareness of the social dynamics and trust relationships of its users, a deep understanding of both hardware and software, as well as the ability to quantify risks, communicate security policies and engage stakeholders. Conceptual models, as tools designed to help make sense of complex issues, can help with some of these problems. In this first Chapter, I summarize several problems often encountered in the risk assessment of socio technical systems and sketch a model-driven solution direction, to be fleshed out in the remainder of the thesis. I also list the individual publications that culminated in this thesis and highlight the societal and theoretical evidence of this research.

1.1 The problem context

Information security risks are risks associated with the use of IT. Thanks to the ubiquity and pervasiveness of computers in modern society, these are are quickly becoming the most prevalent type of risk. As social interactions, business and identities move to the digital domain, old ways of understanding and mitigating information security risk need to be re-thought. The increased availability of hacking tools and tutorials makes cyber-attacks and cyber-fraud easier to perform, while the anonymity provided by the Internet means cyber-criminals are much harder to catch. The computational power of modern computers and the interconnected nature of IT systems open up possibilities for attacks on unprecedented scales. These factors led to cyber-crime related losses of roughly $400 billion in 2014 [4]. These are expected to rise to around $2 Trillion per Year: by 2019 [5].

Even companies whose products or services are physical par excellence rely on IT systems for things like accounting, marketing, and customer or enterprise management. In order for organizations to properly manage risks, they first need to assess them. But information security risk assessment is a complex process for several reasons. First, it requires domain-specific knowledge as well as intimate knowledge about the system and its operation. Second, it involves decision-making based on incomplete information and often unquantifiable return on investment. Third, it tries to capture a snapshot of a moving landscape, with new vulnerabilities being discovered weekly. Fourth, the results need to feed into existing enterprise processes and communicated back to stakeholders.

All of this is aggravated by the fact that most information systems are embedded in larger socio-technical systems in which users become attack vectors. Some form of social engineering (i.e. manipulating people) is thought to have been used in two-thirds of attacks by hackers, activists and nation states [6]. A 2017 survey by the Business Continuity Institute found that phishing and social engineering remain the top driver of cyber-disruption to organizations [7]. Therefore, to obtain a complete overview of the risk landscape surrounding the development, implementation, operation, and maintenance of an information systems, it is helpful to view its stakeholders and IT components together as a *socio-technical system* [8–11]. Socio-technical systems (STS) theory recognizes the interplay between people and technology, thereby supporting the identification and analysis of a wider variety of information security risks – such as social engineering, procedural, and fraud risks – but also raises challenges in modeling the complete system.

1.2 Research goal

The goal of this research is to *improve information security risk assessment in a model-driven way without unnecessary quantification*.

This goal can be decomposed into several Research Questions (RQs):

RQ1 How can the effort and resources required to perform an IS risk assessment be reduced?

RQ2 How can the defensibility, understandability, and re-usability of risk mitigation deci-

sions be improved?

RQ3 How can IS risk assessments be better integrated with established enterprise processes?

1.3 Motivation and relevance

Due to the diversity of information technology and dynamic nature of risk, there is no one-size-fits-all Information Security Risk Assessment (ISRA) method. The variety of information systems, from consumer applications to cloud infrastructures, means there is a wide variety of (potentially conflicting) requirements for risk analysis methodologies and tools. This results in a large ecosystem of mostly high-level guidelines. In order to be operationalized, these guidelines need to be interpreted and contextualized, which usually requires expert knowledge. However, many enterprises are not willing or able to hire such experts and their employees might not have the skills or knowledge required by certain risk assessment frameworks. Furthermore, data required by many generic risk assessment methodologies - such as likelihood or impact estimations - might not be quantifiable, for example because the events are very rare. Finally, most modern information systems are in fact socio-technical systems, which introduces new attack vectors and new perspectives to consider. Consequently, there is a need for lightweight, qualitative, and flexible information security risk assessment methods.

Conceptual models are extensively used in computer science to describe, explain and understand complex software and systems. Therefore, conceptual models play an implicit role in IT risk management and risk assessment activities. In some risk assessment methodologies, such as CORAS [12], conceptual models play a central role. But the majority of risk assessment methods does not come with pre-defined modelling languages and many do not mandate the use of models at all. However, conceptual models exhibit several desire-able features:

- Models abstract away unnecessary details
- Models can be represented visually
- Formal models support automation
- Informal models can handle qualitative data.

Socio-technical models which are easy to construct can make it easier for domain experts and stakeholders to construct accurate models of their organization, without being modelling experts. Intuitive, understandable models could then serve to better inform risk assessment processes.

Improving the defensibility and understandability of risk mitigation decisions is critical when these decisions have to be explained, for instance when requesting resources for counter-measures, when trying to show compliance, or in the aftermath of a cyber-attack. In addition, since many security mechanisms pertaining to socio-technical systems are in fact policies that have to be communicated and understood to be effective, the ability to convey the rationale behind them has the potential to increase awareness.

Conducting the risk assessment based on established enterprise modelling paradigms, such as process modelling, value modelling, or architecture modelling has the potential to reduce the effort required to obtain models of the Target of Assessment (ToA) as well as increase their understandability. In addition, it allows us to map risks and countermeasures back to these models, thereby feeding the results of the assessment back into the governance, risk and compliance workflow.

1.4 Research methodology

This thesis will, therefore, investigate (1) how to streamline the socio-technical modelling process required to model an organisation for the purpose of information security risk assessment, (2) how argumentation models can be used to support risk assessment, as well as communicate its results and (3) how a risk assessment can be conducted based on established enterprise modelling paradigms. These three topics relatively broad and differ in terms of scope, domain, and applicable research methods. Therefore, rather than designing an over-arching research methodology, I first investigate each of the three topics in isolation, using a variety of research methods. For each topic, I suggest one or more solution directions which I validate by means of case studies, experiments, surveys, or technical action research as applicable. Each solution direction is concretized in a separate Chapter which also describes and motivates the respective research methodology.

1.5 Thesis outline

First, in Chapter 2 I dive deeper into how specific modelling paradigms, both from the field of security and from other fields can inform, augment and extend socio-technical risk assessment. In Chapter 3 I discuss a series of experiments aimed at investigating whether using physical tokens to construct so-called "tangible models" of socio-technical systems can make the modelling process easier and more collaborative. In Chapter 4 I look at whether argumentation models can support the risk assessment process by encoding the rationale behind security decisions. In Chapter 5 I show how argumentation models capable of maintaining a living overview of risks and mitigation provide support for collaborative risk assessment. In Chapter 6 I introduce an extension to the *e*^{3}*value* modelling ontology which empowers analysts to quantify risks in terms of their business impact. In Chapter 7 I show how this extension can be used to generate and rank business risks, such as the risk of fraud. In Chapter 8 I present a method for deriving a value model from a process model and show how the resulting mapping can be used to identify potential sustainability issues of service delivery processes. Finally, in Chapter 9 I draw conclusions with regard to the role of tangible models, argumentation models, and value models in information security risk assessment.

1.6 Publications

This section lists work published by the author during his doctoral research (2013-2017). Fig. 1.1 positions the various scientific papers with regard to research discussed in this dissertation. The list below contains all publications grouped by type but only the ones relevant to the topic of this dissertation are included in Fig. 1.1.

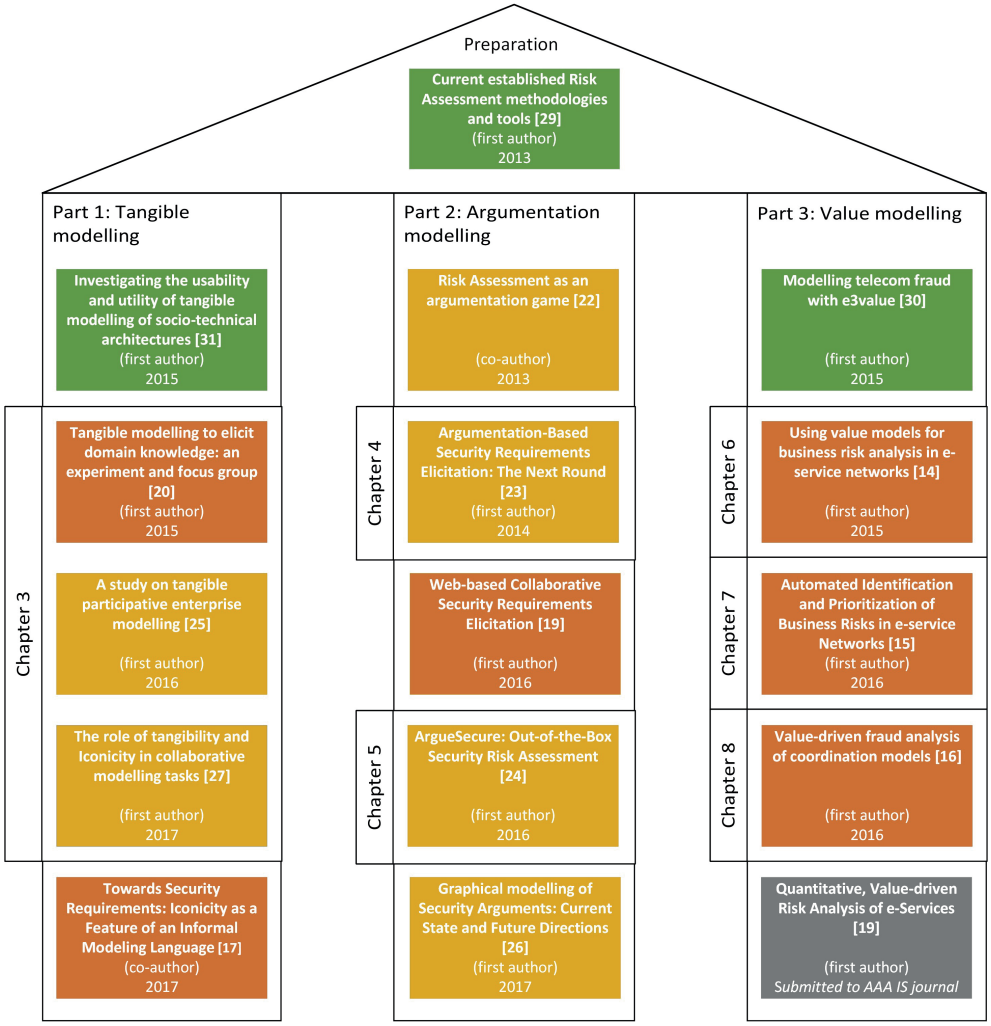

Figure 1.1: Overview of publications relevant to this dissertation (technical reports in green, workshop papers in orange, conference papers in red and journal articles in gray)

Publications in international journals

[13] Quantitative, Value-driven Risk Analysis of e-Services
Authors: D Ionita, J Gordijn, A Yesuf, RJ Wieringa
Venue: American Accounting Association's Journal of Information Systems [submitted]
h-index: 19
Year: 2017/2018

Publications in international conferences

Full papers

[14] Using value models for business risk analysis in e-service networks
Authors: D Ionita, RJ Wieringa, L Wolos, J Gordijn, W Pieters
Venue: IFIP Working Conference on The Practice of Enterprise Modeling (PoEM)
h5-index: 10
Year: 2015

[15] Automated identification and prioritization of business risks in e-service networks
Authors: D Ionita, RJ Wieringa, J Gordijn
Venue: International Conference on Exploring Services Science (IESS)
h5-index: 10
Year: 2016

[16] Value-driven risk analysis of coordination models
Authors: D Ionita, J Gordijn, AS Yesuf, R Wieringa
Venue: IFIP Working Conference on The Practice of Enterprise Modeling (PoEM)
h5-index: 10
Year: 2016

[17] Towards security requirements: Iconicity as a feature of an informal modeling language
Authors: A Vasenev, D Ionita, T Zoppi, A Ceccarelli, R Wieringa
Venue: 22nd International Working Conference on Requirements Engineering: Foundation for Software Quality (REFSQ)
h5-index: 16
Year: 2017

[18] Threat navigator: grouping and ranking malicious external threats to current and future urban smart grids
Authors: A Vasenev, L Montoya, A Ceccarelli, A Le, D Ionita
Venue: First International Conference on Smart Grid Inspired Future Technologies: (SmartGIFT)
h5-index: N/A
Year: 2017

Short papers

[19] Web-based Collaborative Security Requirements Elicitation.
Authors: D Ionita, R Wieringa
Venue: 24th International Working Conference on Requirements Engineering: Foundation for Software Quality (REFSQ)
h5-index: 16
Year: 2016

[20] Tangible Modelling to Elicit Domain Knowledge: An Experiment and Focus Group
Authors: D Ionita, R Wieringa, JW Bullee, A Vasenev
Venue: 34th International Conference on Conceptual Modelling (ER)
h5-index: 16
Year: 2015

[21] Outlining an "Evaluation continuum": Structuring evaluation methodologies for infrastructure-related decision making tools
Authors: A Vasenev, L Montoya, D Ionita
Venue: First International Conference on Smart Grid Inspired Future Technologies: (SmartGIFT)
h5-index: N/A
Year: 2017

Publications in international workshops

[22] Risk assessment as an argumentation game
Authors: H Prakken, D Ionita, R Wieringa
Venue: International Workshop on Computational Logic in Multi-Agent Systems (CLIMA)
Part of: 12th International Conference on Logic Programming and Nonmonotonic Reasoning (LPNMR)
Year: 2013

[23] Argumentation-Based Security Requirements Elicitation: The Next Round
Authors: D Ionita, JW Bullee, RJ Wieringa
Venue: IEEE 1st International Workshop on Evolving Security and Privacy Requirements Engineering (ESPRE)
Part of: 22nd IEEE International Requirements Engineering Conference (RE)
Year: 2014

[24] ArgueSecure: Out-of-the-Box Security Risk Assessment
Authors: D Ionita, R Kegel, A Baltuta, R Wieringa

Venue: IEEE 3rd International Workshop on Evolving Security and Privacy Requirements Engineering (ESPRE)
Part of: 24th IEEE International Requirements Engineering Conference (RE)
Year: 2016

[25] A study on tangible participative enterprise modelling
Authors: D Ionita, J Kaidalova, A Vasenev, R Wieringa
Venue: 3rd International Workshop on Conceptual Modeling in Requirements and Business Analysis (MReBA)
Part of: 35th International Conference on Conceptual Modeling (ER)
Year: 2016

[26] Graphical modeling of Security Arguments: Current State and Future Directions
Authors: D Ionita, M Ford, A Vasenev, R Wieringa
Venue: The Fourth International Workshop on Graphical Models for Security (GraMSec)
Part of: 30th IEEE Computer Security Foundations Symposium (CSF)
Year: 2017

[27] The role of tangibility and iconicity in collaborative modelling tasks
Authors: D Ionita, D Nazareth, A Vasenev, F van der Velde
Venue: ER Forum on Conceptual Modelling: Research in Progress
Part of: 36th International Conference on Conceptual Modeling (ER)
Year: 2017

Doctoral symposiums

[28] Context-sensitive Information security Risk identification and evaluation techniques
Authors: D Ionita
Venue: 22nd IEEE International Requirements Engineering Conference (RE)
h5-index: 23
Year: 2014

Technical reports

[29] Current established risk assessment methodologies and tools
Authors: D Ionita, PH Hartel, W Pieters, R Wieringa
Year: 2013
Cited in: N/A

[30] Modelling telecom fraud with e3value
Authors: D Ionita, SK Koenen, RJ Wieringa

Year: 2014

[31] Investigating the usability and utility of tangible modelling of socio-technical architectures
D Ionita, R Wieringa, JW Bullee, A Vasenev Authors:
D Ionita, R Wieringa, JW Bullee, A Vasenev
Year: 2015

1.7 Summary of contributions

The core contributions of this work consist in the methodological application of several modelling paradigms to socio-technical information security risk assessment. The resulting observations are useful for developing more powerful risk assessment frameworks in the future. The proposed tools, all documented, freely available and open-sourced can already be used to supplement or complement risk assessment efforts.

Theoretical contributions include: additions to the body of knowledge pertaining to group modelling behavior grounded in cognitive theories (Chapter 3), several conceptual models of risk argumentation (Chapters 4 and 5), and the e^3*fraud* ontological extension (Chapter 6) with its associated risk analysis approaches (Chapter 7).

2 Background

This chapter summarizes previous work relevant to the topic of model-driven risk assessment. First, it introduces some unique challenges that stakeholders face when assessing the risks pertaining to a socio-technical system. Then, it discusses the types of models current risk assessment methodologies make use of: Target of Assessment models used to inform the assessment, and models of risks used to encode the results of the assessment. Finally, it introduces several modeling frameworks not designed explicitly for risk assessment, on which the novel techniques proposed in the following chapters build upon.

2.1 Information security risk assessment (ISRA)

Information Security Risk Assessment is concerned with the *identification*, *analysis*, and *evaluation* of risks that the owner, operator or user of an information system or a piece of software might face, as well as with the identification of *risk mitigations* [32]. A ISRA is usually preceded by *context establishment*, whose aim is to define the goals, and scope of the assessment as well as to gain sufficient understanding of Target of Assessment (i.e. the parts and aspects of the system that are the subject of the risk assessment) [33]. Therefore, the model of the Target of Assessment, whether a mental model or a formal one, serves as input to the risk assessment process while its output consists of an overview of relevant risks and applicable mitigations. The result of a risk assessment is a ranked list of risks, potentially accompanied by a respective list of possible mitigations. Risks are often inter-related and they need to be operationalized in terms of vulnerabilities and quantified in terms of their business impact.

Conceptual models are compositions of concepts and relationships used to help people know, understand or simulate a subject the model represents. To this end, they are often used by teachers, designers, scientists, and engineers to provide accurate, consistent and complete representations of a target system [34]. Conceptual models may be physical objects or diagrams, but most often rely on mental models constructed via a process of conceptualization and generalization. In this respect, conceptual models are abstractions of real world systems, processes or states of affairs. They can therefore play an important role in assessing risks: any risk assessment is based on a conceptual model of the target of assessment and aims to produce a conceptual model of its risk landscape.

Besides target of assessment models and risk models, other types of conceptual models may also play a role in assessing information risk. For instance, process models which describe the behavior of the ToA or how users interact with it might help in revealing new types of exploiting the system, such as by means of social engineering, or by exploiting the order of activities. Value models which describe revenue flows may be useful to assess vulnerability to fraud or to quantify the business impact of specific risks. Finally, argumentation models which describe the rationale a claim can support the risk assessment process by formalizing the rationale behind security decisions, thereby increasing their defensibility, informing future decisions and helping show compliance.

2.2 Conceptual models used in ISRA

With regard to the risk assessment process, two broad categories of models can be identified: models of the Target of Assessment serving as input and models of risks produced as output. In this section, I describe several different modelling paradigms previously used to describe either the input to a risk assessment or the output. For each paradigm, I zoom in on one or more specific modelling languages.

2.2.1 Target of Assessment models (input)

In Chapter 3 of this dissertation, I investigate factors which may help streamline the construction of ToA models. To this end, I select different modelling languages used to model information systems. In order to control for possible effects of the language and its domain, I select languages from three different fields: architectural models from engineering, socio-technical models from computer science and enterprise models from management sciences. Since the goal is not to compare languages, but to see how the treatments proposed affect a given language, the selection is based solely on familiarity.

2.2.1.1 Architectural models

Architectural models describe the physical or digital architecture of a software or system and are therefore the most common models used to perform a cyber-risk assessment. Examples include class diagrams, network diagrams, wiring diagrams and building blueprints. Architectural models have the advantage of being well known and extensively used in the development and management of software and IT systems. Therefore, they are well understood and often readily available.

However, architectural models leave out the social layer, for example roles, relationships and individual profiles. Considering that social engineering plays an increasingly large role in successful cyber-attacks (two thirds according to a recent survey [6]), architectural models have limited utility in security risk assessment and often need to be complemented with knowledge about the individuals involved in the deployment, usage, and maintenance of the target of assessment.

IRENE

IRENE is a architectural model-driven risk assessment technique for smart grids. The method comes with its own modelling language, designed to be used in stakeholder workshops in order to collaboratively create a model of the Target of Assessment. It is therefore intended to be usable by nontechnical domain experts. I used this language in my collaborative ToA modelling experiments described in Chapter 3.

2.2.1.2 Socio-technical models

Most IT systems are in fact socio-technical systems. This is because humans are involved in the development, usage, and maintenance of the system. From a risk perspective, humans provide new attack vectors [35, 36]. Social engineering (i.e. the psychological manipulation of people into performing an action or divulging confidential information) is increasingly used to undermine information security technology [37, 38]. Therefore, risk assessment methods have started to consider the human factor [39–41]. To achieve this, the social layer and the technical layer have to be well defined and linked. Socio-technical models attempt to represent both the social layer and the technical layer in an integrated model and are therefore a natural fit [8, 42]. CORAS was among the first risk analysis techniques to define a specialized UML-based

socio-technical modelling language to describe the ToA [43, 44]. The CORAS modelling language is also is discussed in more detail in Sect. 2.2.2 later in this chapter.

TREsPASS

TREsPASS (Technology-supported Risk Estimation by Predictive Assessment of Socio-technical Security) was an EU funded project aiming to develop more or less automated model-driven risk assessment tools. One of the results of the project was a modelling language capable of representing architectural (both physical and digital), as well as social aspects of the Target of Assessment in a single model. The model was intended to be detailed enough to support a thorough risk assessment of cyber-risks, but also physical risks, such as the risk of breaking-and-entry. One of the main limitations of the TREsPASS approach is gathering the data required to construct the model. In this dissertation, I attempt to mitigate these effects using the collaborative ToA modelling approach presented in Chapter 3.

2.2.1.3 Enterprise models

An Enterprise Architecture (EA) consists of various aspects of an enterprise (e.g., a private company, government department, academic institution, other kind of organization, or part thereof). Enterprise modelling (EM) is the coherent description of these aspects, required to enable communication among stakeholders and guide any kind of transformation processes [45]. Enterprise modelling languages are therefore able to represent things such as business processes, business rules, concepts, information, data, vision, goals, and actors that make up an EA [46]. In short, an enterprise model is a "representation of the structure, activities, processes, information, resources, people, behavior, goals, and constraints of a business, government, or other enterprise" [47]. Since enterprise models provide insight into an organisation's structure, processes, and underlying IT, they can be used as a basis for security risk assessment [44, 48].

Several enterprise model-driven risk assessment techniques exist. Most notably, the Zachmann Framework [49] was used by many researchers as a basis for security engineering [50–52]. The German IT Baseline protection manual relies on assessing the IT infrastructure together with relevant organizational aspects [53]. Suh and Han use a business model to identify security requirements on information assets depending on their business function [54].

4EM

The 4EM methodology consists of an EM language, as well as guidelines regarding the EM process and recommendations for involving stakeholders in moderated workshops [55]. 4EM sub-models include Goals, Business Rules, Concepts, Business Process, Actors and Resources and Technical Components and Requirements models and are usually constructed by involving various stakeholders into moderated modelling workshops. In this dissertation, I use 4EM to validate the collaborative ToA modelling approach described in Chapter 3.

2.2.2 Models of risk (output)

In Sect. 2.2.1 above, I discussed modelling the target of assessment. But from a risk analysis perspective, modelling what can go wrong is far more important. Models of risk essentially

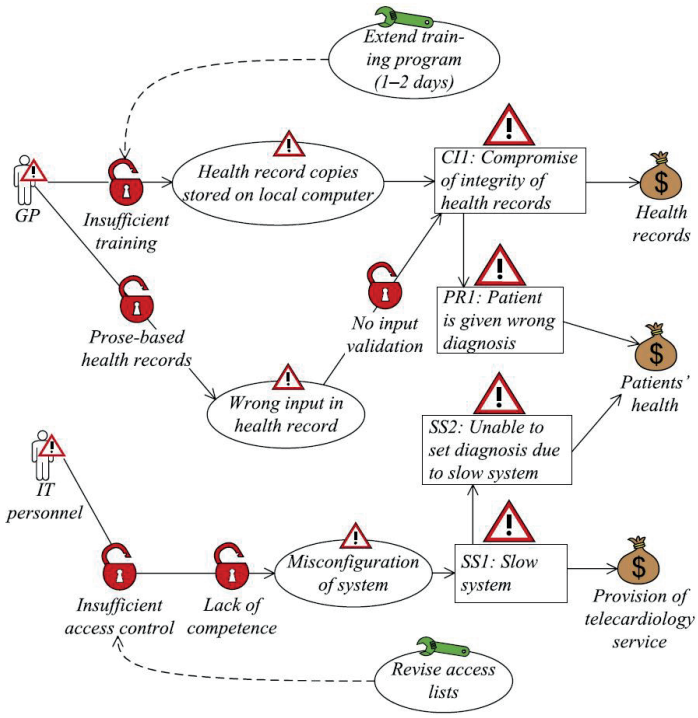

Figure 2.1: Example of a CORAS "treatment diagram". Source: [1]

formalize the results of a risk assessment. They need to paint a complete, correct and understandable picture of vulnerabilities and risks, as well as to provide actionable risk mitigation advice. Models of risk and countermeasures may even serve as assurance [56, 57] or proofs of compliance [58].

In practice, risk assessments usually aim to produce ranked lists of risks [59]. Recently, techniques drawing from goal modelling and safety risk analysis have been proposed to better structure these lists. I discuss two prominent ones below.

CORAS

CORAS is a model-driven risk analysis methodology. It defines its own UML-based modelling language, able to construct "asset" diagrams, "threat" diagrams, "risk" diagrams and "treatment" diagrams. Asset diagrams describe the target of assessment, but also help with estimating the impact of risks identified later on. Threat diagrams support risk identification and likelihood estimation by exploring the attacker's perspective Risk diagrams builds upon the asset and threat diagrams in order to present an overall risk picture. Finally, treatment diagrams enrich the risk diagram with risk mitigation possibilities for risks deemed unacceptable. An example of a treatment diagram for electronic medical records is shown in Fig. 2.1: the open locks represent vulnerabilities, the exclamation marks are threat scenarios or risks, the green wrenches represent mitigations and the $ bags are assets.

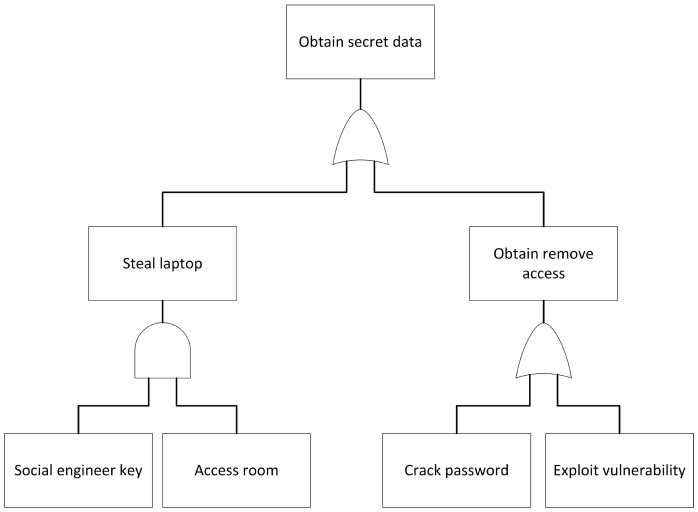

Figure 2.2: A example of an attack tree. Source: [2]

Attack Trees

Attack trees are a formal risk modelling approach which iteratively decomposes risks into combinations of atomic actions or events that have to occur in order for the risk to materialize. See Fig. 2.2 for an example. The approach is inspired by fault trees which similarly decomposed failures into series of lower-level events. Events (or actions in the case of attack trees) are composed using AND/OR gates. Attaching probabilities to these gates allows analysts to assess the total overall risk level, but also to perform root cause analysis in case a failure (or attack) occurs. However, these quantitative analyses require accurate data on likelihoods of leaf nodes. While in the case of safety these values can be obtained from historical data or by sample testing, there is no way to obtain accurate predictions of the frequency of attacks. This is because the motivation of attackers can change, but also because new vulnerabilities are discovered almost every day. These are known as zero-days and once they are disclosed publicly, the volume of attacks can increase by 5 orders of magnitude [60].

2.3 Other conceptual models potentially relevant for risk assessment

In Chapters Chapter 4 through Chapter 8 of this dissertation, I introduce several novel risk assessment techniques which make use of conceptual modelling paradigms not traditionally used in risk assessment. This section provides some background into these paradigms. For each one, I zoom in on one or more modelling languages, the choice of which is presented in the respective session. Later in the thesis, I will use some of these specific languages to demonstrate the proposed model-driven risk analysis techniques.

2.3.1 Business process models

Business process models describe how a business works, in terms of sequences of activities executed by specific business units or organizations. A single process model shows how a business accomplishes a mission, activity or task; many process models are required to fully describe the inner workings of most real-world organizations [61]. Even a single process can be quite complex, involving multiple people, groups, and systems performing a variety of tasks, either in parallel or sequentially. Sometimes, tasks are repeated, and many business processes include points where decisions which affect the flow have to be taken. Moreover, the process has to react to events and sometimes coordinate with other processes or systems.

There exist a large variety of techniques to document processes, ranging from flowcharts to Gantt charts and from Data Flow Diagrams to UML. For business process modelling, two established notations currently stand out: The *Business Process Model and Notation (BPMN)* and the *Business Process Execution Language (BPEL)*. The BPMN notation [5], is designed to appeal to technical users while being understandable to business users as well. BPEL [62], on the other hand, is mainly targeted at web service developers and lacks a standard graphical notation. Several approaches for translating between BPMN and BPEL have been proposed [63–65], but they have mainly served to expose fundamental differences between BPMN and BPEL [66,67]. I use BPMN in this dissertation because of its standardized notation and because it is the most used in practice.

BPMN
contains four types of elements. I briefly explain each element below, based on the example of Fig. 2.3:

Flow objects are the main components of a BPMN diagram:

- **Event:** An event is represented by a circle and denotes that something happens. The icon in the circle denotes the type of event: *start events* ("Goods to ship" in Fig. 2.3), *intermediate events* or *end events* ("Goods available for pick-up" in Fig. 2.3). Events and can be further specified as type *catching* or *throwing*.
- **Activity:** An activity is represented by a rectangle with rounded corners and denotes something that must be done. There a total of eight activities in the diagram of Fig. 2.3.
- **Gateway:** An activity is represented by a diamond shape and is used to fork or merge paths. In Fig. 2.3, the only labeled gateway is "Mode of delivery". The two gateways with a circle inside are inclusive (i.e. OR), while the two with a plus sign are exclusive (i.e. AND).

Connecting objects show relationships between components in a BPMN diagram:

- **Sequence flow:** A Sequence flow is represented by a solid arrow, and simply shows the order in which activities are to be performed.
- **Message flow:** A message flow is represented by a dashed line and shows the message being exchanged by actors or departments. There are no message flows in Fig. 2.3.

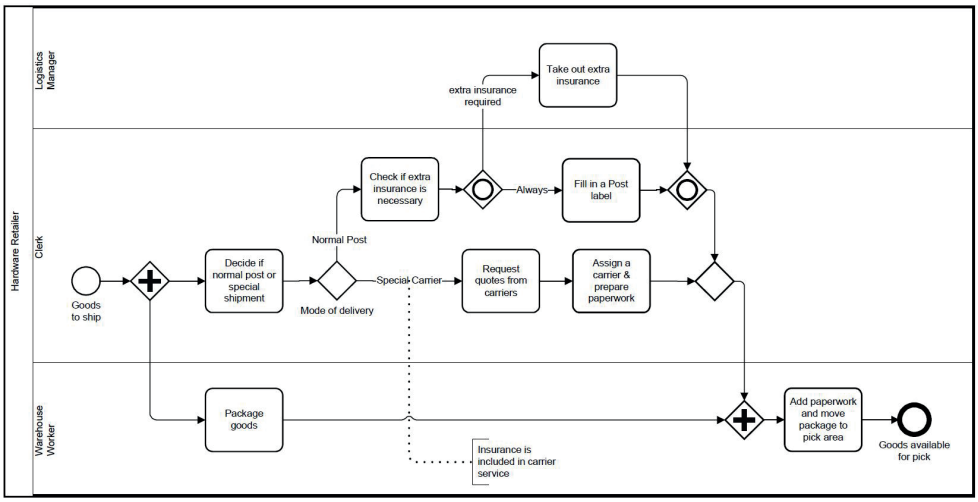

Figure 2.3: Simple BPMN model. Source: [3])

Association: An association is represented by a dotted line and is used to associate an Artifact to a Flow Object. In Fig. 2.3, the "Insurance is included in carrier service" Artifact is associated with the "Special carrier" sequence flow.

Swim Lanes are visual mechanisms of organizing and categorizing activities:

Pool: A pool represents a major participant in the process, which can be further decomposed in components (i.e. lanes), such as departments, roles or individuals. The diagram of Fig. 2.3 contains a single pool: "Hardware retailer.

Lane: A lane represents an individual actor, function or role and is depicted as a rectangle stretching the width and height of the pool. There are three lanes in Fig. 2.3: "Warehouse Worker", "Clerk" and "Logistics Manager"

Artifacts allow data and information to be included in a BPMN diagram:

Data object: Data objects represent data that might be required or produced by an activity.

Group: A group is represented by a rounded-corner rectangle with dashed lines and is used to group activities.

Annotation: An annotation is simply a note that gives the reader more information about the mode/diagram/component. In Fig. 2.3 "Insurance is included in carrier service" is an annotation.

Business processes that involve two or more profit-and-loss responsible business actors cooperating in order to create or exchange value are known as coordination processes. Therefore, a BPMN model with more than one pool is considered a coordination process model, as

it involves two or more independent entities. BPMN coordination process models form the basis of the sustainability assessment technique described in Chapter 8.

2.3.2 Value models

Value (co-creation) modelling was developed for the purpose of showing that a business model involving multiple parties in a value constellation is profitable [68]. Value models abstract away technical and operational aspects, such as IT architecture and business coordination processes, and focus solely on representing creation and exchange of economic value. As such, value models are used whenever assessing the profitability of a planned or existing business network is a critical success factor, such as during service innovation or re-engineering [69].

According to Andersson et al. [70] and Samavi et al. [71], there are three established approaches to value modelling. Namely (1) the Business Model Canvas (BMC) [72], (2) the Resource / Event / Agent (REA) ontology [73] and (3) e^3*value* [68]. The BMC take the viewpoint of a single enterprise and regards the other entities involved as third parties. It disregards the structure of the value constellation and does not allow profitability assessment. REA and e^3*value* were both designed to capture the exchanges of economic resources which occur in a network of economic actors [74], such as services, products or money. The two ontologies share strong conceptual similarities and a direct mapping is possible [70]. Since many e-services are provided by a network of collaborating enterprises, e^3*value* and REA are better suited for modelling them. However REA requires each transaction to affect both the stock and the funds of both actors involved. E-service networks also involve the exchange of intangibles such as knowledge or experience [75]. Furthermore, e^3*value* allows for quantification of revenues and expenses as a result of customer needs, and software supported analysis of these financial figures. Therefore, we opt for e3value as the value modelling ontology of choice.

e3value
describes a business in terms of actors which exchange value objects via value transfers during a fixed period of time:

Actors are profit-loss responsible entities, such as organizations, customers and intermediaries. In the example of Fig. 2.4, the “Online shop” and “Courier” are actors.

Market segments represent a group of actors of the same type. In Fig. 2.4, “Customers” are a market segment.

Value objects are things of economic value. In Fig. 2.4 “MONEY”, “SERVICE” , and “PRODUCT” are all value objects.

Value transfers are transfers of value objects, such as a payment or the delivery of a service. In Fig. 2.4, all of the blue lines between actors are value transfers.

Economic transactions are atomic groups of two or more (reciprocal) value transfers. This means that when a transaction has started, it can be assumed to be completed. Uncompleted transactions cannot occur in the profitability analysis of a value model. In

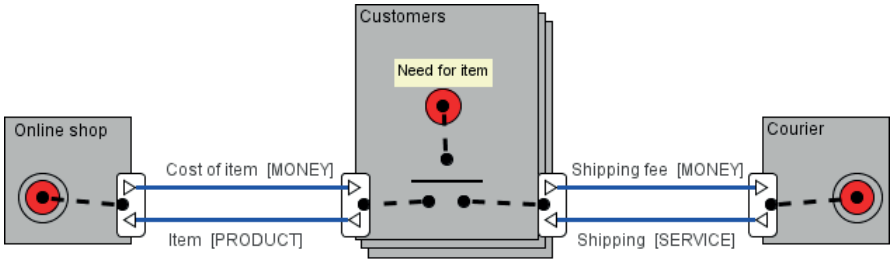

Figure 2.4: Simple e^{3}value *model*

Fig. 2.4, there are two such groups, namely "Cost of item" in exchange for "Item" and "Shipping fee" in exchange for "Shipping".

Dependency paths are chains of economic transactions, starting from a consumer need. In Fig. 2.4, the dependency path starts from "Need for item", then splits. Dependency paths do *not* represent processes [76]. They merely indicate that in the contract period, a consumer need triggers a certain combination of economic transactions, without saying when, how or in which order these transactions are performed.

Each value object has an associated monetary value (for each actor). Each consumer need has an associated occurrence rate (per contractual period). Both the monetary value and the expected occurrence rate need to be estimated by the user before any computations can be carried out. Together, these numbers can be used by the tool to estimate the financial result of each actor per contractual period. *e^{3}value* is a quantitative approach. Each value object has an associated monetary value (for each actor). Each consumer need has an associated occurrence rate (per contractual period). Both the monetary value and the expected occurrence rate need to be estimated by the user before any computations can be carried out. Together, these numbers are used by the tool to estimate the financial result of each actor per contractual period. Instead of hard values, *e^{3}value* also supports Excel-like formulas and referencing. Therefore, values can depend on other values.

A core concept of *e^{3}value* is the principle or reciprocity which says that something should always be provided in return. In other words, value transfers in one direction should always be accompanied by at least of value transfer in the opposite direction. Formally, this means that for an *e^{3}value* model to be valid all economic transactions should contain at least two transfers, one in each direction and that either all the transfers in a transaction occur, or none at all. It is important to note that an *e^{3}value* model assumes that all actors trust each other and all transactions occur as specified.

e3value serves as the basis for the value-based business risk quantification and automated identification techniques described in Chapter 6 and Chapter 7, respectively.

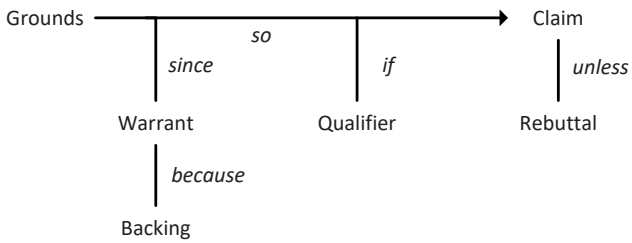

Figure 2.5: The Toulmin argument structure

2.3.3 Argumentation models

Stephen Toulmin laid the foundations for modeling arguments in his 1958 book The Uses of Argument [77]. He proposed subdividing each argument into six components (as shown in Fig. 2.5): a central *claim*, some *grounds* to support that claim, a *warrant* connecting the claim to the evidence, a factual *backing* for the warrant, a *qualifier* which restricts the scope of the claim and finally a *rebuttal* to the claim. He later identified applications of his framework in legal reasoning [78].

In the late 1980's and early 90's, argumentation models started being used to support design decisions. Specifically, the emerging field of design rationale began investigating ways to capture how one arrives at a specific decision, which alternate decisions were or should have been considered, and the facts and assumptions that went into the decision making [79]. In 1989 MacLean et al. [80] introduced an approach to representing design rationale which uses a graphical argumentation scheme called QOC (for Questions, Options and Criteria) - depicted in Fig. 2.6. The QOC is a semiformal notation which represents the design space around an artifact in terms of Questions used to identify the key issues, Options which provide possible solutions to these issues, and Criteria for choosing the best solution. Buckingham Shum et al. [81] later showed how the QOC notation can be used as a representative formalism for computer-supported visualization of arguments, with applications in collaborative environments. Mylopoulos et al. [82] introduced Telo, a language for representing knowledge about an information system intended to assist in its development. Similarly, Fischer et al. [83] claim that making argumentation explicit can benefit the design process itself.

Soon, modeling of arguments found even wider applications in decision making - especially when related to critical systems - where they started being used to make expert judgment explicit, usually by means of so-called 'cases' [84]. Safety cases, for instance, are structured arguments, supported by evidence, intended to justify that a system is acceptably safe for a specific application in a specific operating environment [85]. These arguments should be clear, comprehensive and defensible [86]. Two established approaches to safety cases are the CAE (Claims Arguments Evidence) notation [87] and the GSN (Goal Structuring Notation) [88].

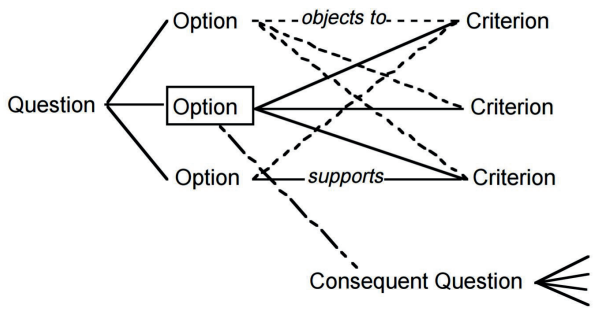

Figure 2.6: The Questions, Options and Criteria (QOC) graphical argumentation scheme

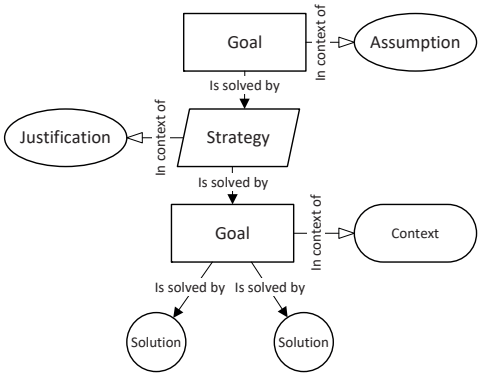

Figure 2.7: The Goal Structuring Notation (GSN)

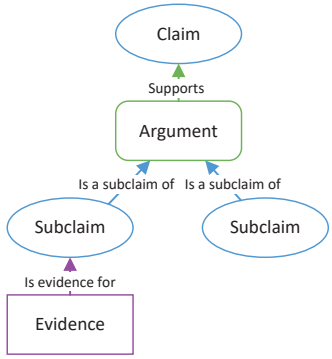

Figure 2.8: The Claims Arguments Evidence (CAE) notation

Both approaches prescribe a graphical representation of the argumentation structure but differ in terms of what this structure contains. The CAE was developed by Adelard, a consultancy, and views safety cases as a set of *claims* supported by *arguments*, which in turn rely on *evidence*. Although these concepts are expressed using natural language, the cases themselves are represented as graphs and most implementations suggest their own graphical symbols. Fig. 2.8 shows the CAE representation used by the Adelard's own ASCE tool [89]. The GSN (Fig. 2.7) was developed by the University of York and provides a more granular decomposition of safety arguments into *goals, context, assumptions, strategy, justifications* and *solutions* [88]. The arguments are also represented as a graph, with one of two types of links possible between each pair of nodes: (1) a decompositonal *is solved by* between a goal and one or more strategies or between a strategy and one or more goals, as well as (2) a contextual *in context of* between a goal, strategy or solution and an assumption, justification, or context. The notation comes with a well defined graphical language which - according to

its creator - attempts to strike a balance between power of expressiveness and usability [86].

Other, more general representations such as concept maps [90], mindmaps [91] or generic diagrams can of course also be used to represent and share knowledge, including arguments [92]. These representations have no (formal or informal) argumentation semantics and I ignore them in the rest of the chapter.

2.3.3.1 Argumentation in security

The success of safety cases has inspired other similar approaches, such as trust cases [93], conformity cases [84] and, in the field of security, assurance cases [56, 57] used to show satisfaction of requirements and misuse cases [94] used to elicit security requirements. Similarly, argumentation schemes for design rationale have been adapted to provide support for security decisions. Recently, argumentation modes have been used to encode the entire risk assessment process, from risk identification to countermeasure selection. This subsection provides an overview of these applications.

Arguing satisfaction of security requirements

Assurance cases are an argumentation-based approach similar to safety cases. They use structured argumentation (for instance using the GSN or CAE notations) to model the arguments of experts that a system will work as expected. However, while safety cases only make claims pertaining to the safe operation of a system, assurance cases are also concerned with other important system functions, in particular security and dependability [95].

Haley et al. [96] laid the groundwork for an argumentation framework aimed specifically at validating security requirements. It distinguishes between *inner* and *outer* arguments. Inner arguments are formal and consist mostly of claims about system behavior, while outer arguments are structured but informal and serve to justify those claims in terms of trust assumptions. Together, the two form a so-called "satisfaction argument".

Supporting the elicitation of security requirements

Misuse cases - a combination of safety cases and use cases - describe malicious actions that could be taken against a system. They are used to identify security requirements and provide arguments as to why these requirements are important [94].

Rowe et al. [97] suggest using argumentation logic to go beyond formalizing domain-specific reasoning and automatically reason about security administration tasks. They propose decomposing each individual argument into a Toulmin-like structure and then representing defeasibility links between the arguments as a graph. This would allow both encoding unstructured knowledge, and applying automated reasoning, for example by using theorem provers. They suggest two applications: attack diagnosis, where experts argue about the root-cause of an attack, and policy recommendation, where security requirements are elicited.

Haley et al. [58] built their conceptual framework for modeling and validating security requirements described in [96] into a security requirements elicitation process, which can help distill security requirements from business goals. The same authors later integrated their work on modeling and elicitation of security requirements into a unified framework for security requirements engineering [98]. The framework considers the context, functional

requirements and security goals before identifying security requirements and constructing satisfaction arguments for them. However, it does not consider the risks the system may or may not be facing when not all security requirements are satisfied, or when not all security goals are achieved.

Argumentation-based risk assessment

Franqueira et al. [99] were among the first to propose using argumentation structures to reason about both risks and countermeasures in a holistic fashion. OpenArgue supports the construction of argumentation models. Their proposed method, RISA (RIsk assessment in Security Argumentation) links to public catalogs such as CAPEC (Common Attack Pattern Enumeration and Classification) and the CWE (Common Weakness Enumeration) to provide support for security arguments using simple propositional logic. The method does not consider the possibility that a security threat may not be totally eliminated. Later, Yu et al. [100] integrated the RISA method and Franqueira's argumentation schema into a unified argumentation meta-model and implemented it as part of a tool - OpenRISA - which partly automates the validation process.

Prakken et al. [22] proposed a logic-based method that could support the modeling and analysis of security arguments. The approach viewed the risk assessment as an argumentation game, where experts elicit arguments and counter-arguments about possible attacks and countermeasures. Arguments derive conclusions from a knowledge base using strict or defeasible inference rules. The method is based on the ASPIC+ framework [101] and uses defeasible logic. This restricts its usability in practice.

Prakken's solution inspired a simplified approach, which used spreadsheets to encode and analyze the arguments [23]. Each argument was decomposed into only a *claim* and one or more supporting *assumptions* or *facts*. Similar to Prakken's approach, any argument could counter any other argument(s) and formulas (this time built-into the spreadsheets) were used to automatically compute which arguments were defeated and which were not.

The argumentation-based risk assessment methods described above served as inspiration for the argumentation-based risk assessment technique described in Chapter 4, as well as for the collaborative risk assessment described in Chapter 5.

Part II

Tangible modelling

3

Collaborative modelling of the Target of Assessment

Based on three peer-reviewed papers: *Tangible Modelling to Elicit Domain Knowledge: An Experiment and Focus Group* [20], *A study on tangible participative enterprise modelling* [25], and *The role of tangibility and iconicity in collaborative modelling tasks* [27].

The results of any model-driven risk assessment are dependent on the quality of the ToA model. Specifically, the more correct and complete the ToA model is, the less likely it is that relevant risks might are left out or mis-evaluated. In an attempt to streamline ToA modelling tasks, this chapter explores how features of the modelling language and of the modelling process affect the quality of the resulting model. Since modelling of socio-technical systems often requires the involvement of multiple stakeholders, I am especially interested in cases where the ToA model is constructed through a collaborative effort. To this end, the chapter describes a series of collaborative modelling experiments with students of various backgrounds, and with different modelling languages and provides interpretations of the results in terms of established cognitive theories and related work.

3.1 Introduction

In order to fully assess the risks pertaining to an information system, sufficient information about the system (i.e. Target of Assessment) is needed. But extracting and integrating information from a multitude of stakeholders from different fields, with different knowledge, some of which don't speak formal languages can be a challenge, especially when are also concerned with the social layer. Parts of this challenge are inconsistent descriptions provided by various stakeholders or documents and the intricate nature of an organization's internal socio-technical structure.

Model-driven risk assessment requires accurate and complete models of the Target of Assessment. To this end, knowledge might have to be collected from domain experts who are rarely modelling experts and don't usually have the time or desire to learn a modelling language. Furthermore, the participation of stakeholders with different backgrounds and expertise in the modelling effort is often needed. This puts strong constraints on the modelling language: It should be understandable by all stakeholders involved in modelling, even if they are not familiar with modelling languages, and it should promote their participation in the modelling effort. However, conceptual models are usually represented in (software) modelling tools using abstract graph-like structures containing boxes, arrows, and other symbols. The problem with these abstract representations is that the domain experts whose input or feedback is needed to construct an adequate model may be unfamiliar with the notation.

It is well known that *iconicity*, the resemblance of model elements and the domain of the model, can enhance understandability and learnability of signs [102–104]. However, most research on iconicity did not investigate its effects in the context of collaborative modelling.

There is also evidence that *tangible* modeling languages, by which I mean languages whose concepts are represented by physical, graspable tokens, such as plastic fiches or Lego pieces, have beneficial effects on collaborative modelling efforts [105–107]. This contrasts with what I call virtual languages, which consist of symbols on paper or on a screen or smartboard. In a comparative study, tangible models were produced faster and were of higher quality than virtual ones [108].

Based on these findings, I hypothesize that tangibility and iconicity are correlated with the understandability and usability of socio-technical languages used to construct Target of Assessment models. Therefore in this chapter I describe the findings of three experiments (see Table 3.1) aimed at identifying the effects tangibility and iconicity of the modelling language has on collaborative modelling tasks:

- The first experiment (experiment 1) compares a tangible iconic modelling language to a virtual abstract one.

- The second experiment (experiment 2) compares a tangible abstract modelling language to a virtual abstract one.

- The third experiment (experiment 3) is of a 2x2 factorial design, in which the modelling language was either tangible or virtual, and either iconic or abstract.

	Experiment 1	Experiment 2	Experiment 3
Participants	8 technical students	38 management students	20 psychology students
Task	collaborative architecture modelling	collaborative enterprise modelling	collaborative architecture modelling
Target	two student associations	fictional organisation	campus
Language	TREsPASS	4EM	IRENE
Treatments	tangibile iconic, virtual abstract	tangible abstract, vistual abstract	tangible iconic, tangible abstract, virtual iconic, virtual abstract

Table 3.1: Overview of the three tangible modelling experiments

I then combine the results of these three experiments with existing theory, in a process of analytical induction, explained in Sect. 3.2 below.

The experiments provide evidence that iconicity not only improves understandability but also modelling speed and model quality and that tangibility promotes collaboration, by facilitating uniform participation of all group members. The experiments also provide preliminary evidence that tangibility magnifies the positive effects of iconicity as well as the negative effects of abstractness on understandability, modelling speed and model quality.

3.2 Research methodology

The overarching research method used across the three experiments is known in the social sciences as *analytical induction*. Analytical induction attempts to arrive at a causal explanation by progressively redefining the phenomenon and its explanatory factors as new cases are examined. [109]. Similarly, I formulate a set of hypotheses based on related work, which we test in the first experiment. Then, in subsequent experiments, I iteratively refine these hypotheses based on the results obtained so far.

3.2.1 Validity

The validity of an experiment or series of experiments is usually broken down into internal validity (strength of the conclusions) and external validity (generalizability of the explanations). In our case, external validity is mostly threatened by the differences between students and real-life experts and internal validity is jeopardized by the low sample size. I discuss mitigations below.

3.2.1.1 External validity

All three experiments were performed with students. Students are not necessarily representative of experienced domain experts: students have no shared experience in the organization being

modeled and the supervisor did not lead the modeling session as a real-world enterprise modeling facilitator would do. Furthermore, the samples were not randomly selected, and the treatments were not allocated randomly. To overcome these limitations, and be able to make statements about causality, we use *generalization by analogy*: "If an observation is explained by a general theory, then this observation may also occur in other cases where this general theory is applicable" [110]. Thus, to explain findings we employed the analogic inference rule: the observations predicted correctly using general cognitive theories such as synchronicity, cognitive load, cognitive fit, gamification, and constructive learning (see Sect. 3.2.2) are assumed to be generally applicable to other groups of humans and carried over into the next experiment. After the three experiments, explanations which haven't been contradicted by any of the three sets of results are assumed to be generalize-able.

The cognitive theories identified in Sect. 3.2.2 are equally applicable to both experts and students. If the experimental outcomes can be explained by these general theories and social mechanisms that are present in both our objects of study and in the population –such as evolutionary capabilities of grasping physical objects and the role of construction and participation in group work in learning–, then this provides some support for the claim that similar outcomes will occur for domain experts. Such a claim would of, course, have to be substantiated by further iterations of analytical induction.

3.2.1.2 Internal validity

The sample sizes used in either of the three experiments are insufficient for statistical inference about the population. Instead, we make use of *descriptive statistics* to characterize the sample and determine whether the extent to which the expectations set out in the hypothesis hold.

3.2.2 Theoretical background

In collaborative modelling tasks, stakeholders work in a group to construct a simplified representation of an actual or potential state of affairs. I consider conceptual models, which describe social, physical and/or digital systems as a composition of concepts and relationships. These concepts may be represented by graphical signs on a screen, or by physical signs on a table or another similar surface. In the case of collaborative modelling, the representation is visible to all participants throughout the modelling effort.

Peirce defines a **sign** as "anything which is so determined by something else, called its object" [111]. A sign can be a letter, a written or spoken word, a logo, a Lego block, a diagram or anything else that refers to something beyond itself.

A sign is **iconic** if it perceptually resembles the object it represents [112]. A map can be viewed as an icon of an area, a portrait is an icon of its subject, and a model car can be viewed as an icon of an automobile. In this chapter, I call a sign **abstract**[1] if it bears no likeness to the object it represents but is rather related to it arbitrarily or by some (e.g. social or legal) convention. Abstract signs, therefore, require a dictionary that documents the relationship

[1] Such a sign is called *symbolic* in semiotics, but I prefer the term *abstract* in order to emphasize the lack of resemblance (non-iconicity) rather than its reliance on an interpretative rule, as well as to avoid ambiguous use of the word *symbol* [113, p. 237]

between the sign vehicle and sign object, and hence are a strain on the memory compared to iconic signs. For example, a word is an abstract representation since its meaning is determined by language. Similarly, a box in a UML diagram can only be understood if one is familiar with the UML language.

In addition, I refer to a sign as **tangible** if it can be grasped and manipulated by a hand.e Miniatures, post-it notes, small-scale models and Lego bricks are graspable, but a footprint in the sand is not. Neither are statues, full-scale prototypes of cars, or billboards, even though they are physical. Tangible signs may have embedded intelligence, such as in bricks on a smart tabletop [114] or FlowBlocks, used to build models of system dynamics [115], or they may be simple non-intelligent objects, such as plastic fiches with text printed on them or 3D printed shapes. Conversely, a sign is **virtual** if it is rendered digitally on a two-dimensional display, such as a computer screen, smart board or smart table and can be only be manipulated indirectly via an input device attached to the same machine. For example, a piece of text in a graphical text editor requires a keyboard to manipulate. Similarly, an icon on a smartboard – even though it can be manipulated by hand – cannot be grasped, and is therefore virtual.

To explain some of the results, I also refer to the concepts of cognitive load and cognitive fit. The **cognitive load** of a task is the total amount of mental effort required to perform a task. The theory of cognitive load suggests that performance improves when conditions are aligned with the human cognitive architecture [116]. Miller [117] claims that the ability to remember and discriminate information can be dramatically expanded by adding dimensional stimuli (such as color, sound, material & space). In particular, Hecht et al. [118] demonstrated that adding a haptic signal to visual and audio stimuli enhance perception performance. Both authors claim that tri-modal (visual, auditory and haptic) interaction enables users to absorb a wider range of details. This suggests that tangible signs are easier to understand than virtual signs. **Cognitive fit** is the reduction of cognitive load of problem-solving by fitting the representation of the problem to the problem itself. Vessey [119] showed that when the representation of concepts or information match a task, problem-solving performance for both simple and complex tasks is drastically improved. This suggests that it is easier to construct models using iconic rather than abstract signs.

3.2.3 Related work

Bjekovič highlighted the intimate relationship between enterprise modelling concepts, their representation and the community which uses them [120]. Wilmont adds that individual differences in performance on conceptual modelling tasks cannot be explained by training and experience alone and are intrinsically linked to the activation of cognitive mechanisms related to working memory, executive control and attention [121].

Fitzmaurice et al. [114] experimented with graspable user interfaces, called Bricks, which allow interacting with virtual information through an intelligent tabletop. Bricks affords synchronous manipulation, rather than through a single mouse, and has more spatial persistence than virtual signs, allowing users to make better use of spatial reasoning skills and muscle memory [114, page 447]. This suggests that in a group modelling context, tangible signs with which all participants can interact afford more equal participation compared to interaction

through a single mouse-and-keyboard, and that participants will find tangible models easier to understand and remember than virtual models.

Kim & Maher [122] showed that designers building a tangible model of an office perceived more spatial relationships, and re-framed the design problem more often, than designers building a virtual model. This suggests that tangible modelling results in models of higher quality than virtual models.

Horn et al. [123] compared tangible and graphical interfaces to exhibits in a science museum and found that people were more likely to interact with a tangible interface than a graphical interface, and that the tangible interactions lasted longer. Parmar et al. [124] compared group interaction of rural women with a health information system through an iconic keyboard interface versus an iconic tangible interface, and found that interaction through the tangible interface increased product engagement and social interaction, and improved community decision-making. Zuckerman & Gal-Oz [115] studied how stakeholders built a system dynamics model using a graphical user interface to a modelling tool, and using a tool consisting of abstract tangible signs, called FlowBlocks. Tangible modelling turned out to be slower than graphical modelling, but users reported higher levels of stimulation and enjoyment with tangible than with graphical modelling, deriving partly from the physical interaction with FlowBlocks. Grosskopf et al. [105, 125, 126] experimented with a tool for building business process models with tangible abstract elements (plastic fiches with text drawn on them) and observed that participants spent more time on modelling, and achieved more understanding of the model, than with graphical process modelling, and reported more fun building the model. These results suggest that tangible models are likely to improve participation and collaboration.

All of these studies compare tangible with virtual modelling. Bakker et al. [103] were – to the best of the authors' knowledge – the only ones to investigate iconicity in the context of collaborative modelling. They compared iconic and abstract tangible game pieces on a smart tabletop that represented a map of the game, and found that subjects preferred the iconic pieces, as it afforded better understandability.

3.3 Experiment 1: collaborative architecture modelling with technical students

In this first experiment, I wanted to investigate whether this problem can be mitigated by using a tangible representation of a conceptual model, that I call a **tangible model.**.

The experiment consists of a small-scale usability experiment (Sect. 3.3.1) and a focus group (Sect. 3.3.3), both of which are built around a tangible representation of an existing socio-technical modelling language. Based on Sect. 3.2.2 and 3.2.3, I hypothesized that a tangible collaborative modelling approach can speed up the modelling process and improve the quality of the resulting models when these models need to integrate knowledge from various fields and various stakeholders. I decompose my hypothesis as follows (Fig. 3.1):

H1 Physical representations of a conceptual model are easier to learn and manipulate than

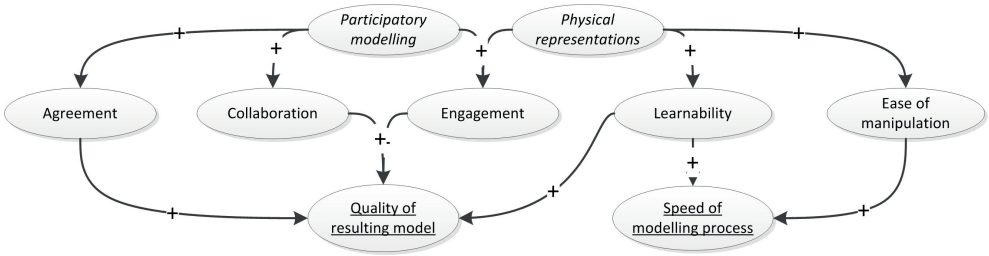

Figure 3.1: Causal graph describing my initial hypotheses. The nodes in italics are the variables I hope to influence. The underlined nodes are the target variables

abstract representations [105, 114, 126];

H2 The participatory aspect encourages collaboration and agreement between stakeholders [105, 114, 122, 124, 126];

H3 Physical and participatory modelling, similar to board games, increases engagement while reducing repetitiveness [115, 123, 124].

3.3.1 Experiment design

I set up an experiment in order to obtain an initial indication as to whether tangible modelling is more usable than abstract modelling and why. I describe the object of study, the treatment, and the measurement procedures, and then analyse the results [110].

3.3.1.1 Object of study

I was interested in comparing the relative usability of tangible and abstract representations of conceptual models for domain experts. For the experiment, we invited Business Administration and Computer Science students at the University of Twente. We divided a total of eight volunteers into two independent groups, each containing a mix of students from various tracks. This choice was motivated by the following external validity considerations: The intended users of complex, multi-layered models are inter-disciplinary teams of domain experts. For enterprise models, these teams typically consist at least of Business Experts and IT stakeholders. The choice of Business Administration and Computer Science student simulates this division of expertise to some extent.

Modelling language

To illustrate the idea of tangible modelling, I created a tangible representation of a socio-technical modelling language developed within the TREsPASS project[2] to model the socio-technical infrastructure of an organization.

The TREsPASS modelling language contains concepts such as: actors, (physical and virtual) locations, (physical and virtual) assets and access policies. The language allows four

[2] `http://www.trespass-project.eu/`.

Concept	Software representation	Tangible representation
Actor	Stickman	LEGO®character
Asset (physical)	Solid circle	LEGO®item[a]
Asset (digital)	Dotted Circle	LEGO®mini-brick[b]
Location (digital)	Box (green)	Card
Location (physical)	Box (yellow)	Box
Access policy	Text-box	Sticky-note
Relationship (position)	Solid line	Physical overlap
Relationship (possession)	Dotted line	Physical attachment
Relationship (containment)	Directional arrow	Physical overlap
Relationship (connection)	Bi-directional arrow	Line

[a] A LEGO®item resembles a real-world object and can be placed in a LEGO®characters' hand

[b] A LEGO®mini-brick is the smallest LEGO®brick available, usually of circular or cylindrical shape and can be placed on a LEGO®characters' head

Table 3.2: Mapping of concepts to representations

types of relationships to be defined between these concepts: position, possession, connection, and containment [127]. Table 3.2 shows a mapping of these concepts to tangible tokens. This is not the only possible mapping, and I use it here for illustration only.

Figure 3.2 shows a TREsPASS tangible modelling kit and Figure 3.3 shows part of a tangible model created using these conventions.

Task

The modelling target should ideally be a socio-technical system the participants are familiar with, such as their own organization. Thus, we asked the student volunteers to create a model representing the physical layout, network infrastructure (both servers and clients), important roles and associated access policies of the two student organizations of Computer Science (CS) and Business Administration (BA). Students from CS might be less familiar with the student association of BA, and vice-versa. This simulates the disjunct knowledge individual domain experts might have with regard to the model and therefore improves external validity of the choice of modelling target. To limit variation due to lack of knowledge, as well as the effect of pre-existing knowledge, each group was provided with a half-page description[3] listing the core components of each association.

Since we are measuring how well the tools fit the task, not how familiar each participant is with the modelling target, the participants were allowed to ask questions with regard to the modelling target at any time.

[3]Description and other handouts are available in full here: `https://surfdrive.surf.nl/files/index.php/s/O96tTZJFjXd2V2w`

OoS validity

While we tried to balance the groups as much as possible, there is still the possibility that some of the participants had more modelling experience, lay-outing expertise or were simply more skilled. This is a threat to internal validity, as it is a possible cause of differences in group outcomes, unrelated to the difference in modelling approaches. To measure this threat to validity, any variations in group behaviour and dynamics were noted throughout the experiment.

3.3.1.2 Treatment design

Each group was shown a brief description of the system and an outline of the task they have to perform. These were identical for both groups. The groups are given as much time as they need to understand this description.

Each group was then given a specification of the modelling concepts of the TREsPASS language. One group received the mapping of concepts to conceptual representations available in the software tool (the first two columns of Table 3.2), the other a mapping to tangible tokens (the first and last column of Table 3.2).

Each group was allowed to ask questions pertaining to the concepts before starting. When ready, each group was given an unlimited amount of time to create a model of the given system – to the best of their abilities – using the only the concepts provided. Once the modelling started, the moderator only intervened when the participants had questions about the modelling target or when the group agreed that they were done. At the end of the experiment, each student was rewarded with a 50 Euro gift-card, with bonus movie tickets awarded by means of a raffle.

Experimenter expectancy is the phenomenon that subjects try to satisfy what they think are the expectations of the experimenter, which could lead to favorable results for tangible modeling in our case. To avoid this phenomenon, we told both groups that their goal is to finish as fast as possible while maintaining consistency with the system description.

A threat to validity we were unable to mitigate was the quality of the software tool. The ability of the TREsPASS tool to manipulate diagrams, as well as the developer's choice on how to represent the concepts directly impacts its usability and thus may have biased the results in favour of the tangible approach. To find out if this threat has materialized, we need to repeat the experiment with another tool.

3.3.1.3 Measurement design

I was interested in the usability of tangible modelling versus computer-based modelling. I operationalize usability in terms the four indicators [128]: *Learnability* is operationalized by the time needed to understand the modelling language, and the number of question asked with respect to it. *Efficiency* is operationalized by the inverse of the total time needed to construct the model. *Correctness* is the number of errors at the end of the modelling process. We distinguish three types of errors: (1) Placing an element where none was expected, (2) Missing element where one was expected and (3) Using a wrong concept to represent an

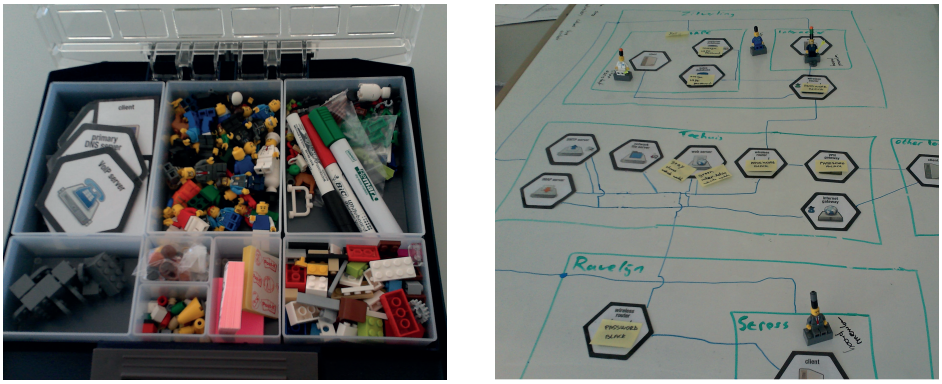

Figure 3.3: Part of a tangible TREsPASS model

element. *Satisfaction* is measured via an exit-questionnaire, containing questions on ease of task, time on task, tool satisfaction and group agreement, to be answered on a 5-point semantic differential scale with labelled end-points [129].

In order to make sure that we are indeed measuring the effects of the method, and not something else, we need to control variation due to other causes. Such causes might be internal (for example, due to the improper description of the concepts), or contextual [130]. We controlled for internal causes by providing the same definitions to all groups, irrespective of the method used. Contextual causes include variation across the subjects applying the method, the environment during application or other aspects related to the context in which the method was applied. We tried to minimise these by using the same room, layout, and instructions for both groups. Variations might still appear due to different skill levels within each group, which are harder to control for.

3.3.2 Results and analysis

Pictures of the models created by each group are shown in Figs. 3.2 and 3.3.

We observed several improvements when using the tangible method. The sample size is insufficient for statistical inference about the population, but we can still use descriptive statistics to characterize the sample and determine whether the expectations set out in the hypothesis hold. Tables 3.3 and 3.4 lists the measurements taken. While the time needed to learn the concepts was similar, the number of questions when applying them was almost double for the abstract-representation group. The tangible group finished 52% faster and with half as many errors. The tangible group also indicated 54% higher satisfaction with the tool provided and 12.5% more agreement with the resulting model while perceiving the task, on average, as being 25% easier and 24% faster. Furthermore, the tangible group did not divide tasks, suggesting increased collaboration. Detailed measurements and observations are available in an internal report [31].

Further, qualitative observations, made by the supervisor of the experiment:

Group	Time	Questions	Time	Errors
Virtual modelling	11min	7	1h42min	7
Tangible modelling	12min	4	1h7min	3

Table 3.3: Measurements

Group type	Perceived agreement	Perceived satisfaction	Perceived difficulty	Perceived duration
Virtual	4	2.75	4	4.25
Tangible	4.5	4.25	3	3.25

Table 3.4: Self-reported measurements

- The tangible modelling group worked quite fluently, and did not spend too much time on any particular phase.
- The software group split up into two sub-groups in an attempt improve the process but later had trouble defining the relationships between concepts.

Some group members provided reasons behind the usability scores reported in the exit questionnaire. Many rationales of the group using the abstract notation mentioned weaknesses of the Graphical User Interface or the way concepts are represented in the software tool as obstacles, and this confirms the bias against the abstract notation due to the quality of the software tool, that may be present in this experimental setup.

However, this group also mentioned other factors that impeded their modelling: (1) cluttering due to small screen size and tendency to try fitting everything into the screen to avoid scrolling, (2) lines and arrows overlapping with each other or with objects (due to inability to easily trace custom line shapes), (3) when many components are added, the model becomes hard to understand and debug and (4) making changes is hard and can have adverse effects on the understandability of the model. These factors remain present even if another tool were used, and this lends some support to the claim that the difference in outcomes was not only due to the quality of the software tool of the TREsPASS notation but also to the advantages of tangible modelling.

The tangible group provided less explanations for their scores, mostly discussing the following: (1) not all elements were used/useful, (2) the model exploded in size, (3) the modelling itself was not difficult if the system description is clear and (4) most of the time was spent discussing details with the team.

3.3.3 A focus group to assess utility

The experiment was conducted with students. In order to gain additional insight from practitioners, I conducted a demo session at BiZZdesign[4], a company providing consultancy on the enterprise architecture modelling. Eight consultants participated in a 2-hour demo and

[4] www.bizzdesign.nl

workshop aimed at generating feedback and discussions on the topic of tangible modelling for enterprise architectures. After a presentation describing the overall approach, introducing the TREsPASS language and its tangible mapping, they were asked to collaboratively create a tangible TREsPASS model and later to envision the possible usefulness of a similar approach to ArchiMate.

The participants of the focus group indicated several application scenarios where tangible modelling of enterprise architectures might prove useful:

- Architecture modelling sessions with domain experts (not modelling or architecture experts). Tangible modelling will allow less technical people tend to have a stronger impact on the model, as they can now manipulate the concepts themselves, and do not rely on a "modeller" to parse their input.

- Early stages of design where different types of stakeholders have to come up with an architecture; the participative aspect increases collaboration and encourages imagination.

- Models built with the goal of increasing awareness and feeling of involvement of employees with regard to the internal structure of the company. Nontechnical people can more easily understand the tangible model, which could be displayed somewhere in the company. Potentially, employees could be allowed to tweak it, thus taking enterprise architecture out of the architecture department.

3.3.4 Conclusions of Experiment 1

Before the experiment, I used cognitive theories to predict the results. If the results are in line with the prediction, it means the theories explain the results and we can therefore expect similar results in similar contexts. In this case, the tangible modelling approach outperformed conceptual modelling on almost all of the measured indicators. The tangible representations also resulted in slightly increased learnability. Users of tangible modelling also reported lower task difficulty and higher satisfaction with the tools. Overall, this reduction of the modelling effort led to a significant reduction in both the perceived and measured time needed to construct the model. Furthermore, the tangible group did not divide tasks, suggesting increased collaboration. While we were unable to measure discussion, we did observe a decrease in the number of errors in the final model. The tangible group also reported higher agreement with the resulting model. This is in line with our hypotheses that a tangible modelling approach can be easier to learn and use (H1), encourages collaboration and engagement (H3) while fostering discussion and ultimately agreement (H2).

While I was unable to control for all other contextual causes (such as participant skill or limitations of the software tool), both participants in the experiment and practitioners recognized the value of a tangible modelling approach over a tool-supported approach with abstract notations when the modellers are not familiar with the modelling language. Because these positive effects can be explained by general theories of human cognition, and are similar to the results reported by other researchers [105, 108, 125, 131], I expect similar benefits in similar situations.

However, there are limits to generalizability. I have evaluated our approach only on a small scale. Due to practical reasons such as the limited availability of physical tokens or the space to place them on, I do not expect our approach to be scalable to large systems or organizations. I further restrict the scope of our generalization to situations where formal analyses and a strict syntax adherence are secondary. Finally, the focus group indicates that a tangible modelling approach may be especially useful at the start of enterprise architecture processes, when awareness and commitment of domain experts are required. I intend to explore this further by replicating this experiment with business analysts and an enterprise modelling language.

3.4 Experiment 2: collaborative enterprise modelling with management students

A variety of techniques exist which rely on enterprise models for risk assessment [44, 50–54]. Traditional EM approaches involve an enterprise modelling expert who constructs an EM by interviewing domain experts, analyzing documentation and observing current practice, and validates the resulting model with stakeholders. Models constructed by such a *consultative* approach tend to exhibit low quality and poor commitment [46].

Recently, practitioners and researchers have advocated the potential of *participative* EM approaches, both in terms of promoting stakeholder agreement and commitment, as well as in producing higher quality models [132, 133] In other studies, *tangible modelling* approaches – in which physical tokens represent conceptual models – were found to be faster, easier and more interactive compared to a computer-supported approaches, where diagrams on a screen were manipulated [20, 105, 108]. In this section, I extend the first experiment to the EM domain by combining participative EM and tangible modeling in a hybrid approach.

I do so by means of an empirical study in a graduate EM course in which we compared the effect of using a tangible modelling set with the use of computerized tools. The results were encouraging, as the tangible modelling groups showed a higher level of collaboration, produced better results, and scored higher on post-tests. On the other hand, they felt that it took longer to produce models and reported slightly lower levels of agreement. I discuss possible explanations and implications of these results and indicate several avenues for further research.

3.4.1 Experiment design

The research goal of this experiment is to study effects of employing a tangible approach to EM compared to conducting computer-based modelling sessions. This section describes our research design following the checklist provided by Wieringa [110]. I translate this research goal into the following research question:

> What are the effects of introducing tangible modelling as part of participative EM sessions?

The effects we concentrate on are the *quality* of the models, as well as the *difficulty*, *degree of collaboration*, and *efficiency* of the modelling process. Furthermore, I am interested in the

educational value, namely the relative *learnability* with regard to 4EM. Measurement design is presented later in Sect. 3.4.1.3.

3.4.1.1 Object of study (OoS)

The tangible enterprise modeling experiment was carried out with graduate students of an enterprise modeling course at Jönköping University, with assistance from teaching staff of the Computer Science and Informatics. Students were asked to form groups no larger than five members. Although all sessions were supervised, the supervisor did not lead the sessions (as an EM practitioner would do), but just observed and provided feedback with regard to the correct application of the 4EM method. Therefore, *objects of study*, i.e. the entities about which we collect data, are EM sessions performed by students. The *population* to which we wish to generalize consists of enterprise modeling sessions carried out by domain experts.

OoS validity

The student groups were self-formed. As a result, EM experience or knowledge of participants may be unequally distributed. Besides, measuring effects of tangibility can be hampered. While some groups may consist of very conscientious students, others could contain uninterested ones. Moreover, some students may be shy and thus could collaborate less with their group. Nevertheless, as all of these phenomena may exist in the real world as well, these aspects (arguably) also make our lab experiment more realistic in terms of external validity. Specifically, participation, modelling effort, and quality of the models produced by different groups may differ. These differences may be caused by multiple factors in addition to the difference between tangible and computer-based modelling. To take these possible confounding factors into account, we tried to make the presence of these phenomena visible by performing most measurements on individuals instead of on groups and by observing group behavior, dynamics and outliers.

3.4.1.2 Treatment design

Participants were first presented with a description of a real-world anonymized case of a sports retailer company. Each group was then given five weeks to perform a business diagnosis of the retailer by constructing three out of the six 4EM sub-models, namely the goal, concepts and business process viewpoints. The groups were instructed to perform as much of the modelling as possible together, during weekly, dedicated modelling sessions (4 hours session a week). Treatment was self-allocated: Groups were allowed to choose between tangible or computer-based modelling sessions, as long as there was an even split. The tangible modelling groups were given a large plastic sheet, colored paper cards and pens to create the models. Different colors of paper cards were representing different types of elements — goals, problems, concepts and processes, similar to 5.1 of [55]. Cards could be easily attached to the plastic sheet and moved if necessary. These groups were instructed to make use of the cards when collaboratively building the models, and create digital versions of models after that. By contrast, the computer-based modelling groups (allowed to use a diagram tool of their choice) started working directly on a computer.

	Factor	Indicator	Type	Scale
Result	Model quality	Semantic quality	Group	1(poor) - 5 (excellent)
		Syntactic quality	Group	1(poor) - 5 (excellent)
Process	Difficulty	Perceived difficulty	Individual	1(very easy) - 5(very dif-lt)
	Collaboration	Observed collaboration	Group	1(very low) - 5 (very high)
		Perceived agreement	Individual	1(none) - 5 (very much)
	Task efficiency	Observed pace	Group	1(very slow) - 5 (very fast)
		Perceived duration	Individual	5 / 10 /15 / 20 / >20 hours
Edu. value	Learnabilty	Exam questions on 4EM	Individual	0-15
		Final report grade	Group	F (fail) - A (excellent)

Table 3.5: Operationalized indicators and measurement scales

Treatment validity

While in real-life situations, the modelling technique might sometimes be prescribed, it was noted that free choice of the preferred notation to be used in EM activities and its effects on ease-of-use and understandability is desirable and worth investigating [134]. Our experiment is similar to situations where modellers have the freedom to choose their tools, and dissimilar to situations where the modelling technique is prescribed to them. Noticeably, the choosing of tools may hamper external validity of this study. In addition, internal validity may be threatened by the fact that participants were informed about both available treatments. This may cause an observer-expectancy effect, where participants change their behavior based on what they think the expectations of the experimenter are. In an attempt to mitigate this, we did not inform participants about the goal of the research nor of the measurements.

3.4.1.3 Measurement design

I am interested in comparing the effects of tangible modelling versus computer-supported modelling on the quality of the result, on the modelling process, and in connection to their educational potential. Table 3.5 provides an overview of the operationalized concepts, their indicators and respective scales. I explain my choices below.

The quality of a conceptual model is commonly defined on three dimensions: syntax (adhering to language rules), semantics (meaning, completeness, and representing the domain) and aesthetics (or comprehensibility) [135, 136]. In this study we measured the *semantic quality* and *syntactic quality* of the resulting model and omitted measuring aesthetics due to its highly subjective nature. Semantic and syntactic qualities were estimated by the supervisor on a 5-point semantic difference scale by comparing the final models with the case description and 4EM syntax, respectively.

With regard to the modelling process, relevant factors are difficulty, amount of collaboration, as well as the overall task efficiency. Difficulty is a purely subjective measure [137] and was therefore measured as *perceived difficulty* via individual on-line questionnaires distributed at the end of the course. The questions (available at `https://surfdrive.surf.nl/files/index.php/s/ixW4JlmtXma6OlE`) were linked to a semantic difference scale, and provided room for optional free-text explanations. Collaboration — the amount

interaction between group members — is crucial for creating a shared understanding of a representation [138]. We indirectly measured collaboration by means of two indicators: *observed collaboration* (estimated by the supervisor throughout the five sessions) and *perceived agreement* (measured the same on-line questionnaire). Task efficiency is the amount of time to produce the final, digital model. In our case, because the task was spread across several weeks and groups may have worked at home, we could not directly measure the time groups spent. Therefore, we operationalize task efficiency in terms of *perceived duration* (measured via the online questionnaire) and *observed pace* (progress achieved during the dedicated modelling sessions, as estimated by the supervisor).

Finally, to evaluate the educational value of a tangible modelling approach, we looked at the final results of the students. As indicators, we use *final report grades* and students' performance on two *exam questions on 4EM*. The final report grade was decided by the supervisors and lecturer together, while exams were graded by the course lecturer, who otherwise did not take part in this study.

Measurement validity

Potential issues with measurement validity might occur due to the qualitative and self-reported nature of the data (internal causes), as well as the loosely controlled environment (contextual causes). Potentially, different scales could confuse the respondents. For instance, '1' corresponds to 'poor' in one case and to 'very easy' in another. Furthermore, the fact that model quality, observed collaboration, task efficiency, observed pace, and the final report grade were all assessed the supervisor of the modelling sessions who was also one of the experimenters, may also influence validity. The lack of randomization in group forming and the fact that groups they were allowed to choose their diagram tool should also be taken into account. To preserve measurement validity, we tried to reduce mono-operation bias by operationalizing each measured concept in terms of two different indicators where possible. We also attempted to minimize mono-method bias by combining self-reported and observed values where possible.

3.4.2 Results and analysis

Measurements

We gathered data on the work of 38 students from Information Engineering and Management (School of Engineering) and IT, Management and Innovation (School of Business), who formed eight groups of three to five students. Although self-assigned, exactly half of the groups opted for "physical" (i.e. tangible) modelling. Every group submitted a report containing final, digital versions of their model (constructed in a tool of their choice), as well as justifications of their design decisions. No student dropped out of the modelling sessions, but only 23 filled in the online questionnaire and 26 took part in the exam. This reduces the relevance of statistical measures such, so we do not compute significance. We use statistics descriptively instead to provide an indication of the effects.

Results per group (Table 3.6) show a higher degree of collaboration and a faster pace of the tangible groups. We observed that these groups tended to communicate more and make better use of the dedicated modelling sessions, while computer-based groups tended to divide tasks and occasionally skip sessions. Also, tangible groups produced models with slightly less syntactic quality, but with a higher level of content correctness. *Individually* (Table 3.7),

Group type	Count	MEASURED			OBSERVED	
		Semantic quality	Syntactic quality	Final report grade	Pace	Collaboration
Tangible	4	4 (σ 0.82)	3.75 (σ 0.5)	see Fig. 3.4	4 (σ 2)	4 (σ 1.41)
Non-tang	4	3.5 (σ 0.57)	4.25 (σ 0.5)	see Fig. 3.4	3 (σ 0.8)	2.75 (σ 1.25)

Table 3.6: Group measurements, aggregated[5] per group type

Respondent group type	Count	MEASURED	SELF REPORTED		
		Exam questions on 4EM	Perceived difficulty	Perceived agreement	Perceived duration
Tangible	12	8.76 (σ 4)	3.08 (σ 1.08)	3.83 (σ 1.03)	see Fig. 3.5
Non-tang	11	8.15 (σ 3.86)	3.55 (σ 0.82)	3.91 (σ 0.54)	see Fig. Fig. 3.5

Table 3.7: Individual measurements, aggregated[5] respondent group type

participants from tangible groups reported slightly lower perceived agreement (by 2%) and lower difficulty (down 13%). Furthermore, such participants sometimes reported of longer durations than their peers from groups using only a computer. Fig. 3.5 shows that more tangible modelling participants than computer-based modelling participants perceived duration as being more than 20 hours. Regarding the educational effect of using tangible models, we have noticed a significant improvement on both measured indicators for learnability. The reports submitted tangible groups were scored consistently higher than others (see Fig. 3.4). Furthermore, tangible modelling students obtained, on average, 7.5% higher on questions related to the 4EM method and its application.

[5]Full anonymized results available at: `https://docs.google.com/spreadsheets/d/1RB74Gk1O-G43Wv2WdR4c3a-XCwb1E8-5KqhCqRlcKQo`

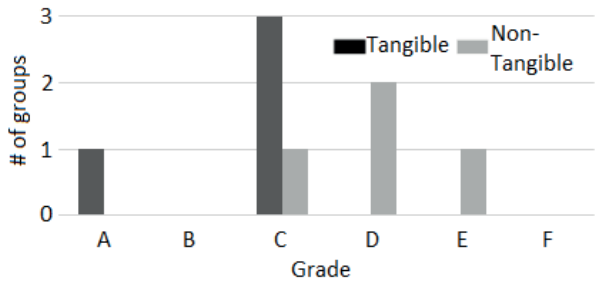

Figure 3.4: Distribution of final report grades

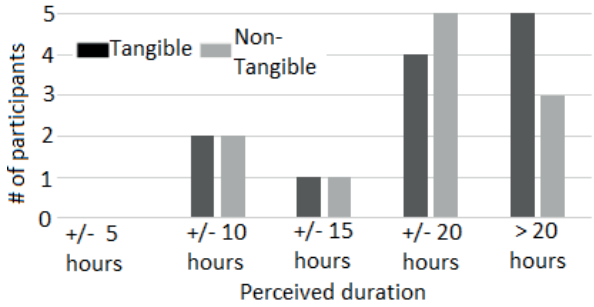

Figure 3.5: Distribution of perceived duration

Discussion

I cannot exclude the possibility that all differences between tangible and computer-based groups are random fluctuations explained by chance alone. Also, since our sample and treatments were not formed and allocated randomly, I refrain from using statistical inference to generalize. However, plausible explanations to interpret the noted differences can be offered.

First, the reduced syntactic quality of tangible models can be explained by the fact that tangible modelling does not constrain the syntax of models as strictly as computers do. Thus, some students might have used this freedom to construct models that are not syntactically correct.

Second, our explanation for the higher semantic quality of tangible models is that the tangible groups interacted more (without dividing tasks) and seemed to work harder (higher pace and longer perceived duration). This can be that tangible modelling supported participation by providing the fun-factor. The perceived duration might have also been influenced by the fact that after completion of tangible models, the students had to enter them into a software tool.

Third, the lower perception of difficulty and better exam results of tangible groups can be explained by the theory of constructivism, which says that learning is most effective when people jointly create tangible objects in groups.

Finally, the slightly lower perceived agreement within tangible groups may be explained by higher levels of collaboration. Due to less subdivision of tasks, tangible modelling forced groups to promptly discuss disagreements. It is also possible that the computer-based groups had lower *actual* levels of agreement without noticing this. Since they divided tasks among members and discussed less than the tangible groups, they may have overlooked some disagreements or misinterpretations. While our data does not exclude this possibility, it does not support it either. More research is needed to test this hypothesis.

Generalizability

Given the limited availability of statistical data, this study employs generalization by analogy: ”If an observation is explained by a general theory, then this observation may also occur in other cases where this general theory is applicable” [110]. In this case, the observations were predicted and can be explained in terms of social or psychological mechanisms such as synchronicity, cognitive load, cognitive fit, gamification, and constructive learning (see Sect. 3.2.2). Therefore, we can expect to observe similarities in cases where groups of humans perform similar tasks, such as stakeholder modelling workshops.

3.4.3 Conclusions of Experiment 2

Implications for research.
Earlier research shows that tangible modelling promotes collaboration because of synchronicity, manipulability of physical tokens, and increased fun while leading to better results due to the joint construction of physical models [20]. The results of Experiment 2 are consistent with these reports to a large extent. A notable difference is the perception of increased duration, which contrasts with both literature and the findings of Experiment 1, where tangible modelling was observed to be faster than computer-based modelling. Results also show that collaborative modelling may increase the effort required for modelling, contrary to [132] and [139]. One explanation for this difference is that the previous experiment ([20]) used iconic physical tokens, i.e. objects that resemble the entities being modelled, which made them easier to understand. To test this explanation, I need to compare tangible modelling with iconic tokens and with plastic cards in future research. Also needed is a similar real-world experiment with EM practitioners, to verify the external validity of our results. Another interesting direction for further investigation is computer-based participative modelling tools (such as using smart boards and touch screens).

The observed similarity in perceived agreement in the tangible and computer-based groups contradicts the results of Experiment 1, where tangible modelling with iconic objects was found to lead to higher agreement among subjects. I speculate that computer-based modelling, where tasks are divided and engagement in the modelling process is less intense, may actually lead to lower agreement among modellers without them being aware of it. This hypothesis could be tested by repeating the experiment described in this section and adding a post-test that measures the level of agreement among modellers.

Implications for practice.
Our results suggest that tangible enterprise modelling could be a useful tool for building consensus of stakeholders with diverse backgrounds and little EM experience. This is particularly useful in the early stages of enterprise modelling, when the goal is to improve the quality of the business [20, 46]. Our results also suggest that tangible EM has a positive educational effect by providing higher understandability and improved learnability.

3.5 Experiment 3: collaborative architecture modelling with psychology students

Experiments 1 and 2 compared various combinations of tangibility and iconicity, and it is not clear whether the effects of these two variables have always been distinguished well. To distinguish these effects, I attempted to systematically analyze these two variables in a factorial experiment (*Experiment 3*), reported in this section. Based on the above, we formulate two hypotheses, to be tested in this third experiment:

- H1: Iconicity improves understandability.
- H2: Tangibility improves collaboration.

	Iconic	Abstract
Tangible	Whiteboard, magnetic objects, markers	Whiteboard, magnetic cards, markers
Virtual	MS Visio, icons, lines	MS Visio, boxes, lines

Table 3.8: The four toolsets

3.5.1 Experiment design

This experiment was carried out at the University of Twente, in collaboration with the department of Cognitive Psychology and Ergonomics.

3.5.1.1 Object of Study

Our object of study consists of groups of people collaboratively building a model of an existing system or of a new design. Our sample consists of four groups of five psychology students collaboratively building a model of their university's campus and then updating that model to represent their view of the campus' future.

3.5.1.2 Treatment design

The four treatments
We are interested in two variables: tangibility and iconicity. Therefore, we have four treatment groups, each using different representations (signs) of the same modelling language, as shown in Table 3.8.

Modelling language
The underlying language used in this experiment is a simplified version of the IRENE language for modelling smart cities [140], adapted to include elements specific to university campuses. The IRENE language is designed to be used in stakeholder workshops and is therefore intended to be usable by nontechnical domain experts.

Tasks
After filling in a demographics questionnaire[6], participants were randomly allocated to one of four groups. Each group was taken to a separate room and given one of the four toolsets described in Table 3.8, accompanied by a document describing the semantics and syntax of each modelling element. These textual descriptions were identical for all groups.

The groups were then asked to familiarize themselves with the toolset and the descriptions, requesting clarifications if needed (*Task 0*).

Once each group declared that they understood the language, they were given their first modelling task: to build a model of the current campus, as accurately and completely as possible using the tools provided (*Task 1*). The groups could take as long as they like to build the model, after which participants received individual questionnaires[7].

Each group then received a second modelling task: to change the model they just constructed, by adding or removing elements based on how they think the campus should look

[6] The questionnaires are available in full: `https://surfdrive.surf.nl/files/index.php/s/QzscwpSO6Xf02w0`

[7] See footnote 6

like in the future (*Task 2*). As guidance, they were given a fictional budget of one million Euro and a list of prices for each element. Once the group declared they are satisfied with the model, each participant received a final questionnaire[8] containing the same questions as for task 1 but with an additional question on overall enjoyment.

During the two modelling tasks (Task 1 and Task 2) the students were allowed to ask factual questions about the campus but not about the modelling language. Each task was timed and, in addition, the two modelling tasks were videotaped.

3.5.1.3 Measurement design

To evaluate H1, we need to measure understandability and model quality, and to evaluate H2, we need to measure collaboration and task efficiency. In addition to these, we also measure satisfaction, so that we measure the three components of usability as defined by Hornbaek [141], namely quality of outcome, task efficiency and user satisfaction. Our operationalized measurements are listed in Table 3.9 and explained below.

Language understandability is evaluated by measuring the *time* taken by each group to read and declare that they understand the language, and the *number of questions* they have during this time. In addition, we measured *perceived language understandability* and *learnability* of the language, as reported in the individual questionnaires distributed after each of the modelling tasks.

Quality of the outcome (i.e. of the resulting model) is one of the three usability factors listed by Hornbaek [141]. It is commonly measured as: **semantic quality** (how well the model represents the domain) and **syntactic quality** (how well the model adheres to the prescribed syntax) [135, 136]. We operationalize semantic quality as *incorrectness* (buildings represented in the model but not present on the campus)and *incompleteness* (buildings that are present on the campus but not represented in the model). Syntactic quality is operationalized as the number of *syntactic mistakes.* These quality measurements were performed first by two of the authors independently and then discussed.

Task efficiency, Hornbaek's second usability dimension, is operationalized in terms of *time taken to complete task,* as well as several indicators for perceived effort [141]: *perceived difficulty* and *perceived time* to complete task. Both are measured via questionnaires after each modelling task.

Hornbaek's last usability dimension is **satisfaction**. We measured *perceived tool satisfaction* and *perceived enjoyment* via the same questionnaires.

Collaboration has to do with the relative effort each participant expended in communicating with the others and resolving differences [141]. Furthermore, higher intra-group agreement is thought to be indicative of better collaboration in group problem-solving [142]. Therefore, we operationalize collaboration in terms of two indicators: amount of discussion and agreement. The *amount of discussion* is measured by (1) annotating the video recordings and aggregating the average number of words and turns per participant per minute for each group as an indicator of the individual contributions to the discussion [141, 143] and (2) by computing the coefficient of variation of words per participant as an indicator of how evenly the discussion was spread [143, 144]. Finally, the level of *perceived agreement* of each participant with their group's result is queried via questionnaires at the end Tasks 1 and 2.

[8] See footnote 6

Dependent variable	Indicator	Scale
Language understandability	Effort to understand language (Task 0)	# minutes # questions
	Perceived language understandabi-lity (Task 1 and Task 2)	1 (Easy to understand) - 5 (Difficult to understand)
	Perceived language learnability (Task 1 and Task 2)	1 (Very easy) - 5 (Very hard)
Semantic quality	Incorrectness (Task 1)	|M – D| (model elements not in domain)
	Incompleteness (Task 1)	|D – M| (domain elements not in model)
Syntactic quality	Deviation from syntax (Task 1)	# syntactic mistakes
Task efficiency	Time to complete task (Task 1 and Task 2)	# minutes
	Perceived difficulty (Task 1 and Task 2)	1 (Very easy) - 5 (Very difficult)
	Perceived time to complete task (Task 1 and Task 2)	1 (Very little time) - 5 (Too long)
Satisfaction	Perceived tool satisfaction (Task 1 and Task 2)	1 (Very satisfied) - 5 (Very unsatisfied)
	Perceived enjoyment (Task 1+2)	1 (Boring) - 5 (Fun)
Collaboration	Amount of discussion (Task 1 and Task 2)	# words per minute # turns per minute Coefficient of variation for words per participant
	Perceived agreement (Task 1 and Task 2)	1 (Don't agree) - 5 (Fully agree)

Table 3.9: Measured indicators

3.5.2 Results and Analysis

Figure 3.6 shows the models built by the students during Task 1 and Table 3.10 and Table 3.11 summarize the averaged values per group for Task 1. I omit the results of Task 2 as this task was creative rather than descriptive in nature and therefore cannot be easily compared with previous experiments. The complete set of measurements and observations may be examined at `https://goo.gl/RP6VM1`.

Group type	Difficulty	Time	Satisfaction	Learnability	Agreement	Understand-ability	Enjoyment (incl. Task 2)
TI	3.4	2.8	2.4	2.2	4.2	2	3.8
TA	3.4	3.4	3.4	2.8	3.2	3.2	4
VI	2.8	2.4	1.4	1.6	4.75	1.8	4.5
VA	3.4	2.8	2.4	2	4.6	2.2	3.6

Table 3.10: Self-reported measurements (on task 1, unless otherwise specified)

Group type	Ques-tions asked	Time (task 0)	Time (task 1)	Turns/ min.	Words/ min.	CV words/ part.	|M – D|	|D – M|	Syntactic mis-takes
TI	0	3	33	18	105	*0.45*	2	0	1
TA	0	2	55	27	88	*0.27*	5	4	3
VI	2	3	40	34	107	*0.37*	4	3	0
VA	4	9	49	26	92	*0.67*	6	2	0

Table 3.11: Objective measurements (on task 1, unless otherwise specified)

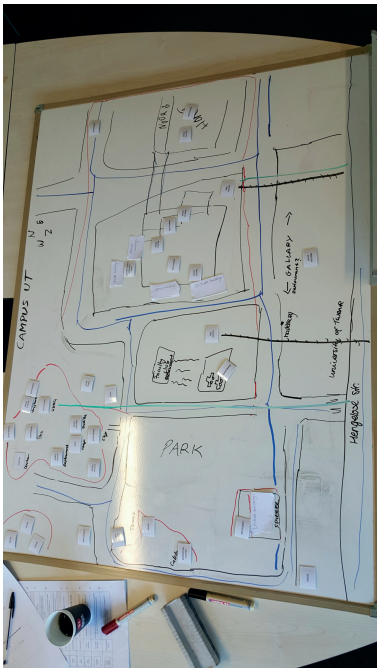

(a) Tangible iconic group

(b) Tangible abstract group

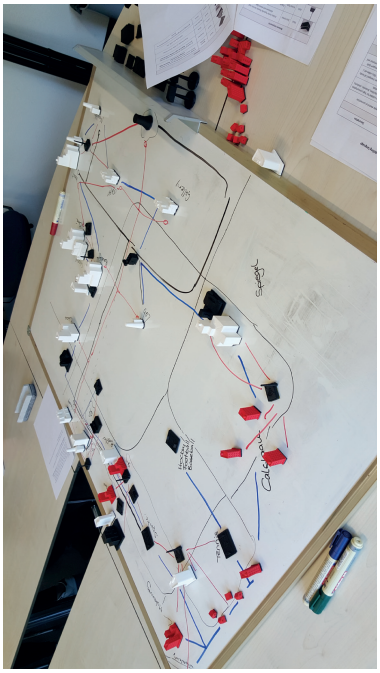

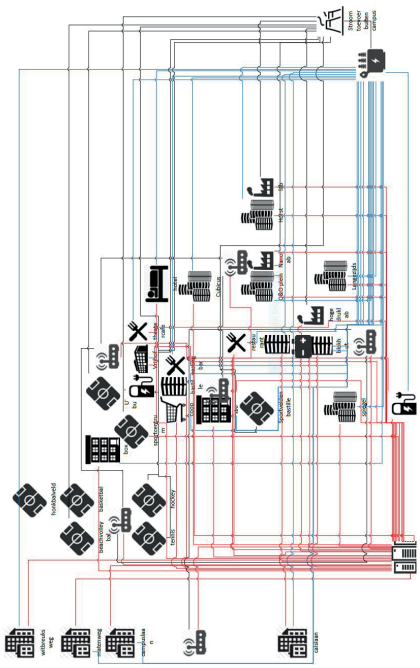

(c) Virtual iconic group

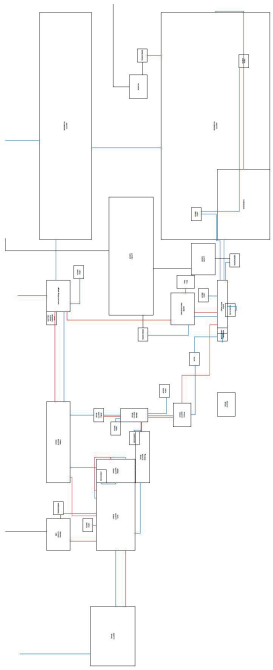

(d) Virtual abstract group

Figure 3.6: Models produced during Task 1

Language understandability

The results show that iconicity improves understandability, and abstractness decreases it. This agrees with our expectation in H1 and can be explained by the theory of cognitive fit. The effect of tangibility on understandability is less clear: the two tangible groups reported lower understandability and learnability than the corresponding iconic groups, contradictory to the observations of Fitzmaurice [114] and Luebbe [126] that tangibility improves understandability. One explanation might be that the subjects of this experiment - psychology students - had little to no modelling experience and were unfamiliar with parts of the domain and therefore benefited from the syntactical constraints built into the software tool. This suggests that to predict the effect of tangibility on language understandability in future experiments, we may have to include the variables *experience with modelling and design* and *familiarity with the domain.*

The measurements also suggest an interaction between tangibility and iconicity: Tangible iconic signs are perceived as slightly less understandable than virtual iconic signs, but tangible abstract signs are perceived as considerably less understandable than virtual abstract signs.

Task efficiency

Iconicity sped up the modelling process, and abstractness decreased it. This agrees with the theory of cognitive fit. Tangible modelling was perceived to be more difficult by our subjects than virtual modelling. This contradicts observations of earlier Experiments 1 and 2, where tangible modelling was perceived to be easier than virtual modelling [20, 25]. However, this contradiction may be explained by the flip side of our above explanation: Experiments 1 and 2 where done with students of computer science and technical management science, who have been trained in modelling and design, and felt more comfortable with the modeled domain, a smart campus and an enterprise architecture, respectively. Current measurements do not support nor rule out this explanation, and future experiments should include the variables *experience with modelling and design* and *familiarity with the domain* to test this explanation.

Tangibility magnified the effect of iconicity on task efficiency. Tangible iconic models were built faster than virtual iconic models, and tangible abstract models were built slower than virtual abstract models. This interaction may explain why in Experiment 2 [20], the tangible iconic group built models twice as fast as the virtual abstract group. The interaction is also consistent with the observations of Zuckerman et al. [115] and of Lübbe [126] that tangible abstract groups took longer to build a model than a virtual abstract groups.

Quality of product

While iconicity had no consistent effect on model quality, tangibility had a clear effect on syntactic quality: All tangible models contained syntactic mistakes, but the virtual models contained no syntactic mistakes. A possible explanation for this is that our modeling tool enforces syntactic constraints, while our tangible modeling tool enforced no syntactic constraints.

In terms of semantic quality however, we could not clearly separate the effects of tagibility and iconicity. Tangible iconic models were more complete and correct than any of the other models, an effect observed in both Experiments 1 and 2 [20] and in this latest, third experiment. However, tangible abstract models in this experiment had lower semantic quality than the corresponding virtual models, whereas in Experiment 2 [25] they were slightly better. A possible explanation of this apparent contradiction is the apparent uncertainty of the subjects involved in Experiment 3 in the face of the freedom afforded by our tangible modelling

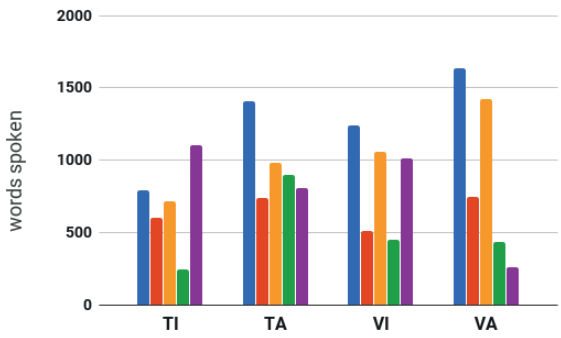

Figure 3.7: Words per participant. Each bar represents a different participant.

languages, combined with their lack of knowledge about the domain. This uncertainty may be aggravated by the preference of naive (nontechnical) users to represent systems with iconic diagrams rather with abstract diagrams, compared with the preference of technical students trained in modelling and design languages, to represent systems in abstract diagrams [145].

Satisfaction

Tool satisfaction was lowest for the tangible abstract group and highest for the virtual iconic group. In Experiment 2 [25] this was the reverse. This is consistent with our proposed explanations above that the subjects of our latest, third experiment were thrown off by freedom afforded by our tangible modelling tools, and find it more difficult to make abstract models than to make iconic models. Task enjoyment was highest for the virtual iconic group. This partly contradicts previous conclusions ([20, 115]), that tangible modelling is *always* more enjoyable than virtual modelling. Task enjoyment was highest for the virtual iconic group. We may again explain this in terms of lack of modelling experience and preference of naive users for the guidance provided by virtual tools, and for iconic modelling tools.

Collaboration

The iconic groups spoke more (words/min) than the corresponding abstract groups, despite similar or less turn-taking. Iconic groups also exhibited higher agreement. Tangibility promoted more equal participation (lower CV) for groups working with abstract signs, but slightly less equal participation in iconic groups. However, a closer analysis of the video revealed that the tangible iconic group contained a “silent” participant (see Fig. 3.7) which drove up the CV. A clear effect of tangibility compared to virtuality in Experiment 3 as well as in Experiment 2 [25] is that virtuality promotes task division, while tangibility does not.

This is easily explained by the physical setup of the virtual modelling groups, which all used a single keyboard with mouse and confirmed by the number of words per participant (Fig. 3.7), which shows that the virtual groups were dominated by one or two participants. Analysis of the videos showed that these were the people grabbing the keyboard and/or mouse. Iconicity resulted in more perceived agreement than in abstract models. This effect is slight for virtual models, but is large for tangible models. So tangibility is also a magnifier for this effect for iconic models.

3.5.3 Discussion

The positive effect of iconicity on understandability can be explained by the theory of cognitive fit. Improved understandability, in turn, explains higher modelling speed and higher agreement, assuming that modelers have similar conceptual models of the domain.

Tangibility magnifies these effects because tangible signs are graspable, allowing modelers to use their spatial reasoning skills [114]. This is in line with the cognitive theory of grounded conceptual knowledge [146]. In this view, human cognition does not initially (and primarily) develop by formal instruction, but by the interactions between perception and action with which the human child explores its environment. Then, and in later life, repeated forms of these interactions result in the neural brain patterns on which conceptual knowledge, as used in language and reasoning, is based [147]. Combining graspability and iconicity would more easily reactivate these patterns, and hence the concepts they represent, thereby contributing to an increased and faster understanding.

In addition, all participants have equal access to tangible signs, and this can speed up modeling. However, this effect is moderated by other variables, such as whether the modelers are familiar with the domain, and may even be influenced by their personality - eg. willingness and ability to participate actively in a group - or higher order cognitive processes - such as relational reasoning and abstraction [121].

Note that these explanations hold even though the subjects Experiment 3 *perceived* tangible modelling as harder than virtual modelling. The increased freedom of tangible modelling made them feel uncertain. Combined with the natural tendency of humans to adapt the linguistic structure they use to describe something based on the context and their situational goal [148], this can lead to more syntactical mistakes in tangible models. Still, even though they claimed to have no familiarity with the domain, the models of the iconic tangible group were semantically complete and contained the least mistakes of the four experimental groups.

3.5.4 Conclusions of Experiment 3

We conclude that our experiments provide support for H1 (Iconicity improves understandability), but that only provide partial support for H2 (Tangibility improves collaboration). Instead, iconicity turned out to improve group discussion and perceived agreement, which are indicators for collaboration. Tangibility mostly amplifies these effects of iconicity, but also supports more equal participation of group members than is possible with virtual models using a single mouse-and-keyboard setup.

In addition, we found that modeling experience of participants, familiarity with the domain, and personality of the participants moderate these effects. Future studies should control for these variables.

The constructs outlined in this chapter and the relationships between them highlighted by hypotheses H1 and H2, together with observations of previous experiments provide support for constructing and updating theories related to group modelling of complex systems.

3.6 Validity

Or sample sizes were small and not randomly selected. Therefore, I could not meaningfully use statistical inference so the reasoning was case-based. Specifically, I tried to explain

the results of a series of experiments in terms of existing theories in a process of analytical induction [110, 149], as explained in Sect. 3.2. As indicated at length by Znaniecki [150], who coined the term "analytical induction", this has been a common pattern of reasoning in the experimental physical sciences. In this process, the researcher defines a class of phenomena, studies cases in which this phenomenon occurs, and provides explanations of the observations, relating these explanations to existing theory. This process is repeated, where each the researcher may select cases where he or she expects the theory to be confirmed or disconfirmed. Based on the observations done in a case, existing theory may be refined to account for all of the cases observed so far. What is new in our approach is that our cases are not observational but experimental.

The major threat to internal validity is that the differences between groups that we have observed are too small to infer causality. Despite this being also true for Experiments 1 and 2 [20, 25], all three experiments have identified differences that point consistently in the same direction, and that can be explained in terms of existing theories such as that of cognitive fit and grounded conceptual knowledge. This does not definitively prove our explanations, but it does increase their internal validity. To provide more support for our explanations of the phenomena, more experiments need to be performed and interpreted in an ongoing process of analytical induction.

External validity is the support for our generalization to a wider population of stakeholders who collaboratively make models of a domain. We, therefore, provide interpretations in terms of general cognitive theories which can be assumed to be valid for students and non-students alike.

3.7 Conclusions and future work

In *Experiment 1* [20] two groups of modelling novices built models of the physical layout and IT architecture of a university campus using tangible iconic signs or virtual abstract signs, respectively. The tangible iconic group built a model twice as fast as the virtual abstract group, with a quality twice as good. Tool satisfaction was higher for the tangible iconic group as well.

In *Experiment 2* [25], eight groups of modelling novices collaboratively constructed enterprise models using either tangible abstract signs or virtual abstract signs. Again, the tangible groups produced a model of higher quality. In both Experiments 1 and 2, the virtual groups divided tasks among themselves, whereas in the tangible groups all participants had the same task, indicating a higher amount of collaboration in the tangible groups. However, agreement about the model was only slightly higher for the tangible group in Experiment 1 and slightly lower for the tangible group in Experiment 2.

In *Experiment 3* tangibility promoted equal participation, and iconicity had a beneficial impact on understandability, modelling speed and model quality. Tangibility magnified the effects of iconicity.

Overall, the three experiments provide evidence that iconicity not only improves understandability, but also modelling speed and model quality and that tangibility promotes collaboration, by facilitating uniform participation of all group members. The experiments also provide preliminary evidence that tangibility magnifies the positive effects of iconicity as well as the negative effects of abstractness on understandability, modelling speed and model quality. These effects appear to be more pronounced in individuals with a technical background, but

this should be substantiated in further research.

An open issue that we have not investigated is how to facilitate entering a tangible iconic model in a computer, once it has been built by a group of stakeholders. Luebbe [126] did this for tangible abstract models by camera, but for tangible iconic models, intelligent tangible signs on a smart tabletop may be a more feasible option. Another issue for future research is that complex, socio-technical systems consist of physical, social and virtual elements, not all of which can be represented in an iconic way. In these cases, models will likely consist of iconic as well as abstract signs, and the issues we have observed with abstract signs manipulated by nontechnical experts come into play.

Part III

Argumentation modelling

4

Argumentation based risk assessment

Based on *Argumentation-Based Security Requirements Elicitation: The Next Round* [23]

In this chapter I present results from my search for a scalable argumentation-based information security RA method. I start from previous work on both "heavy-weight" formal argumentation frameworks and "light-weight" informal argument structuring and try to find a middle ground. An initial prototype using spreadsheets is validated and iteratively improved via several case studies. Challenges such as scalability, quantify-ability, ease of use, and relation to existing work in parallel fields are discussed. Finally, I explore the scope and applicability of the approach with regard to various classes of Information Systems while also drawing more general conclusions on the role of argumentation in security.

4.1 Introduction

Eliminating all risks pertaining to the usage of an information system is impossible and risk treatment decisions have to be selective: some risks can be mitigated in several ways while others will have to be accepted. The ability to trace back previous decisions is important if they have to be defended or revised, or if a new risk assessment needs to be performed. Firstly, the decision maker may have to justify risk mitigation decisions made earlier, for instance in the case of a successful attack [151] or to satisfy the "reasonable security" requirements of regulators [152]. Second, the ever-changing IT landscape forces decision makers to frequently revisit decisions pertaining to their information systems. In fact, the new European GDPR (General Data Protection directive) explicitly requires data controllers and processors to ensure "ongoing" confidentiality, integrity, availability, and resilience of processing systems and services [153, art 32(1)(b)]. Third, related systems may face related, but not identical risks and therefore, the ability to reuse (parts of) the arguments made for similar systems facilitates decision making in the future [154]. By recording the argumentation behind the identified risks and selected countermeasures, risk assessments can be re-visited when an attack takes place, extended when new risks surface, and re-used in related products or contexts. Altogether, this highlights a need to document decisions concerning information security and the rationale behind them.

4.2 Related work

With respect to previous research, security arguments can be compared to safety cases [85, 86, 89], in that they summarize the reasons why, and the extent to which, a system is thought to be acceptably secure (or safe). Several techniques for modeling security arguments exist, some inspired from legal argumentation (Toulmin-like argumentation structures [23, 98]), others from formal methods (deontic logic [99], defeasible logic [22]). These are described in Sect. 2.3.3.

4.2.1 OpenArgue/OpenRISA

OpenArgue is an argumentation modeling tool featuring both a syntax editor and a graphical editor, which comes with the ability to derive an argumentation diagram from a textual specification [100]. OpenArgue assumes security requirements are known at the time of analysis and focuses on identifying ways by which these requirements could be invalidated. This means all arguments are linked to a specific security requirement. It benefits from syntax highlighting as well as a built-in model checker which can identify formal inconsistencies in the argumentation diagram. OpenArgue has a simplified Toulmin intra-argument structure consisting of a central *claim*, supported by *grounds*, the relevance of which is supported by *warrants*. However, OpenArgue allows specifying rather complex inter-argument relationships: arguments can rebut or mitigate one or more other arguments by challenging either their grounds or their warrants. This can lead to inter-twined graphical representations of the argumentation model that are hard to understand. This effect is amplified by the fact that the tool does not come with a custom editor but rather uses a generic Eclipse UML editor and thereby poses significant usability and scalability issues. Figure 4.1 shows the the

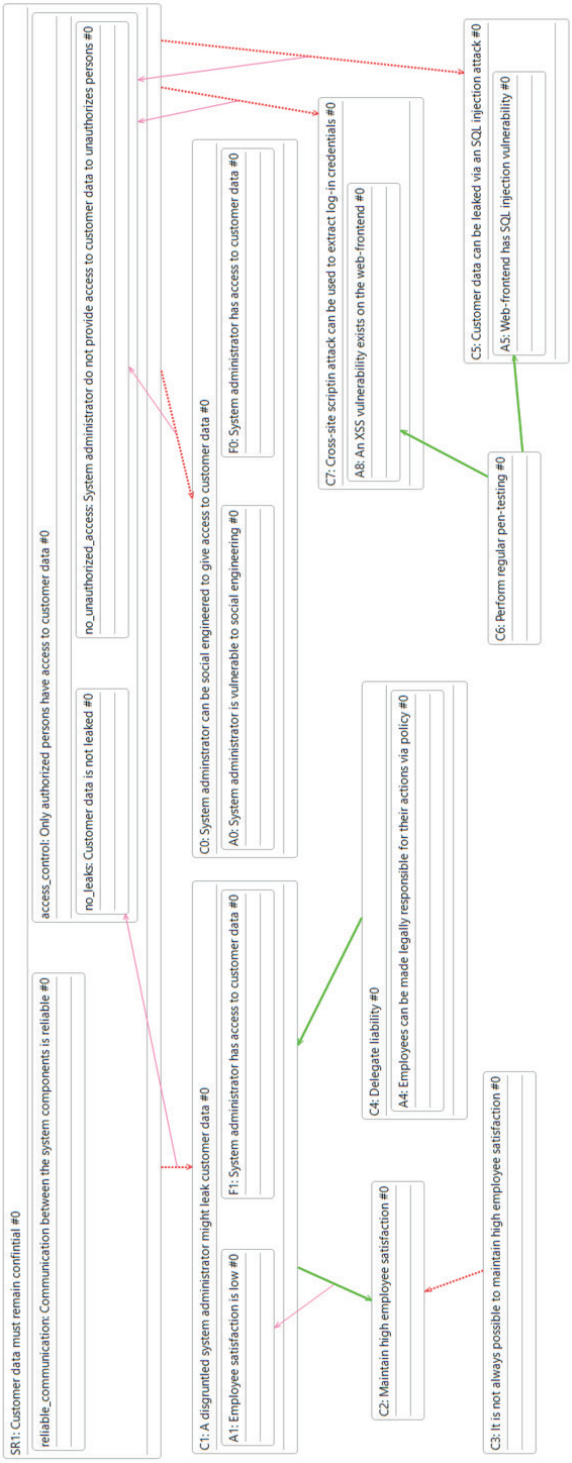

Figure 4.1: OpenArgue - sample assessment

argumentation model of a risk assessment linked to the requirement that customer data must remain confidential, constructed using the OpenArgue tool. The boxes are claims, and the inner boxes are. Red arrows show rebuttal links and green arrows represent mitigation relationships.

OpenRISA is an extension of OpenArgue which can, in addition, check the argumentation model against online knowledge bases and verify that the risks identified are valid rebuttals.

4.3 Proposed approach

The argumentation-based risk assessment proposed in this chapter is essentially an attempt to find a middle ground between Prakken's formal risk assessment approach [22] and the less formal Toulmin-based approaches ([98, 99, 155]). From the former, I keep the four types of inter-argument relationships and overall structure of an argumentation game but strip away most of the logical formalism that was part of the ASPIC framework Prakken's approach was based on. From the latter, I use the intra-argument structure, slightly simplified in order to mitigate complexity.

A major distinction from all previous approaches is the use of a tabular format as both input and storage of arguments where these are entered from top to bottom as they are introduced along the game (see Table 4.1). Unlike Prakken [22], the proposed approach does not differentiate between different types of counter-arguments and I assume the inference between fact plus assumptions and the claim to be implied. Unlike Toulmin-based approaches, the proposed approach supports four types of formally specified defeasibility relationship between arguments that allow for semi-automated reasoning as well as the computing of several types of reports or summaries. Furthermore, I introduce a way by which components can be referenced so that potentially conflicting arguments can be highlighted.

The assessors alternate between playing "defenders" and "attackers". Each "team" then takes turns formulating arguments either for or against the security of the system. In Table 4.1, I show a part of the risk assessment conducted during the second validation round, namely the assessment of a Home payments system for elderly individuals. The architecture of the system is shown in Fig. 4.2. Each row represents such a turn and describes one argument, starting with the attackers. Attacker arguments—marked by an A in the first column—describe Risks (in terms of possible attacks or vulnerabilities), while defender arguments describe ways to mitigate such Risks (by introducing countermeasures, transferring or accepting them). In Table 4.1, the first two pairs of rows describe an attack and countermeasure, respectively. The last columns of the D (defender) rows includes an R for ''reduction" and a T for ''transfer" which denote the type of risk decision: a risk reduction and a risk transfer. The last row describes an attack with no countermeasure, which signifies an accepted risk. The rebuts column shows which argument or part thereof is being challenged, the IN/OUT column shows the status of the argument (defeated or undefeated).

Player	Claim		Assumptions		Facts		Rebuts	Status	Flags
A/D	ID	STRING	ID	STRING	ID	STRING	ID	IN/OUT	T/R
A	C0	Listen in to Bluetooth: gather authentication or user data	A0	Bluetooth signal can be received outside	F0	Range of Bluetooth is 10m		OUT	
D	C1	Authentication data is encrypted	A1	AES encryption is good enough	F1	Bluetooth with 2.1 (AES) encryption	C0	IN	R
A	C2	User socially engineered to wire money	A2	Attacker can gain user's trust;	F2	–		OUT	
D	C3	Social Engineering is user risk	A3	–	F3	End-user agreement transferring liability for SocEng attacks	C2	IN	T
A	C4	User credentials can be stolen by peeking through the window	A4	Apartment located on bottom floor(s); Curtains open	F4	–		IN	

Table 4.1: Snapshot of an argument game for a Home Payments System (Sect. 4.5)

There is both an internal and an external argument structure:

- Internally, each argument consists of three parts: a claim, one or more assumptions and one or more facts. Each part is given a unique ID. The facts are either physical facts or known technical specifications of the target infrastructure. Assumptions are important parts of the argument that the assessors are not certain of. The claim is the core conclusion of the argument.
- Externally, there exist defeasibility relationships between arguments. That is, each argument can rebut (i.e. attack) one or more previous arguments by invalidating the claim itself or one of its assumptions. However, facts cannot be invalidated.

The two structures described above allow that each argument directly rebuts a part of any previous argument. The *Rebuts* column in Table 4.1 points to the ID of the part being rebutted. To represent the resulting states, I adopt part of Dung's abstract argumentation framework [156], in which each argument can at any moment in time be in one of two states: IN or OUT. Once an argument is successfully rebutted (that is, the opposing team proposes a valid counterargument), it becomes OUT, with the counterargument being IN. This can continue recursively, applying the following rules:

- An argument is IN if all its counterarguments are OUT. IN arguments have not been successfully defeated in the argument so far.
- An argument is OUT if it has a counterargument that is IN. OUT arguments have been successfully defeated in the argument so far.

To test out the effectiveness and applicability of these ideas in case studies, I implemented the method as a spreadsheet containing underlying formulas for recursively determining the argument state, which is represented using colours (red for OUT, green for IN).

I initially assumed the following loose mapping from argument states to Risk states:

- Attacker arguments that are IN at the end of the game are *accepted* (retained) Risks (e.g. the last row in Table 4.1)
- Attacker arguments that are OUT at the end of the game are *Reduced*, totally *eliminated* (e.g. the first two rows in Table 4.1), or *transferred* (e.g. middle rows in Table 4.1)

A secondary functionality is relating arguments to system architecture. The risk assessors start from the architecture of the Target of Assessment (and possibly its context) and enter a list of architecture components (nodes or connectors in the diagram) in the spreadsheet. Arguments are then automatically tagged with the labels of architecture components that they refer to as they are typed. This makes it easier for the assessors to identify potential conflicts in their statements as they are making them. Such a conflict occurs when a fact or statement is in contradiction with a previous fact or statement or if it is impossible to realize due to a previous statement about the same component. This labelling also helps avoid inconsistent views about the system among the assessors.

As stated above, facts cannot be invalidated, so it is important they are mutually consistent. Some facts may be properties of the Target of Assessment postulated by the defenders' team. This technique, therefore, assumes it is in the power of the risk assessors to make decisions about the ToA architecture, and that these decisions will be implemented. For the system developers these facts become requirements.

4.4 Research Strategy

Based on previous work, including a survey of common Risk Assessment methods, tools and frameworks [32], I tried to identify possible limitations of current approaches and the scope for improvement.

To ensure utility and usability, I iteratively validated and improved the artifact via three Case Studies. This has resulted in four iterations of the approach, each supported by spreadsheets:

1. Reduce complexity of ASPIC-based approach of Prakken et al [22] ⇒ 1st version
2. Improve the method after a Case Study of a Home Payments System with students ⇒ 2nd version
3. Improve the method after a Case Study of a Home Payments System with experts ⇒ 3rd version
4. Improve the method after a Case Study on Cloud-based Infrastructures ⇒ 4th and latest version; it is this version that is described above.

The Case Studies were chosen for their diversity and their relation to the TREsPASS research project[1], the project this research was part of. They are described in the following sections.

4.5 Case Studies 1 and 2: The Home Payments System

4.5.1 Case Description

This case study is centered on customer privacy protection and is owned by one of the project partners. The system consists of set-top boxes located in customer's homes, some centralized servers and personalized NFC-active bank cards. The set-top boxes are connected to the TV and allow the user to perform various financial operations (including but not limited to online banking, allocation of funds, payment of bills and online shopping) from the comfort of their home, by using the card as a means of authentication. A basic architecture diagram is shown in Fig. 4.2.

This case study is intentionally under-specified for two reasons: (1) the system, developed by Consult Hyperion, is still at the prototype stage; and (2) this allows for more freedom with regard to the design decisions that can be taken during the assessment. Thus, I am also looking to test if the approach might be used during product design phases, where Security Requirements elicitation is more crucial than after implementation.

4.5.2 Case-Specific Observations

This case study was performed twice: a pilot round in which the assessors were IT Security PhD students and a second round with senior Information Security researchers. Table 4.1 shows part of one of the assessments. The following observations were made during both sessions:

[1] https://www.trespass-project.eu/

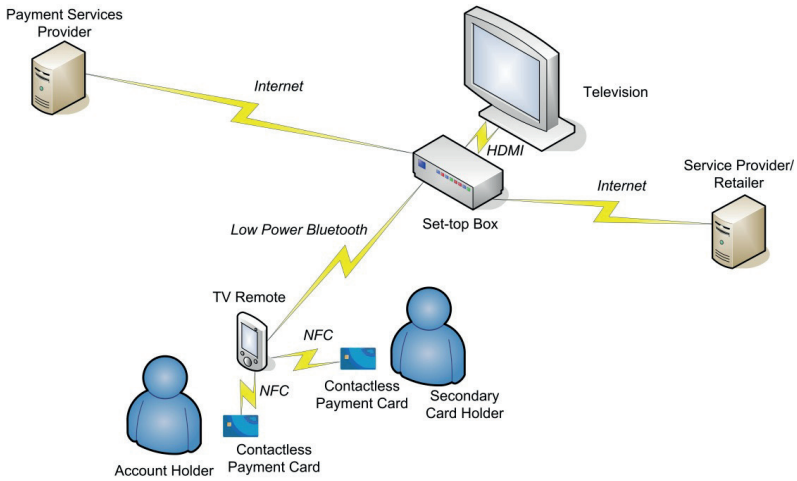

Figure 4.2: Home Payments System

Assumptions are non-exhaustive

Attacker assumptions are usually about the system and its context, while the defenders usually make assumptions about the attacker's profile and skills. A common problem with assumptions is that even when asking the participants to make them explicit, there are always some that remain hidden. Hidden assumptions cannot be explicitly attacked. This can be overcome by stating an opposing assumption as part of the counter-argument.

Reduced Risks are not the same as eliminated Risks

Attacker arguments which are out OUT at the end of the game signify eliminated or reduced risks (e.g. middle rows in Table 4.1). While for eliminated Risks, the attack is completely prevented, in the case of reduced risks, although the impact or likelihood have been sufficiently reduced, the attack itself might still be possible. To represent this, defender arguments that only partly mitigate the Risk (to an acceptable level) are flagged "R" (i.e. reduced).

Transferred risks should be clearly marked

Transferring Risks is a treatment option available during Risk Assessments (usually accomplished via insurance, end-user agreements, etc.). This means that the attack is still possible but liability for the potential negative consequences has been transferred. To support this, arguments that transfer the consequences of a Risk are flagged "T" (i.e. transferred).

Separate teams are better.

Allowing participants to take turns playing attacker and defender leads to them already having a counter-argument in mind when stating an argument, and thus subverts the argumentation dynamics of the game. Separate, fixed teams do not only mitigate this, but also instil a level of competitiveness between the two teams, resulting in better-formulated arguments.

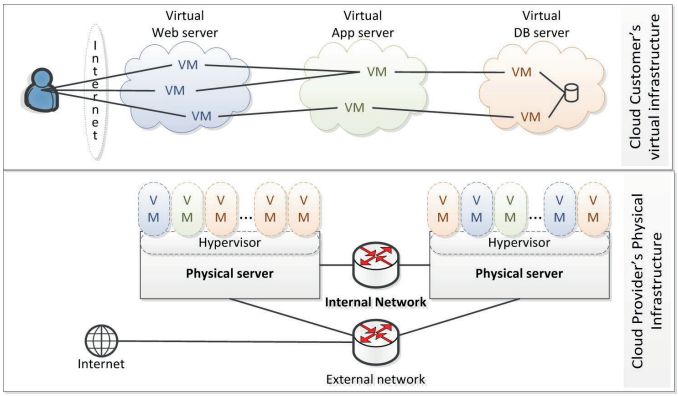

Figure 4.3: IaaS Cloud architecture

4.6 Case Study 3: The Cloud-Based Infrastructure

4.6.1 Case Description

Cloud-based implementations are now commonplace, with various service providers outsourcing their IT infrastructure to the cloud. Such virtualized infrastructures give rise to completely new categories of Risk, as well as new requirements with regard to identifying and mitigating such risks. As such, in order to explore the limits of applicability of the new method, a Risk Assessment was conducted in collaboration with IBM Research Zurich. As the target for assessment, a generic IaaS infrastructure was imagined. An overview of this infrastructure is shown in Fig. 4.3. It consists of two infrastructure layers:

- A physical layer, owned by the Cloud Provider. This consists of some servers, connected to each other via an internal network, and to the Internet via an external uplink. Each server runs a number of virtual machines belonging to the Cloud Customers which are managed via an interface called a Hypervisor. The entire physical infrastructure is managed by the Cloud Administrator, who can also access and manage the virtual machines (e.g. resource allocation) via an SSH connection to the Hypervisor console.

- A virtual layer, consisting of a large number of virtual machines, networks, and databases. Each Cloud Customer owns and controls a sub-set of virtual machines. A Customer Administrator is usually responsible for configuring and managing these for each Cloud Customer.

4.6.2 Case-Specific Observations

This case study was conducted with junior researchers on virtual infrastructures from IBM Zurich. The observations listed in the first two Case studies were confirmed during the third study. However, many new observations surfaced, mostly specific to Cloud-based infrastructures:

The introduction of a third, virtual, layer in-between the physical and logical domains, together with the dynamic nature of this layer makes assessing Cloud-based infrastructure a more complex undertaking compared to traditional Information Systems. When you add to this the introduction of the cloud provider as a new stakeholder and the mixed ownership of components across these three domains, the separation between system and context fades away and the amount of incomplete or missing information w.r.t the infrastructure explodes. In essence, any cloud scenario needs at least three, only partially overlapping views: the Cloud Provider's view, the Cloud Customer's view, and the end-user's view. Since there is usually little shared knowledge across these three stakeholders, and they have different goals, not only are there unique risks for each but even the common risks are ranked differently.

As described in Cloud Risk Assessment documents from ENISA [157] and the CSA [158], the fact that cloud customers do not usually have any control of or information about the physical infrastructure and resource allocation itself gives rise to a new host of vulnerabilities, ranging from resource exhaustion to collocation exploits.

Because none of the stakeholders have the ability to directly influence the components owned or managed by the others, countermeasures for Cloud-specific Risks are mostly implemented via SLA clauses. These commonly have expiration dates and even time constraints for implementation due to contractual periods, making them significantly different to the more technical countermeasures the technique was designed to handle. This is because there is no clear transfer or mitigation of the Risk. Instead, partial transfer of risk by means of such SLA clauses is common. The way these clauses are written and how compensation is specified determines the degree to which Risk is transferred. The proliferation of organizational entities with heavy reliance on SLAs dictating the relationships between such entities also makes assessments more complex, as well as making the method less applicable.

Despite the above difficulties, the participants were able to complete the Risk Assessment. However, the results looked significantly different from those of the previous case study. First of all, the attacker's facts were almost always missing. The same was true for the defender's assumptions. The attacker's assumptions usually implied the existence of a known vulnerability while the claims mostly referred to relationships between Vulnerabilities and Risks that are described in [157] and [158]. Therefore, it seems that while the method is flexible enough to be applied to such different scenarios, it does not offer significant added value in cases like this.

4.7 Discussion

4.7.1 Relation to Group Decision Support Systems

In each round, only one argument is described. This means the team has to agree on the argument being presented to the other team before submitting it. This suggests a parallel with group decisions support systems (otherwise known as group decision rooms), where a software tool mitigates the interactions between various stakeholders in order to help achieve a consensual opinion. Such concepts could also be applied to the approach described in this chapter, thus increasing its (perceived) utility by providing the users with extra functionality.

4.7.2 Relation to Design Rationale

The Security Requirements resulting from the assessment could be in the form of technical countermeasures, but can also specify policies or more general design decisions w.r.t the target system. Especially for systems under development or prototypes, these Security Requirements together with the claims they support are very similar to "design rationale" in the sense that they describe and motivate the desired properties of the system; in this case, in terms of the risks they are preventing.

The principles of capturing "design rationale as a by-product" of other related decision-making tasks, described in [159, Chapter 4.3] could be applied in order to evolve the approach in this direction as well as providing a more consistent method of storing the arguments.

4.8 Validity and Scope

In order to avoid problems related to participant bias, repeated testing and maturation, completely different panels of experts were used for each of the Risk Assessment sessions. None of the participants had seen or heard about the method before the session so as to further avoid selection bias. Only observations that have been confirmed in at least two of the three sessions are described in this chapter.

Furthermore, subsequent iterations produced improved effects, with the exception of the assessment of Cloud-based infrastructure. This leads us to believe that the effects are indeed produced by the tool, although the tool has limited applicability when dealing with (partially) virtualized and/or outsourced infrastructures.

Improper sample selection (that is, the participants are not sufficiently similar to the intended user population) has been mitigated by making sure that all participants in the assessments have at least a Master's degree in Computer Security or a related field, and are currently actively conducting research in the field of Information Security.

4.9 Applicability

The flexible argument structure and lack of security-specific features, in theory, make the method described in this chapter applicable - in theory - to a wide range of scenarios that are based on brainstorming and require trace-ability. Still, when validating the approach described in this chapter, and especially in the third case study of Sect. 4.6, some limitations were identified. These were partly due to the significant overhead of formulating each argument according to the template. When the statements or decisions consist of a central claim supported by an assumption and/or fact, there is little added benefit in attempting to decompose it any further. Such cases do not benefit from describing the defensibility relationship between the various arguments because there is no back-and-fourth rhetorical discourse or the argument's inner workings are of little significance to the conclusions.

The added benefit of using argumentation is mostly visible when the architecture is known because it allows traceability to components (further explained in Sect. 4.10). This was observed especially in the cloud case, where Risks, and even Attacks, Vulnerabilities are seldom described at a technical level and most countermeasures come in form of a policy or

SLA provisions. Thus, in such cases, there is no real added value compared to simply listing each applicable Risk and suggesting one or more countermeasures.

4.10 Conclusions and future work

This chapter described a tool-supported approach to collecting the rationale behind a risk assessment. This also helps structure the discussion. The tool can be used on its own, or in combination with checklists or external security consultants.

In all three case studies, most case study participants agreed that the method proposed in this chapter is feasible, provided that the obvious limitations of spreadsheet editors are overcome. In particular, they agreed bookkeeping is reduced compared to the ASPIC-based approach and could be further reduced by designing a custom software tool and using such an argumentation-based approach requires minimal training and no experience. For each session, a 10-minute presentation was sufficient for the participants to be able to start using the method. The average duration of each session was 2 hours.

However, one of the main findings is that, despite expectations, there does not seem to be a deep argument game during these assessments. This is because for each identified Risk or Attack, a suitable mitigation is found (either via a countermeasure or by transferring or accepting the Risk). The "attacker" team can either accept the defender's argument and move on or try to subvert the countermeasure by describing a slightly altered attack path. However, such an altered attack could, to all intents and purposes, also be viewed as a new round instead of a counter-argument (or rebuttal). So each round of the game contains at most two rounds: an attacker argument followed (optionally) by a defender counterargument.

Furthermore, the method is not affected by the presence or absence of quantitative values of likelihood and impact of Risk. This makes us optimistic that flexible Risk Assessment and/or Security Requirements elicitation tools, which can work with both quantitative and qualitative values, can be developed.

Architecture diagrams need to be more or less definite and known during the Risk Assessment as they provide crucial input. Furthermore, the scalability of Risk Assessment methods and tools increases inversely with the complexity and ambiguity of architecture. This also applies to the approach proposed in this chapter as, during the experiment, I noticed that if information about components is missing or incomplete, participants are unable to provide facts to support their claims, resulting in a drastic decrease in the utility of Toulmin-like argumentation structures.

I observed that arguments, in the context of Risk Assessment and/or Security Requirements Elicitation, take most often take the form of traceability links between the requirements, vulnerabilities, components and attacker profiles. In this respect, they outperform the use of checklists. However, the added overhead raises the question of what level of traceability is necessary and sufficient and how that level can be provided without overburdening the process. While the approach described in this chapter come closer to this desired equilibrium than ASPIC and OpenArgue for some types of assessments, more research is required in order to more precisely determine at what level the method's cost is outweighed by its benefits.

As the features required go beyond what is normally achievable via spreadsheets, a dedicated software tool would likely bring significant increases in the usability and scalability of the approach, while potentially adding extra functionality. I discuss this in the following

chapter. Furthermore, due to limited time, only two real-world cases were used for validating the argumentation based risk assessment method described in this chapter. To properly assess the potential benefits and limitations of the process, as well as its wider applicability, more case studies need to be conducted on other cases. Finally, the relationship between the proposed approach and current approaches in the field of design rationale and group decision support systems, highlighted in Sect. 4.7.2 and Sect. 4.7.1 below, is worth exploring. Automated reasoning, for instance, may extend the applicability and scope of argumentation-based risk assessment. Other ideas and approaches from these more advanced fields could be used to develop a dedicated argumentation-based risk assessment tool.

5

Collaborative risk assessment supported by a shared argumentation model

Based on *ArgueSecure: Out-of-the-box Risk Assessment* [24]

Most established security risk assessment methodologies aim to produce ranked lists of risks. But ranking requires quantification of risks, which in turn relies on data which may not be available or estimations which might not be accurate.

As an alternative, in the previous chapter, I discussed argumentation-based risk assessment. In this chapter, based on practitioner feedback, I introduce a second iteration of this method accompanied by two dedicated tools: an online, collaborative web-portal and an offline version. In the first iteration, the focus was on end result: a list of security requirements supported by structured arguments. This second iteration uses a simplified argumentation model to promote a participative and collaborative risk assessment process. The result is a new risk assessment framework I call ArgueSecure, which is geared towards the collaborative construction of a graphical risk argumentation model. It uses a tree structure which intuitively encodes argument traces, therefore maintaining traceability of the results and providing insight into the decision process.

5.1 Introduction

The argumentation-based risk assessment approach described in Chapter 4 was successful at encoding the links between components, risks, and countermeasures needed to understand the rationale behind an information security risk assessment. However, one of its main limitations was that it was not usable in real-time during a risk assessment session involving multiple stakeholders.

But assessing the risks pertaining to an information system often requires the involvement of multiple stakeholders, such as internal IT specialists, cyber-security experts, and budget-responsibles to decide on mitigations. During brainstorming sessions, stakeholders may collaboratively identify risks (including new and hybrid risks) [160] or discuss and agree on security requirements [94]. Furthermore, studies show that often project teams – rather than in-house experts or external consultants – take part in risk analysis meetings and that qualitative approaches are generally preferred by such groups [161, 162].

In this chapter, I refine the approach of Chapter 4, with the hopes of improving its usability in collaborative risk assessments. To this end, I propose ArgueSecure: a light-weight, flexible, qualitative risk assessment and security requirements elicitation framework, consisting of a set of dedicated tools and an associated method. The method employs a similar argumentation structure to the spread-sheet approach described in the previous Chapter and takes into account the lessons learned from the Case Studies it was applied to. The tools are designed to be capable of capturing and encoding the key arguments put forth during a qualitative risk assessment in real-time. These arguments serve a dual purpose. Firstly, they provide support for the results of the assessment, whether risks or countermeasures. Second, they promote reusability and can be used to construct a knowledge base of such risks and countermeasures.

The ArgueSecure tools come in two flavors: a Web server that can be deployed as an intranet or Internet portal and an offline Java tool. Both tools are open source, and work out-of-the-box with minimal configuration: the Web server is available as a deployable VM while the Java tool is provided as a single-file portable executable with import and export functionality. The goal of the offline version is to provide bookkeeping for risk assessment or security requirements engineering sessions. The goal of the online version is to, in addition, allow stakeholders and experts to engage in a risk assessment without being in the same room and even without being available the same time while maintaining full traceability between security requirements and risks. The development of an online version is based on the assumption that the stakeholders whose input is required for eliciting security requirements might not be available to participate in a dedicated session. In this chapter, we describe the evolution of the ArgueSecure method and present our experience with developing and evaluating these tools.

The remainder of this chapter focuses on the lessons learned in iteratively developing and evaluating these tools and the underlying framework. Section 5.2 introduces the first version of the dedicated ArgueSecure software tool and draws conclusions with regard its applicability, utility, and usability; Sect. 5.3 describes how these conclusions led to a re-design of ArgueSecure as an online portal and how this portal was again evaluated on the same criteria; Finally, Sect. 5.4 summarizes the lessons learned throughout developing these tools and discusses implications for practice as well as the potential for future work.

5.2 Collaborative risk assessment with ArgueSecure offline

The first dedicated ArgueSecure tool (now ArgueSecure-offline) implemented the same argumentation-based method as our spread-sheet based tool - described in Chapter 4 - and was intended to be used during dedicated security requirements elicitation sessions. The application is built to be usable in real-time during a session and the GUI is designed to work on low-resolution screens so that it can be easily projected. However, unlike the spread-sheets described in Chapter 4, each risk assessment (i.e. each list of risks and mitigations) now follows, tree-like structure:

Cat : A category of risks

- **R1** : A risk
  - ⚔ Claim made by an attacker about the existence of an attack path.
    - **A** An assumption of the claim.
    - **A** Another assumption of the claim.
  - ⛨ Claim made by a defender, that partly or completely defeats the attacker's claim by pointing out that an attacker's assumption is probably, or certainly, false.
    - **A** An assumption of the defender's claim, e.g. about a mitigation that already exists or that will be implemented.
  - ⚔ Renewed claim of the attacker that bypasses the defenders' argument.
    - **A** An assumption of this renewed claim.
- **R2** Another risk.

Cat : Another category of risks. [etc.]

This structure provides an visual representation of the identified risks and also shows the relationships between risks, attacks, and mitigations. In line with previous work, each risk is treated as its own argument game with "attackers" and "defenders" taking turns until the risk is either accepted, reduced, eliminated or transferred [23]. Each turn consists of a single claim which - except for the first - rebuts a previous claim. Defender's claims can refer to ToA[1] components or architectural decisions that reduce or eliminate a risk, but can also refer to decisions, disclaimers, or policies that transfer the risk or potential loss to another party through a contract (e.g., a hold harmless clause) or to a professional risk bearer (e.g. to an insurance company or to a customer). Transfer claims are marked using an arrow instead of a shield.

The buttons and text are large enough to be visible from across the room when projected on a large screen (see Fig. 5.1 for a screenshot). The tool is designed to be usable exclusively via the keyboard so as to support real-time book-keeping of the session. Save/load functionality allows assessments that span multiple sessions and even distributed assessments (by sending the file via e-mail, for instance). Together with various exporting and reporting features, it also supports reusability and dissemination of elicited security requirements.

[1] We use Target of Assessment (ToA) to refer to the software, system, or project which serves as the subject of a risk assessment

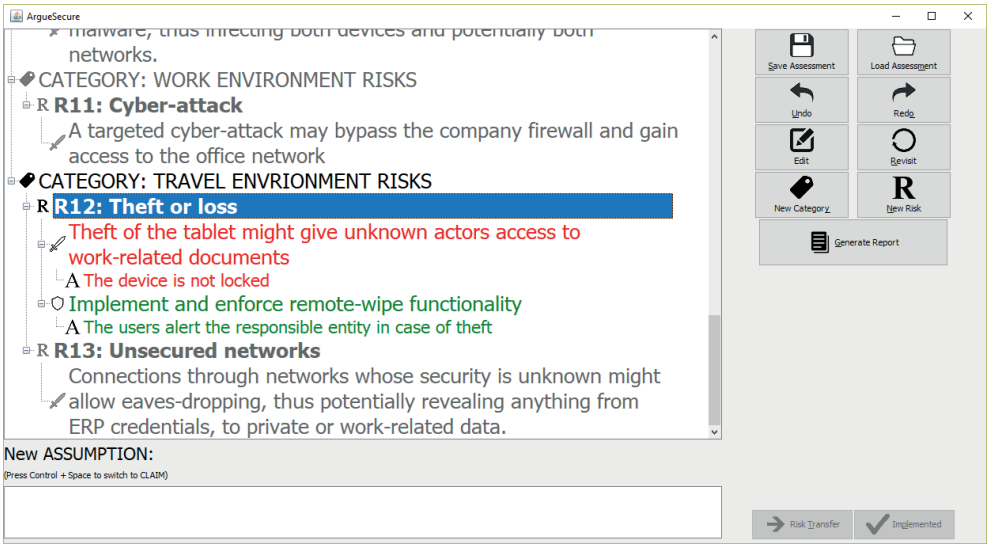

Figure 5.1: Screen-shot of ArgueSecure-offline

5.2.1 Deployment and usage

The application is provided as a single, self-contained executable.

Conducting an ArgueSecure RA requires little preparation. Any number of stakeholders, domain experts and/or security experts can participate but should be split up into two teams: attackers and defenders. The method assumes the participants possess pre-existing knowledge of the Target of Assessment. Ideally, but not mandatory, some sort of system model or diagram should be agreed upon by the participants. The preferred workflow is as follows:

1. Create a new category and give it a name or choose an existing category
2. Create a new risk under this category, and provide a brief name/description of it
3. Each risk starts with an attacker argument, describing an attack path or refining the risk. Each argument consists of a claim, supported by one or more assumptions.
4. Each attacker argument may be countered by a defender argument, describing a counter-measure mitigating the risk.
5. This back-and-forth rhetoric can continue until:
   (a) The attacker team is unable or unwilling to counter the last defender argument. This means the risk has been eliminated.
   (b) The defender team is unable or unwilling to counter the last attacker argument. This means the risk (or residual risk) has been accepted.
6. If other risks can be identified under this category, go back to step 2.
7. If a new category of risks can be identified, go back to step 1.

At any time during the assessment, defender arguments can be visually marked as ”implemented” if they describe existing risk countermeasures and/or ”transfer” if they describe a risk transfer.

5.2.2 Validation and lessons learned

We evaluated the usability and utility of the offline tool by a pilot study, a focus group and a case study.

5.2.2.1 Pilot study

The pilot study was carried out within the PISA[2] (Personal Information Security Assistant) project to obtain an initial overview of risks faced by employees working remotely or from home. After instantiating an assessment with several known risks, the tool, together with this draft assessment was sent via e-mail to various domain experts which were asked to complete the assessment as they see fit. This was because relevant stakeholders were unavailable at the same place at the same time (P0). Unfortunately, most e-mail servers block executable files and many computers do not have the Java Virtual Machine installed (P1). This has significantly hindered adoption to the point of falling back to a simple text editor. The participants who did contribute always added new risks to the assessment, and never elicited attacker claims rebutting previous defender claims (P2). Furthermore, assumptions were never substantiated (P3). Essentially, the tool was used as a running list of potential risks.

5.2.2.2 Focus group

The focus group consisted of security stakeholders of a major Dutch bank and was used to gather feedback with regard to the usability and utility of the ArgueSecure-offline tool. The goal was to collaboratively perform an ArgueSecure risk assessment of a new home banking authentication device. However, since planning a dedicated session with both security experts and responsible management was not possible (P0), the assessment was conducted in two phases: first, security experts created a list of risks and attacks; then, during a shorter meeting with bank stakeholders, decisions were made on which countermeasures to implement and which risks to accept. We observed that, similar to the pilot study above, renewed claims against elicited defences were rarely introduced (P2). While the tool was generally perceived as useful, participants also indicated that usability and scalability become issues as the depth of the tree increases (P4). However, unlike the pilot study, we were physically present during the meeting and therefore able to encourage participants to express their unstated assumptions.

5.2.2.3 Case study

The case study was aimed at identifying the limitations of the ArgueSecure approach when applied by two security practitioners to a fictional scenario involving ATM security. A facilitator was present to manipulate the tool as the two brainstormed about risks and countermeasures. We observed that despite instructions by the facilitator, the division into attacker and defender teams was not respected, with both participants eliciting attacks as well as defences (P5).

[2] http://scs.ewi.utwente.nl/projects/pisa/

After a one hour session, the participants were also asked to fill out a questionnaire about their experience. The restrictive cardinality of the approach was highlighted as the main weakness. Namely, the inability to (1) map multiple attacks to the same risk (P6), (2) have a risk belong to zero or more categories (P7) and (3) state that an elicited defence mitigates several attacks or risks (P8).

5.3 Web-based risk assessment with ArgueSecure online

Solis et al. [163], Cheng et al. [164] and Seyff et al. [165] claim that requirements engineering is becoming more and more a collaborative effort by distributed stakeholders. We combined this general insight in the future of RE with our experience with phenomena P0 and P1 by developing ArgueSecure-online: a distributed, web-based risk assessment and security requirements elicitation portal with real-time collaboration functionality. Its goal is to allow busy stakeholders to contribute to the security requirements elicitation process of a software and/or system in a flexible manner, (thereby avoiding phenomenon P0) and without having to download an executable (thereby avoiding phenomenon P1). The tool allows users to collaboratively or privately build and maintain structured lists of risks and mitigations for software and/or systems. The tool maintains the structure of the offline version, with some key differences:

- A defence claim can no longer be rebutted as this was rarely done (P2) and posed scalability issues (P4);
- Assumptions have been dropped as they were rarely used and commonly misused (P3);
- The separation between attackers and defenders has been dropped: any participant can elicit either risks and attacks or defences at any time (P5);
- Each risk can now consist of multiple alternative attacks (P6);
- Top-level categories have been dropped and replaced with node-level tags to further decrease the depth of the tree and allow filtering (P4) while permitting many-to-many mapping of risks and even individual attacks or defences to categories (P7);
- A single defence can now mitigate several attacks (P8).

A risk assessment in ArgueSecure-online also follows a tree-like structure. Similar to the offline version, the root node is the assessment itself, further decomposed into risks, then attacks and finally defences. The tree is only three instead of five levels deep due to categories becoming tags and assumptions being dropped:

Risk : a perceived weakness (such as a vulnerability) or threat (such as undesirable situation) of the system considered by the risk assessment. A risk is associated with a single risk assessment.

Attack : a specific attack path associated with a risk (such as a method of exploiting a vulnerability or producing undesirable effects). An attack is associated with one or more risks.

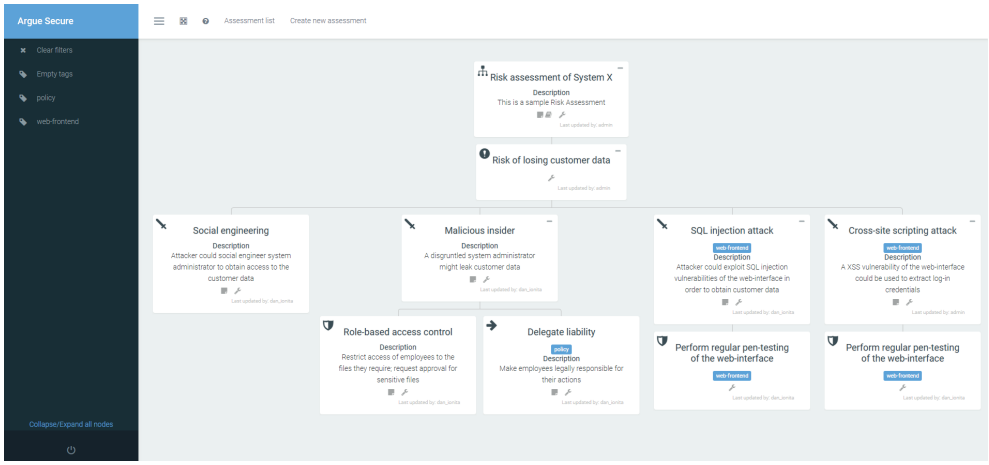

Figure 5.2: Screen-shot of ArgueSecure-online

Defence : a security requirement (such as an architectural change or policy measure) that mitigates specific attack paths. Additionally, defences may refer to the transfer of the risk to a third-party and may be marked accordingly. A defence is associated with one or more attacks.

Each node, independent of its level, consists of a name, accompanied by a description and a set of optional notes.

Although a single defence may mitigate several risks thereby transforming the tree into a cyclic graph, for visualization purposes this graph is presented as a tree: leaf nodes which have two parents are simply duplicated. However, changes to any instance of the defence will propagate to all other instances.

A screen-shot showing a sample risk assessment is shown in Fig. 5.2 and a demo version can be found online at `ArgueSecure.ewi.utwente.nl`.

5.3.1 Deployment and usage

The ArgueSecure application is deployable both as a VM and as stand-alone web-server. The source-code, VM images, and configuration instructions are all freely available on GitHub[3].

Once deployed - locally, on an intranet or on the Internet - the application can be centrally managed via a built-in administrator account.

Regular users can log in individually and are able to create public or private assessments, as well as contribute to the public assessments of other users. Multiple users can contribute to the same assessment simultaneously. Changes are visible in real-time, both as updates to the graphical model as well as notifications. The tool also provides export functionality which prints the assessment as a bulleted list.

[3]https://github.com/hitandyrun/arguesecure-online/

5.3.2 Validation and lessons learned

We evaluated the online version by means of an observational live study and a second focus group.

5.3.2.1 Live study

The live study was carried out entirely online, throughout the duration of REFSQ 2016[4] (a requirements engineering conference) [19]. Participants were provided with individual, anonymized access credentials which they could use to log into a private deployment of the ArgueSecure portal at any time. They were asked to imagine risks related to organizing and participating in a conference and fill in a questionnaire evaluating the tool. They were given no preceding instructions on how to use the tool.

Unfortunately, the questionnaire received only 6 responses. Most respondents claimed to have some experience with risk assessment. Participants, overall, found the interface suitable for brainstorming about risks, although some did point to more flexible alternatives such as Freemind[5] (offline mind-mapping software) and Trello[6] (online project management tool). While, on average, they rated the interface's understandability on first use with a 3 out of 5, after understanding the basic functionality, ease-of-use was scored with a 4 out of 5.

5.3.2.2 Focus group

During the one-hour focus group session, a total of eight information security researchers connected to the ArgueSecure portal using various devices such as laptops, tablets, and smartphones. Each participant contributed to the assessment individually and simultaneously, without being given any instructions in advance. Despite several opportunities for improvement, all participants rated the interface as easy or very easy to understand and easy or very easy to use, with 4 out of 6 claiming it was very easy to perform desired tasks after only a few minutes. Furthermore, the tree representation was generally seen as suitable for brainstorming about risks. Finally, during the focus group, we have seen tags used to map threat agents to risks or to categorize risks based on relevant factors for the particular application.

5.4 Conclusions and future work

Both the offline and online version of the tool, just as the earlier spreadsheet-based version, have been used successfully in real-world risk assessments, but the online version has the advantage of solving problems P0-P8 experienced with the offline version, and both on-line and offline versions have the advantage with respect to the spreadsheet-based tool of ease-of-use and of being more scalable in the number of risks.

Our experience with argumentation-based risk assessment does teach us some more general lessons: Formalized argument structures may be present in other domains, such as in the legal domain, analyzed by Toulmin [77], but they are not used by any of the security or other experts with whom we did risk assessments. Assumptions were not stated, and claims once

[4] https://refsq.org/2016/conference-program/on-line-experiment/
[5] freemind.sourceforge.net
[6] https://trello.com

defeated were not revisited. Warrants and backings were present, but every expert found it a waste of time to document these in the RA as they were common knowledge. A modal qualifier indicating the strength of the support for the claim was not given; here the shared understanding was that no claim was supported with 100% certainty and that the degree of certainty could not be quantified, nor was it worth the effort to estimate it. The decision to invest in a mitigation was made on subjective, unquantified assessments of the severity of a risk and the available budget for mitigations.

This means that an argument stored in ArgueSecure consist of a claim supported by grounds. The grounds, in turn, are stated in terms of known vulnerabilities, the architecture of the ToA, and assumptions about attackers' capabilities. Our conclusion is that approaches to argument-based security that require elaborate argument structures, such as that of Toulmin are not usable in practice.

At the same time, the traceability between mitigations and the grounds for these mitigations was found important by all experts. However, while assumptions can help qualify and clarify a claim, most experts did not make these explicit. The tool cannot check for mal-formed or incomplete arguments, but a good facilitator can help in externalizing tacit knowledge without hampering the process.

The version of ArgueSecure described in this chapter is a result of several iterations of development and validation. It is a derivative of the argumentation-based risk assessment method described in Chapter 4 [23] has evolved accordingly. While this chapter is focused on presenting the lessons learned throughout this process as well as the evolution of the toolkit, it could be interesting to distill a formalized risk assessment method based on our findings, ideally after conducting a larger scale observational case study with the help of the ArgueSecure framework.

Part IV

Value modelling

6

Quantifying business risks using value models

Based on *Using value models for business risk analysis in e-service networks* [14].

Commercially provided electronic services commonly operate on top of a complex, highly-interconnected infrastructure, which provides a multitude of entry points for attackers. Providers of e-services also operate in dynamic, highly competitive markets, which provides fertile ground for fraud. Before a business idea to provide commercial e-services is implemented in practice, it should therefore be analysed on its fraud potential.

This analysis is a risk assessment process, in which risks are ordered on severity and the unacceptable ones are mitigated. Mitigations may consist of changes in the e-service network to reduce the attractiveness of fraud for the fraudster, or changes in coordination process steps or IT architecture elements to make fraud harder or better detectable.

In this chapter I describe a technique for quantitative risk analysis which builds upon *e*^{3}*value* business value models. This allows for impact estimation as well as understanding the attacker's business cases. I show how the *e*^{3}*value* ontology — with minimal extensions – can be used to model analyse known telecommunication fraud scenarios. I(also show how the approach can be used to quantify infrastructure risks. Based on the results, as well as feedback from practitioners, I discuss the scope and limits of generalizability of the approach.

6.1 Introduction

e-Services, commercial services delivered electronically [166], are of vital and increasing importance to society. Examples are internet provision services, telephony services, email services, on-line delivery of music or other content, e-banking, on-line booking, etc. These services are delivered fully electronically, as opposed to many other 'physical' services such as a haircut at a barber. In this chapter, I will once again use telephony services as running examples.

The delivery of e-services is done via an *Information and Communication Technology (ICT)* infrastructure. For instance, modern telephony connections are handled by a complex technical architecture and rely on several information systems, e.g. for billing or call management. Technical vulnerabilities in such infrastructures may cause great concern [167].

However, since e-services are commercial offerings, they have *commercial* vulnerabilities in addition to technical ones. For instance, it is possible to register a telephony subscription using the identity of someone else (e.g. by providing a false proof of identity in the subscription process), resulting in calling for free.

These problems are exacerbated in highly competitive e-service markets such as telecom and on-line content provision, where service providers struggle to increase their market share by pushing new, increasingly flexible service packages with low and sometimes even negative margins. In an effort to reduce time to market, service providers might not have the time or resources to fully assess the potential for loss of each new service package. However, due to the increasingly complex and interconnected nature of e-service provision, these plans often contain loopholes which malicious customers might abuse in order to reduce their costs or even turn a profit. Traditional heavy-weight Governance, Risk, and Compliance (GRC) frameworks are therefore of little use to analyse fraud potential: their models are focused on the technical layout while established methods are mostly concerned with confidentiality, integrity and availability issues and may take days or weeks to apply [32].

Fraud scenarios are best described in terms of value flows (such as money or services) and often disregard the underlying technical infrastructure. As such, in order to quickly estimate the potential for loss of a given service package, as well as identify thresholds useful in drafting "Fair-use policies" and fraud detection heuristics to limit this loss, we need a method capable of modelling business networks. As explained in Chapter 2, the e^3*value* ontology [168] for exploring new e-business ideas, is very well suited for this purpose. To make e^3*value* suitable for risk analysis, I propose the e^3*fraud* extension which introduces three fraud heuristics and differentiates between "ideal" and "sub-ideal" value model. e^3*fraud* conceptualizes risks in a *model*-based way, using a business oriented terminology. This ensures that the approach is usable by IT-oriented stakeholders, while keeping business concerns in mind. I present examples of fraudulent behaviour in the telecom industry and show how to model and analyse them using e^3*fraud*. Furthermore, I show how the approach could be used to quantify the commercial impact of IT risks.

Business value modelling frameworks such as e^3*value* are designed to estimate the profitability of an e-service package, for all actors involved. Consequently, the value-driven risk modelling and analysis approach described in this chapter are based on the assumption that that security risks are in the end *business*-related [169], as they always result in loss of money for one or more entities. It is also worth noting that the proposed methodology explicitly recognizes the notion of a *value constellation* [170]. Many e-services, in fact, are value

constellations because they require multiple profit-loss responsible actors to collaborate in order to produce value for the customer. For example, in the telephone domain, there is a caller, a callee, one or more telecommunication companies (e.g. for transit traffic), parties for billing and selling of prepaid cards, etc..

This remainder of this chapter is structured as follows: In Sect. 6.2 I summarize the approach taken to produce the results presented in this chapter. In Sect. 6.4 I outline the *e³fraud* approach to analyse *commercial* risks in networks of e-services using a case study provided by a telecommunication operator. In Sect. 6.6 I show how the *e³fraud* approach could be used to quantify known *infrastructure* risks. Section 6.8 tackles some of the issues encountered by the authors.

6.2 Research methodology

The approach undertaken follows the traditional Design Cycle [110]. Partners from the telecom industry put forth the need for an approach to conducting a lightwieght risk analysis of new service packages before they hit the market and provided several fraud scenarios for analysis. Investigation of these scenarios revealed that they could be commonly described solely in terms of value exchanges amongst the actors.

Based on previous experience in creating value models and doing profitability analyses of a value constellation, I selected the *e³value* framework as a starting point. The *e³value* approach models a network of end users and enterprises who exchange things of economic value with each other. However, *e³value* is designed for mutually beneficial value models. So I iteratively extended the *e³value* ontology and toolkit so as to accommodate the scenarios in questions (see Section 6.5.3) and provide meaningful output (see Sect. 6.5.4), respectively.

The long-term goal of this research is to facilitate automatic identification, modelling and analysis of business risks related to e-service provision by software tool support. To this end, I used two real-life case studies to demonstrate and validate the modelling conventions and analysis approach. Only one of these cases is discussed in this thesis. The others are discussed in full in D7.3.1 of the TREsPASS project [171]. For further validation, I demonstrated the approach to a telecom provider and gathered feedback about their perception of the potential usability and utility of the approach in practice, which I discuss in Sect. 6.7.

Most of the results shown in this chapter were produced using software tools: for the creation of the models, I used an existing editor (see `e3value.few.vu.nl`) and for running the *e³fraud* analysis I developed a Java extension was created.

6.3 The *e³fraud* ontological extension

The *e³value* method described in Sect. 2.3.2 assumes all actors behave as planned. In other words, *e³value* models assume an ideal world, in that they represent a value constellation with actors that only behave as assumed by the model. The reason for this assumption is that *e³value* is first and foremost intended for business development; during workshops with stakeholders it is already sufficiently difficult to understand which actors are involved, and what they exchange of economic value with each other, without having to consider actors who behave dishonestly. Because of this, I call *e³value* models *ideal* value models.

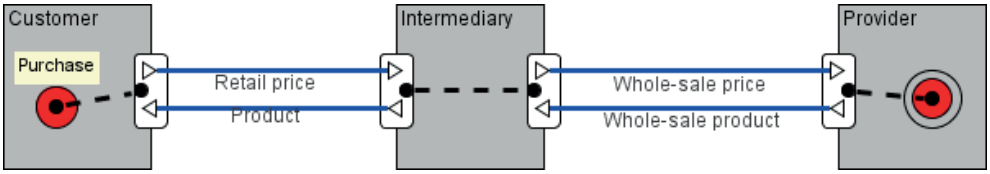

(a) Ideal e^3*value* model

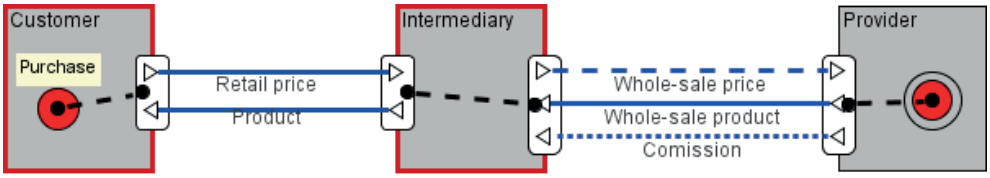

(b) Non-Ideal e^3*fraud* model

Figure 6.1: The e^{3}fraud *extension - graphical notation*

Since the launch of e^3*value*, an extension to model sub-ideal scenarios, in which not all business actors are to be trusted, has been proposed [172]. This extension, named e^3*control*, introduces the concept of *sub-ideal* value models. In a sub-ideal value model, the reciprocity constraint of e^3*value* is dropped in order to be able to represent situations where one or more value transfers do not occur (e.g. customer not paying for a product) or occur incorrectly (e.g. paying an insufficient amount of money) [173]. However, this is insufficient to model more complex types of fraud, such as revenue sharing fraud [14].

To this end, I introduce the e^3*fraud* extension to e^3*value* which allows for the construction of value models where actors violate agreements, contracts, or the law, which I call *non-ideal* value models. In e^3*fraud*, the focus is on the how these affect the expected revenues of other actors in the network. The e^3*fraud* approach extends the e^3*value* ontology with the concept of *collusion*, *non-occurring* value transfers, and *hidden* value transfers. A collusion consists of two actors that act as if they were financially independent, but which are in fact pooling their budgets, revenues, and costs. A non-occurring value transfer is a transfer of value objects that in the e^3*value* model is expected to occur but does not (NB: in e^3*control* it was also possible for transfers to occur partially, but that is not relevant for fraud). A hidden transfer is a transfer of value objects which are unexpected or otherwise hidden from the rest of the value network. The latter two are represented as dotted and dashed lines, respectively, while collusion is represented as a container surrounding each colluding actor. In Fig. 6.1, a hidden "Commission" is paid out with every wholesale purchase, the "wholesale payment" does not occur, and the customer is colluding with the intermediary.

6.4 The e^3*fraud* approach to analysing business risks

The e^3*fraud* approach takes as input an ideal business model and produces a set of sub-ideal business models. Each sub-ideal business model represents a business risk, for which graphs can be generated showing the loss/gain of each Risk.

The e^3*fraud* approach consists of three steps:

1. Construct the ideal business value model in *e³value*, showing the e-service at hand in terms of expected economic value creation and distribution
2. Construct one or more sub-ideal models in *e³fraud*, showing possible fraud scenarios in terms of economic value
3. Analyse financial feasibility and financial impact of the fraud

Steps 1 and 2 are also proposed by Kartseva et al. [173]. However, where Kartseva et al continue with proposing solutions to possibilities to prevent committing a fraud, the *e³fraud* analyses in step 3 the financial feasibility of the fraud for the fraudster, and the financial impact of the fraud on the ToA. In other words: the attack should be profitable for the attacker; otherwise, the attack is not financially feasible. In addition, the attack should be costly for the ToA; otherwise, countermeasures are not financially feasible. This allows stakeholders to assess the severity of a fraud scenario represented by the sub-ideal model and helps decision makers choose which scenarios need to be mitigated.

Furthermore, in Kartseva models [174], sub-ideal model behaviour is represented by value transfers that do not occur (e.g. a customer not paying for a product), or occur wrongly (e.g. paying an insufficient amount of money). *e³fraud* adds the notion of *hidden* transfers: fraudulent behaviour might involve value transfers that some (honest) parties do not expect or cannot observe, but of which they later experience the financial effects. This implies that an *e³fraud* model now takes the *perspective* of an individual enterprise or customer.

In this chapter, to demonstrate the modelling approach, sub-ideal models are constructed manually, but in the following chapter, I will introduce tool support for automatic generation and ranking of sub-ideal models.

6.5 Case study

6.5.1 Scenario Description

In this section, I explain the *e³fraud* approach by means of an easy to understand example, from the field of telecommunication/telephony.

A simple example of fraud in the telephony sector is Revenue Share Fraud (RSF) and involves setting up revenue sharing agreement with one provider, and a flat-rate (unlimited) subscription with another, and then calling yourself. This triggers the payment of interconnection fees from one provider to the other, thus resulting in a transfer of economic value between the providers. Depending on the scale of the operation and the detection capabilities of the provider, fraudsters could pull in up to several million dollars over a weekend [175]. I define Telecom misuse as the contracting or consumption of telecommunication services in a manner that is not in line with the service provider's expectations. Fraud is then any instance of misuse as previously defined with the explicit goal of obtaining financial rewards.

6.5.2 Construction of an ideal business value model

First, I construct first an ideal business value model. Such a model shows what actors transfer in terms of economic value if all actors behave *honestly*. I re-iterate the core concepts of *e³value* as I go along.

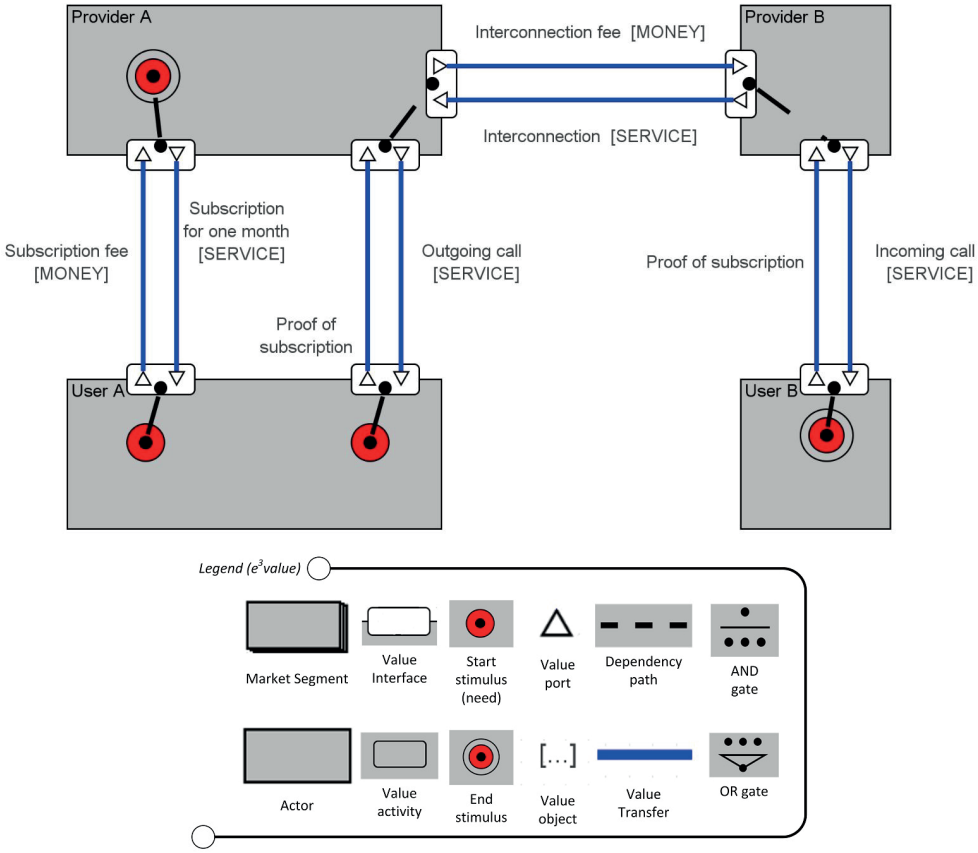

Figure 6.2: Ideal model: User A calls user B

Figure 6.2 shows the ideal business value model for a flat-rate mobile phone subscription. In *e³value*, actors are profit-and-loss responsible enterprises, non-profit organizations, or end-users. In this specific example, the telecommunication providers (provider A and B) and their customers (user A and B) are actors. The *e³value* language also has the concept of a *market segment*. A market segment groups actors that assign the same economic value to value objects and engage in the same transactions. Had I wanted to add more customers to the model, they could have been grouped together into a single market segment. For simplicity, I do not use the notion of market segment yet.

Actors exchange things of economic value with each other, called *value object*s, via *value transfers* (visualized as a straight line *between* actors). Value objects are physical objects or outcomes of services that are experienced by at least one actor in the model as of economic value. In Fig. 6.2 user A obtains a monthly flat-rate subscription ("Subscription for one month") with Provider A and in return for this, the user pays Provider A an amount of money on a monthly basis ("Subscription fee"). The flat-rate subscription entitles the users to perform unlimited telephone calls to any other number. In return for presenting proof of this subscription to the provider, the provider delivers its service, which is a telephone call.

In many cases, the caller and the callee do not have a subscription with the same provider, but rather with two providers, in Fig. 6.2 providers A and B. So, to create a telephone connection initiated by user A to user B, provider A has to interconnect with provider B, since provider B is the operator user B has a subscription with. In other words, provider B delivers an interconnection service to provider A. This service is of value to provider A, because otherwise provider A could not create telephone connections outside its own network.

User B has his own contractual agreement with Provider B. However, this ideal model is built from the perspective of Provider A for whom the structure of this agreement is neither known nor observable and thus not represented. The only transaction between User B and Provider B which Provider A can observe is the telephone call.

An *e³value* model shows how actors do business which each other in a *contract period.* This is a period described by the contracts that describe the value transfers among actors shown in the diagram. An important property of an *e³value* model is *economic reciprocity*. Figure 6.2 shows various *value interface*s, containing *value port*s, transferring value objects (see the legend). The notion of value interface represents economic reciprocity, meaning that *all* value ports transfer objects of value, or *none at all*. For example, when provider A obtains interconnection from provider B, provider A will pay, as described by the contract, in the contract period. The same holds the other way around: If provider B is paid, provider B provides the interconnection service as described by the contract.

Finally, the *e³value* contains the notion of *dependency paths*. Such a path consists of *consumer needs*, value interfaces, value transfers *connection elements* (visualized as straight lines in the *interior* of an actor), and *boundary elements*[1]. A dependency path shows which transfers must happen, as a result of a consumer need. It does not show *when* they will happen, only *that* they will happen in the contract period described by the model. This is sufficient to estimate economic profitability in the contract period.

The technical and business processes by which these transactions are implemented contain a lot more detail and are not shown [76]. It is even possible that the coordination processes that implement the commercial transactions implement a value transfer between actors A and B by

[1] A dependency path also may contain *AND*, *OR*, and *explosion/implosion* elements to represent dependency splits and joins, but for explanatory purposes, these are not used in Fig. 6.2.

means of a coordination process involving actors A, B, and C. An *e³value* model abstracts from these operational details and shows commercial transactions only.

In Fig. 6.2, user A needs to make a call to User B. In exchange for the call, User A pays a sum of money. By following the dependency path, one can see that provider A should obtain an interconnection in order to provide the telephone call, and should pay for this interconnection. Finally, provider B delivers a telephone call service to user B.

For now, it is important to understand that in this *e³value* model *all* transfers on a dependency path are assumed to occur. In other words, the model shows what happens in reality, only all actors behave as agreed and expected. So, actors are always paying, and services are always provisioned. That is why I call such a model an *ideal* model; all actors operate honestly.

6.5.3 Construction of Sub-Ideal Business Value Models

In real-life, actors do not always behave as agreed and expected. They can perform intentionally or unintentionally in a different way. For example, an actor may not pay or pay a wrong amount of money. *e³fraud* covers three types of misbehaviour:

1. Transfers which should happen in the ideal model, do not happen at all in reality;

2. Transfers that happen in reality are not supposed to happen in the ideal model;

3. Actors assumed to be independent in the ideal model may collude.

We construct sub-ideal value models from the point of view of the same actor, in this example provider A, who is the ToA. A sub-ideal model represents the business value model as intended by the fraudsters and is created by changing the ideal model to represent misbehaviour.

Figure 6.3 shows an example of revenue sharing fraud, exhibiting both types of misbehaviours described above, as well as collusion. Rather than two independent end-users, as shown in Fig. 6.2, we have now a collusion between Users A and B. They could be the same person, or several individuals working together.

In addition, the contract between User B and Provider B entitles User B to a *Revenue Share* for every *received*. This is a common arrangement for 0900 numbers, but some regular providers provide a fraction of the interconnection fee as an incentive. Again, since we take the point of view of Provider A, we have no information on how User B obtained a contract with provider B or what the structure of their agreement is. For the fraud analysis, it is sufficient to assume such a bonus is being paid. Since this bonus is unknown to Provider A (the ToA), it is represented using a *dashed* line. Note that user A only uses provider B to *receive* calls.

To make matters worse, in this sub-ideal scenario we assume User A does not intend to pay his monthly fee to Provider A. As it is a non-occurring transfer with respect to the ideal model of the ToA, the *Subscription Fee* value transfer is represented using a *dotted* line.

User A will now place as many calls as possible per month with provider B. As can be seen by following the dependency path, the *same* user A also terminates the call, but with his phone hosted by provider B. For each terminated call, user A receives a bonus. Considering that, in addition, he also intends to default on his payment of the Subscription fee, he is in the position to make a generous profit.

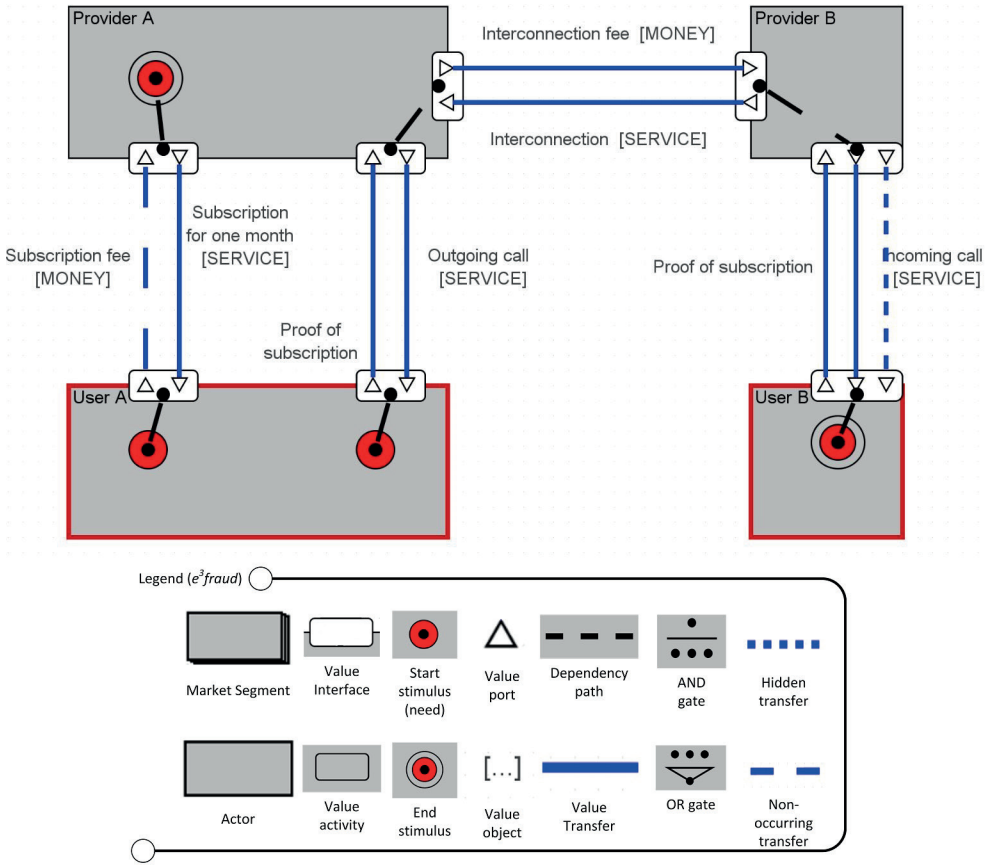

Figure 6.3: Sub-ideal model: User A calls himself and earns money

6.5.4 Financial analysis of the attack

The most important financial factors of Telecom fraud, with regard to Risk Assessment are (1) losses incurred by the provider and (2) motivation (in terms of potential gain) of the attacker. The former allows for estimating the impact of each type of fraud, while the latter is a critical part in estimating the likelihood of such a scenario taking place. Computing loss and again projections therefore provides the analyst with an indication of the likelihood and impact of a fraud scenario. Together, likelihood and impact can be used to estimate an overall risk level associated with each particular scenario [32].

To estimate losses or gains, we need to analyse a pair of models: an ideal model showing the *e*^{3}*value* model of normal usage and a sub-ideal one showing the *e*^{3}*fraud* model of fraudulent usage. Furthermore, in pay-per-usage environments, such as telecom, the magnitude of the risk is dependent on the scale of usage (e.g. minutes called).

A custom software tool was developed that takes as input two models and generates profitability graphs for each sub-ideal case(s) showing the dependency of the profitability with regard to usage. This allows for a visual comparison of ideal vs. non-ideal business cases of the provider as well as regular vs. fraudulent business case for the customers (and potential fraudsters) across a given range of occurrence rates.

The two graphs in Fig. 6.4 are generated by this tool from the models shown in Fig. 6.2 and Fig. 6.3. Realistic and, where available, real, values were used to instantiate the models. The chosen values are based on tariffs charged by Dutch Telecom providers in 2014. For simplicity, I only show the financial result of user A and Provider A (vertical axes), relative to the number of calls made (horizontal axis).

The financial outcome expected by the Telecom provider, in normal usage conditions, is visible in Fig. 6.4a. Here, the user has a fixed cost, the monthly cost of the subscription. The costs of Provider A increase with each call, due to the termination bonus paid to Provider B for the interconnection. Operating costs are not represented here as they are unknown and assumed to be negligible for an individual user, but could be easily included in the model. The sub-ideal case (Fig. 6.4b) is significantly different. Besides the clear loss for the provider, the fraudster's financial motivation is now clearly visible.

While the value model(s) alone do not contain sufficient information to reliably elicit procedural or technical countermeasures on its own, these financial results can be used to discuss checks on the non-occurring transactions or infer thresholds based on break-even points so as to mitigate the Risks.

6.6 Using the *e*^{3}*fraud* approach to quantify technical risks

Real-world security risk assessments typically result in the identification of a list of risks that is too long to be mitigated completely. The assessors must, therefore, prioritize the risks and mitigate only the most "important" ones. Importance is usually estimated by factoring in the attractiveness of a potential attack to the attacker with the amount of loss caused by the attack. In this section, I demonstrate how the *e*^{3}*fraud* approach could be used to complement a Risk Assessment of the infrastructure by facilitating impact as well as the gain estimation of individual risks, based on ranges of variables.

The alternative approach described in this Section takes as input a technical risk described in terms of an ideal model (before the attack) and a sub-ideal model (after the attack) and

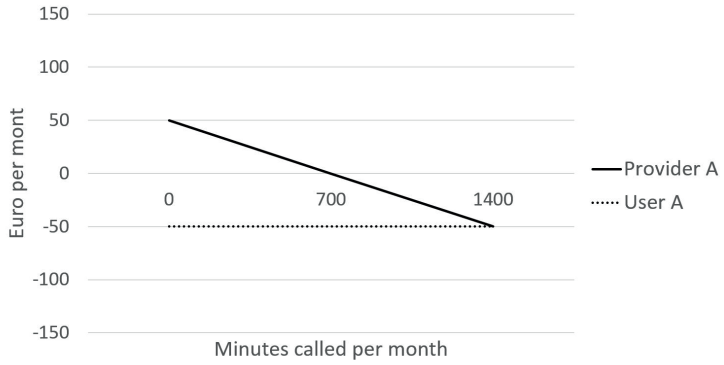

(a) Ideal case

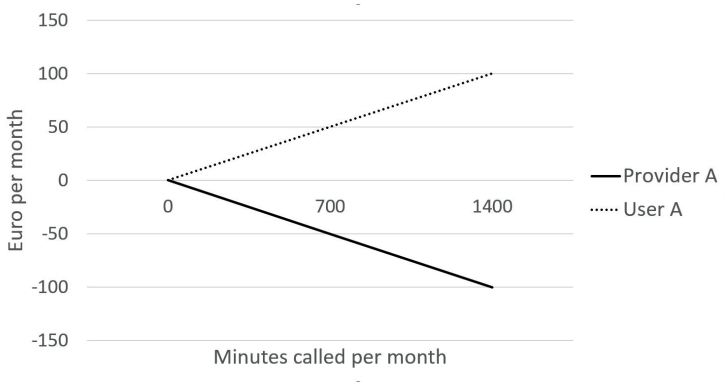

(b) Sub-ideal case

Figure 6.4: Profitability graphs of the RSF scenario

produces graphs showing the financial loss/gain related to the Risk.

6.6.1 Scenario Description

A Private Branch Exchange (PBX) is a telephone exchange or switching system that serves a private organization and performs concentration of central office lines or trunks and provides intercommunication between a large number of telephone stations in the organization [176]. By exploiting vulnerabilities in a company's PBX, fraudsters may obtain access to one or more of an organization's phone numbers, which they can then use for personal, often fraudulent, purposes. Although attacks on the phone infrastructure are not as notorious as the revenue share fraud analysed above, reports show they are as likely to occur as an attack on the data network [177].

There are several ways to attack a PBX. As Kuhn [178] describes, the most vulnerable is the remote access feature. Through this feature, for example, a fraudster can create a special mailbox which redirects him to a phone number of his choice. This number could be either a premium-rate (0900) number owned by a criminal organisation the fraudster is part of or a

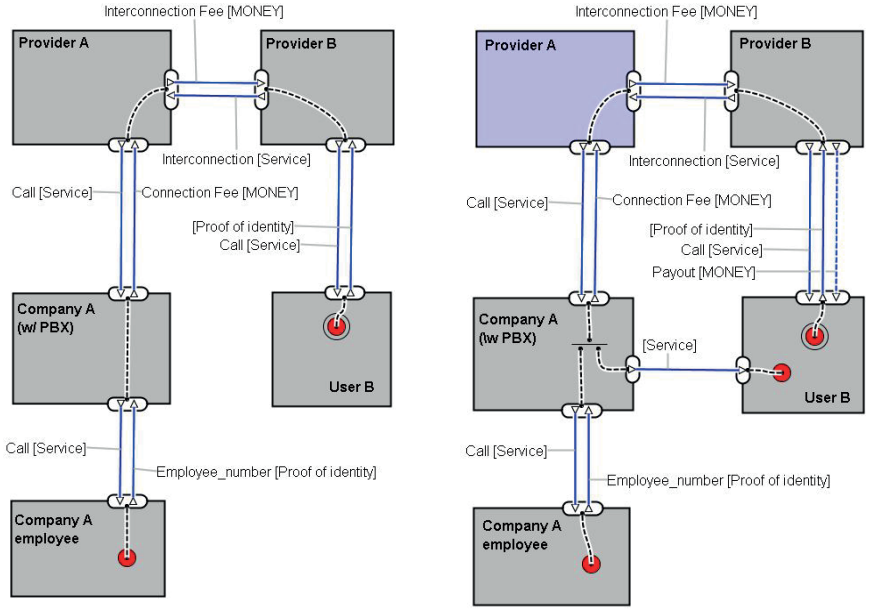

(a) Ideal model: Employee calls User B

(b) Sub-Ideal model: User B uses illicit access to PBX to call himself

Figure 6.5: Models used to analyse the Risk of PBX hacking

number that provides the callee with a revenue share for every received call.

Another option would be to obtain possession of a telephone within the company and start calling his number from there [179]. One way to do this is to blackmail or bribe a company employee. In e^3*fraud*, we abstract from the technical or social means to access a company's phone number, and concentrate on the business model for the attacker.

6.6.2 Construction of Ideal and Sub-ideal Business Value Models

We want to estimate the potential loss the company would face in case of unauthorized access to its PBX, as well as the potential gain a malicious actor with such access could obtain.

To start, we create an ideal model of the value exchanges as perceived and expected by the Target of Assessment, whose perspective we take. In this case, we take the perspective of the Company who owns the PBX. Figure 6.5a shows the ideal model: Employees (Company A employee) may call through the company PBX to external Users (User B). This is the normal usage the company expects, given honest actors.

Next, we tweak the model to show the commercial traces of the risk we want to analyse. For example, Fig. 6.5b shows a sub-ideal model where User B is an attacker. He obtains access to the PBX (dashed line) and exploits this (illicit) access order to make calls through a provider that pays out a reward for every call received.

It is interesting to note that the value model abstracts away from the actual vulnerability being exploited and the process by which access is gained. In order to fully describe the attack,

a coordination model such as an attack tree is needed.

To do a financial analysis of the risk, we instantiate both models with known or realistic values and use the *e³fraud* computation engine to generate profitability graphs showing the financial result of all actors involved for both the ideal and sub-ideal case, similar to the ones shown in Fig. 6.4.

By contrasting the financial result of *Company A* in the ideal model with its financial results in the sub-ideal model, we can quantitatively estimate the *impact* associated with a technical risk. Furthermore, by computing the financial result of the attacker in each sub-ideal case, we can estimate the expected *gain* of each attack. It is worth noting that the costs of setting up the attack are not taken into account and thus we cannot claim to estimate the attack's *profitability* (as *costs-gain*).

Using the analysis described in Sect. 6.5.4, we can produce a multitude graphs, for various ranges of malicious calls.

Tweaking model parameters and re-running the tool also provides an efficient way of conducting sensitivity analyses.

Because *e³fraud* models abstract away from the technical means of conducting an attack, there is reason to believe *e³fraud* models might be re-usable, meaning the overhead of creating such a model for each Risk would not be significant.

6.7 Focus group

The approach was also demonstrated to domain experts working for the fraud department of a leading Telecom provider and received positive feedback. The method was generally perceived as useful, especially for quickly assessing the financial fitness of new plans before they are launched and estimating the impact of new types of fraud on existing plans. Furthermore, the experts saw the profitability graphs as an expressive means of communicating risks to product managers.

However, experts envisage several functional improvements before the *e³fraud* method and toolkit would be usable in practice: the ability to model a larger variety of sub-ideal models (such as ones containing hidden actors) and the possibility of automatically generating and ranking sub-ideal models based on a given ideal model. Finally, a library of model patterns was mentioned as a way to promote (re)usage.

6.7.1 Limitations

While *e³fraud* can be used to help reduce possibilities of fraud on the service level, as well as to quantify some known infrastructure risks, it does not necessarily help to identify an attack on the technical infrastructure, such as a DDOS attack. *e³fraud* is especially applicable in cases where the economics of risk are of particular importance, such as for analysing fraud that takes place at the service level rather than the technology level.

The biggest strength of *e³fraud* models - similar to the *e³value* models they build upon - is also their biggest weakness: they abstract away from any and all procedural and architectural information. But in some cases the order in which the transactions happen is important. For example, its impossible to make a call with a SIM card that hasn't been bought yet. The

how question (critical to process models) does not concern us, but the order in which certain transactions are executed does matter.

Even though the order of execution is not important for the business model, to discuss countermeasures, we need to be able to reason about the coordination processes and IT architectures that can mitigate risks. In general, transformation to or generation of any sort of architecture or coordination model from an *e*^{3}*value* model is not feasible [76]. There exists previous work discussing these relationships for [180–185]. However, none of these papers are about (in-)security or fraud and mostly assume ideal business environments. This motivates a closer study of the relation between *e*^{3}*fraud* models, coordination models, and enterprise architecture which are relevant or useful in the context of Risk Assessment.

Finally, this work is still in its initial stages. As the methodology and tool have been developed and tested on a limited number of telecom fraud scenarios, the approach is somewhat example-driven. Thus, an obvious next step is to model and analyse a larger variety of scenarios so as to further develop and validate the idea.

6.7.2 Generalisability

I have developed and illustrated the *e*^{3}*fraud* approach on a number of cases from telecom service provision. To which extent is this generalizable to other kinds of e-service provision? At the moment I can only speculate about this, but the true test of generalizability is the application of *e*^{3}*fraud* to other kinds of e-service provision networks. I plan to do this in future work. In the absence of any such empirical evidence, I analyse the features of the studied telecom service networks that make *e*^{3}*value* and *e*^{3}*fraud* suitable to identify and analyse business risks.

In the telecommunications sector, information on the technical infrastructure of competitors is unobtainable, which makes *e*^{3}*value* and *e*^{3}*fraud* well-suited to model business risks. Describing the money flows and their triggers is necessary and sufficient to understand such scenarios and not only derive estimates of both impact for the provider and gain for the fraudsters, but also identify countermeasures.

A second characteristic of the telecom sector is the importance of a short time-to-market. The marketing department of a Telecom provider typically wants to launch new services without delay and so any kind of initial analysis of prospective risks arising from proposed products will need to be comprehensive enough to be meaningful and yet quick enough to be acceptable. Once the product is launched, it will be important to identify any unacceptable activity at the earliest opportunity, to minimise the losses associated with this. Initial evidence shows that *e*^{3}*fraud* offers the promise to offer efficient support in risk identification and mitigation. To further improve this efficiency, I am currently working on automatically generating and ranking the sub-ideal models.

Based on this brief analysis of the factors that contribute to the usability and usefulness of *e*^{3}*fraud* in the identification and analysis of risks of fraud in telecom service provision, I speculate that *e*^{3}*fraud* will be equally useful in other cases of e-service provision where information about the technical infrastructure of competitors is unobtainable, time-to-market of new services must be short, and losses created by instances of fraud must be kept within acceptable bounds. I plan to do case studies that provide evidence for this speculation in future research.

6.8 Conclusions and future work

This chapter introduced a conceptual extension to the *e*^{3}*value* ontology that allows using business value models to identify and quantify risks of fraud in e-service networks, where information about the technical infrastructure of the partners in the network is incomplete or even absent. I've shown how *e*^{3}*fraud*, a proof-of-concept extension of *e*^{3}*value*, is able to identify, model and analyse business risks, as well as quantify the business impact of technical or procedural risks. It can be applied to the analysis of existing networks as well as to the risk analysis of newly proposed ways of doing business. It can also be used to estimate the business impact of technical attacks and helps to decide which technical attacks are worthwhile defending against.

The approach presented in this chapter has been successfully applied to four Telecom fraud scenarios, containing a variety of business as well as technical risks. The models were validated with the help of the scenario owners. Results matched existing estimations and by using real values I was able to show that at least one instance Revenue Sharing Fraud is still possible today.

The above brings about the question of deciding when and where each type of model is needed in order to conduct an effective Risk Assessment. Furthermore, a main topic of research for the coming year is investigating how value models can be integrated into existing Risk Assessment methodologies and frameworks. One option would be generating the root node of an attack tree from an *e*^{3}*value* model: each undesire-able transaction could be decomposed into atomic actions needed in order to achieve it, thus forming a (sub-)attack tree. Another, simpler option would be simply using *e*^{3}*value* models for impact and gain estimation, leaving the other analyses to be performed on more technical models.

7

Automated business risk identification using value models

Based on *Automated identification and prioritization of business risks in e-service networks* [15].

Modern e-service providers rely on service innovation to stay relevant. Once a new service package is designed, implementation-specific aspects such as value (co-)creation and cost/benefit analysis are investigated. However, due to time-to-market or competitive advantage constraints, innovative services are rarely assessed for potential risks of fraud before they are brought to market. But these risks may result in loss of economic value for actors involved in the e-service's provision. In this chapter, I introduce an approach that automatically generates and prioritizes undesirable scenarios from a business value model of the e-service, thereby drastically reducing the time needed to conduct an assessment. I provide tool support to generate fraud scenarios from a business model, in which actors may not perform transactions that they committed to, perform secret transactions, or collude with other actors. The tool can rank fraud scenarios on various criteria, such as loss to a service provider or profit to a fraudster. I also provide examples from telecom service provision to motivate and illustrate the utility of the tool.

7.1 Introduction

Many services are *commercial* services. That is, they are of economic value to someone, and are paid for. As a result, end users and enterprises may be tempted to commit fraud or abuse, which I refer to as non-ideal behaviour. Such non-ideal behavior of actors involved in the acquisition or consumption of the service can lead to undesirable losses for the provider or undeserved gains for other actors. Examples include, but are not limited to, misusing the service, bypassing payments and exploiting unintended interactions between services. For example, in the field of telecom service provision, "simboxing" involves acquiring telephone services from multiple providers and setting up a composite service that disguises international calls as local traffic, thereby undercutting termination fees [186].

The problem is exacerbated because many services are in fact electronic services, which are provisioned via the Internet or other digital means [166]. These electronic services are characterized by short time-to-market (typically a few months). But these e-services are provisioned over complex networks, that increases the opportunity for malicious actors to commit fraud or otherwise misuse e-services [187]. Risk assessment thus becomes more complex, and this creates a tension with the desire of marketeers to put out innovative e-services fast. Thus, there is a need to speed up and enhance the capability of e-service risk assessment.

Service innovation commonly consists of three phases: Service exploration, where potential new or improved services are identified; Service Engineering, where one or more of the options are explored in detail; and Service Management, which deals with implementation and continuity [188]. The Service Engineering phase carries particular importance, as errors introduced in the early phases of service design can have significant (financial) consequences later on [189, 190]. E-service risk assessment should, therefore, be done in the service engineering stage. This requires quantifying the cost of misuse and designing prevention or detection mechanisms, which in turn requires projections, usage estimates and financial computations [191]. Doing this in a way that does not unduly slow down service innovation requires efficient tool support.

Business risks for a provider include fraudulent violations of contracts by clients, violations of agreements or terms of service, as well as the creation, by clients, of false expectations with the provider with regard to usage. I define **fraud** as the intentional misrepresentation by a client of his or her intentions, in order to acquire something of value from a provider. Fraud may be legal or illegal. I call the actor performing a fraud a **fraudster**. In this assessment of fraud risk, I sidestep the issue of legality but focus on the potential loss for the provider and potential gain for the fraudster. In other words, I focus on what is observable for the provider (his loss) and on what can be estimated about the fraudster, given a business model (his potential gain). The potential loss is the negative impact that the risk can have on the provider, and the potential gain for the fraudster indicates the likelihood that the fraud will be committed.

In Chapter 6, I introduced *e^{3}fraud* as a model-based approach to assessing business risks in e-service networks [14]. The approach recognized three basic fraud operations, namely not paying for a delivered service, performing a hidden value transfer, and colluding with another actor; and I introduced different ways to estimate loss for a service provider and profit for a fraudster. Non-payment breaks transactional reciprocity (and causes loss) [192], hidden transfers can encourage misuse (by providing hidden gains) [193] and collusion allows

exploiting unintended interactions between atomic services [194]. My goal in this chapter is to scale up these techniques to non-trivial scenarios by *automatically* generating fraud scenarios from a business value model. We will see that for realistic business models, the sheer number of possible variations is staggering, so the ability to filter, rank and group risks so as to zoom in on the riskiest scenarios, becomes critical. This chapter describes a scalable tool support for generating, quantifying and ranking business risks directly from a business model of the given service.

This approach to fraud risk assessment is constructive in the sense that we analyze the architecture of a business model, in particular, a business model represented in *e³value*, to construct possible mechanisms to commit fraud. This distinguishes it from statistical approaches to assess fraud risk [195], which use patterns of past client behavior to assess fraud risks in new business models. Since the new business model has by definition not contributed to the statistical data on which this assessment is based, statistics-based fraud assessment leaves one with unknown and un-estimated risks. Therefore, the approach proposed in this chapter is able to discover fraud scenarios a priori, while statistical models identify fraud a priori.

e³fraud is based on the *e³value* method for representing business models and allows the generation of possible fraud mechanisms in new business models. The *e³fraud* tool (available at `https://github.com/danionita/e3fraud`) can automatically generate misuse scenarios based on configurable heuristics, such as collusion, non-payment and hidden payments. Furthermore, it can group and rank such scenarios on various criteria, such as loss to a service provider or profit to a fraudster. Finally, it can help visualize the financial results across a range of projected usage levels. I illustrate the tool using examples from telecom service provision.

The chapter is structured as follows: Sect. 7.2 introduces the new *e³tool*. Sect. 7.3 describes the application of the approach to a telecom service and showcases the results provided by the tool. Finally, Sect. 7.4 draws some conclusions with regard to the approach, its applicability, and future development.

7.2 The approach and its implementation

7.2.1 Starting point: the *e³fraud* methodology

In the previous chapter I've shown how an *e³value* model can be extended to describe fraud. As a reminder, the resulting *e³fraud* models differ from the original *e³value* model in several ways:

- An *e³fraud* model takes the point of view of one actor in the network, dubbed the **Target of Assessment** or ToA and marked with a thick border. This is needed to define the concept of hidden transactions, introduced below, and to assist with ranking the possible fraud models according to the potential loss for the ToA (further described in Sect. 7.2.2.2).

- Value transfers may not take place and are marked using dashed lines. In Fig. 6.3, the "Subscription Fee" transfer does not take place.

- Hidden transfers may occur between secondary actors, not involving the ToA and are marked using dotted lines. The ToA cannot directly observe these hidden transactions.

- Actors might collude, which means that they pool their budgets. Collusion is usually kept hidden for the ToA.

This means that for any *e³value* model, there can be derived a large number of *e³fraud* models. This search space grows polynomially with the number of actors, value transactions, and value interfaces. The *e³fraud* approach involves building several of these sub-ideal models and comparing their financial outcomes to that of the ideal model in order to estimate the potential impact of each instance of fraud or misuse [14]. Therefore, manually creating these models and re-running the analysis is time-consuming. Furthermore, deciding which scenarios should be mitigated implies comparing a large number of models and this cannot be done manually. In the following section, I describe an approach to delegate the time- and resource-intensive tasks of generation and ranking to a computer. The approach is implemented as a Java tool, *e³tool*.

7.2.2 First implementation: the *e³fraud* tool

As a proof-of-concept, I first developed a prototype tool which incorporates the *e³fraud* extension and is able to generate and sort possible sub-ideal variations of a given value model, based on the three fraud heuristics of *e³fraud*. The generation tool is open-source and publicly available at `https://github.com/danionita/e3fraud`. It performs three tasks:

1. Generating *e³fraud* models, representing various fraud scenarios;

2. Ranking the generated *e³fraud* models, so that the business model designer can zoom in on the riskiest ones;

3. Computing and plotting the profit/loss of each actor across a given usage projection.

7.2.2.1 Generation

Given a valid *e³value* model, this prototype tool generates all combinations of possible deviations: *hidden transactions, non-occurring transactions, and collusions*. As explained in Chapter 6, these patterns can be used as the building blocks for several telecom fraud scenarios [30]. Each valid combination is then instantiated as new sub-ideal model. A sub-ideal model may contain any number of hidden transactions and non-occurring transactions but only one collusion. The number of actors colluding is configurable.

Hidden transactions
are generated in three steps.

- First, identify pairs of transactioning secondary (non-ToA) actors.

- Then, for each such pair, the profit/loss resulting from the dependency path of which this transaction is part, is computed for each actor.

- Finally, for the actor(s) with a positive result, a new outgoing transaction is added: the value of this transaction is a fraction of the positive result. The reasoning behind this is that if an actor A makes some profit off of a dependency path, A might be willing to pass on part that profit to another actor, B, if that would motivate B to generate occurrences of that dependency path. This models an established practice in the services industry called Revenue Sharing [196]. By default, this share can be 1/3 or 2/3 of the profit but this is configurable.

The generation of hidden transactions is thus bound by the number of actors and transactions in the ideal model.

Non-occurring transactions
are created by invalidating individual monetary transfers (that is, transfers marked as type MONEY). The restriction to monetary transfers is to limit state space explosion, but this assumption could be dropped in the future. The user may indicate that certain MONEY transfers will always occur by marking them as type MONEY-SECURE. This can be either because they are initiated by the provider itself, because safeguards are in place or simply to reduce the search space of sub-ideal models. This will prevent these transfers from being invalidated by the generation engine. The generation of hidden transactions is thus bound by the number of monetary transfers in the ideal model.

Collusion
takes place when two actors are acting as one: they pool their budgets and collectively bear all expenses and profit. By colluding, actors might deceive controls and invalidate expectations by appearing independent, but, in fact, working together against the best interests of the provider. Therefore, only secondary actors (not the ToA) can collude. To generate collusions, pairs of secondary actors are merged into a single actor. The number of actors allowed to part of a colluding group is configurable. The generation of collusions is thus bound by the number of actors in the ideal model.

7.2.2.2 Ranking

Depending on the complexity of the initial ideal model, hundreds or even thousands of models might be generated. Many of these might not be possible due to existing controls or might be unlikely because they are not profitable for any of the actors.

To aid with selection and prioritization of risks, the tool provides several ways of ranking and grouping the set of generated models, described below. The prioritization is always carried out from the perspective of a single actor (the Target of Assessment), as described below.

In terms of value creation, non-ideal behavior causes a disruption in the financial result of the actors involved. This means that a non-ideal scenario can (1) cause a loss for the service provider and (2) trigger an unexpected gain for one of its customers or users. As such, the software tool allows ranking based on Loss for the ToA, Gain for the secondary actors and Loss+Gain. The gain of a secondary actors is defined as the difference between the financial result of a that actor in the ideal case versus the sub-ideal case.

Ranking on Loss+Gain ranks risks on negative impact for the ToA (Loss) and profitability for the potential fraudster (Gain). This is similar to the classical definition of risk as Impact times Likelihood of an event, except that I do not use Likelihood of the fraud but Gain to the

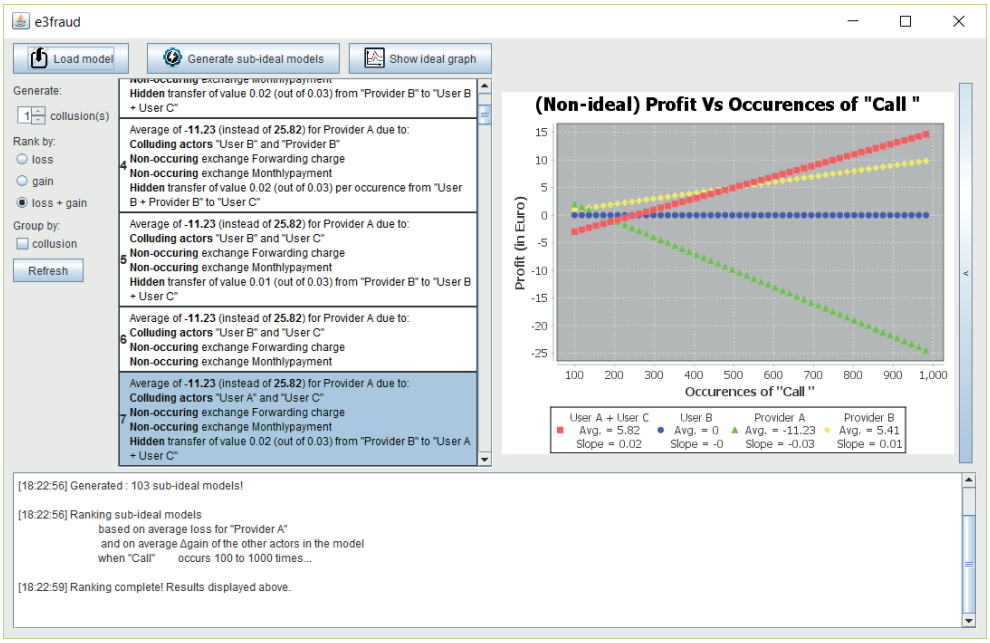

Figure 7.1: Screen-shot of the e^{3}fraud *prototype tool*

fraudster. I use Gain to estimate the attractiveness of a fraud to a potential fraudster. A fraud with a higher gain for the fraudster is more attractive to the fraudster, and therefore more likely than a fraud with a lower gain.

Furthermore, to allow for"what-if" analyses and easier navigation through the long list of sub-ideal models, results can be grouped based on who is colluding with who. Since each group is ranked independently, this allows investigating the riskiest way each pair or group of actors can collude.

7.2.2.3 Visualization

The ranked list of generated non-ideal scenarios is presented by the *e^{3}fraud* tool as a list of textual descriptions. If grouping was selected, the list is nested and collapsible. This facilitates the exploration of the state space.

When one of the non-ideal scenarios is selected, a chart showing the financial implications of the scenarios is shown. This chart uses the occurrence interval given during generation as the X-axis and shows the result of each actor across this interval on the Y-axis. A screen-shot is shown in Fig. 7.1. These representations can be understood by marketeers and product managers without having to learn *e^{3}value* or *e^{3}fraud*.

The results contain several useful pieces of information. Firstly, they show the loss for the ToA across the given occurrence range, as both a plot and an average. Loss is a direct indicator of the potential impact that each of the particular fraud risks can bring about. Secondly, the gain experienced by all other actors in the model is also shown as both a plot and an average. This gain can be used as a proxy for likelihood: the higher potential gain for some actor,

the more likely it is that he will attempt that specific fraud scenario. Finally, the slope of all the plots are estimated, which gives an indication of how the loss and gain of the fraud scales with usage, outside the given range, and therefore how the impact and likelihood vary. Visualizing the result as a plot also allows for easy, visual identification of break-even points and thresholds.

7.2.3 Second implementation: The *e³tool*

In order to use the prototype *e³fraud* tool described above, the user had to first construct a value model using the original *e³value* editor, then export and load them into the fraud generation tool. Furthermore, it did not support viewing or editing the generated models. Therefore, the approach was recently built into an integrated value modelling and analysis tool called *e³tool*, available at `https://github.com/danionita/e3tools` in order to streamline the fraud assessment process. The new tool shows the same ranked list of textual descriptions of possible fraud scenarios, but for each description, *e³tool* also shows a preview of the model and a table comparing the financials of each non-ideal scenario to the ideal case. See Fig. 7.4 in the next section for a screenshot.

The new implementation also adds the ability to select multiple trusted actors and choose which of the three fraud heuristics to be applied, and even customize them. In addition, *e³tool* provides full support for multifactor sorting, where models may be sorted first based on the highest loss, then on the highest gain, or the other way around. To mitigate search space explosion, models with similar financial results are grouped and several filtering options are now available. Furthermore, fraud models which are neither profitable nor cause any loss are hidden by default.

7.3 Preliminary evaluation results

The approach has been applied to several telecom service packages known to be exploitable. For the case study described in Sect. 6.5.1, *e³tool* returns the known revenue sharing fraud model manually constructed in Sect. 6.5.3 as the highest ranked non-ideal scenario. In this section, I use a slightly more complex case to further validate the approach: call forwarding to other networks via post-paid subscription.

A value model depicting the call forwarding service of Provider A is shown in Fig. 7.2. Provider A is our target-of-assessment, i.e. the entity whose risks we are assessing. User A is a customer of Provider A. In this case, he is using a pre-paid SIM card to make calls to some other customer of Provider A, namely User B. Since the post-paid SIM is received upon payment, we model this purchase as [MONEY-SECURE], since it is impossible for the payment to be bypassed. User B, in turn, has a post-paid subscription with Provider A. This subscription involved a fixed monthly payment, plus an incremental payment based on his usage. In this simple example, he does not initiate any calls and has set his device to forward all received calls to User C. User C is a customer of Provider B, and therefore User B will have to pay for the connection from his own Provider, A, to User C in the form of a forwarding fee. Since Provider B is a separate commercial entity, we do not know his contractual structure nor his usage pattern, so we only model the fact that User C can receive calls.

A known fraud scenario (shown in Fig. 7.3) for this case is as follows: User A initiates a

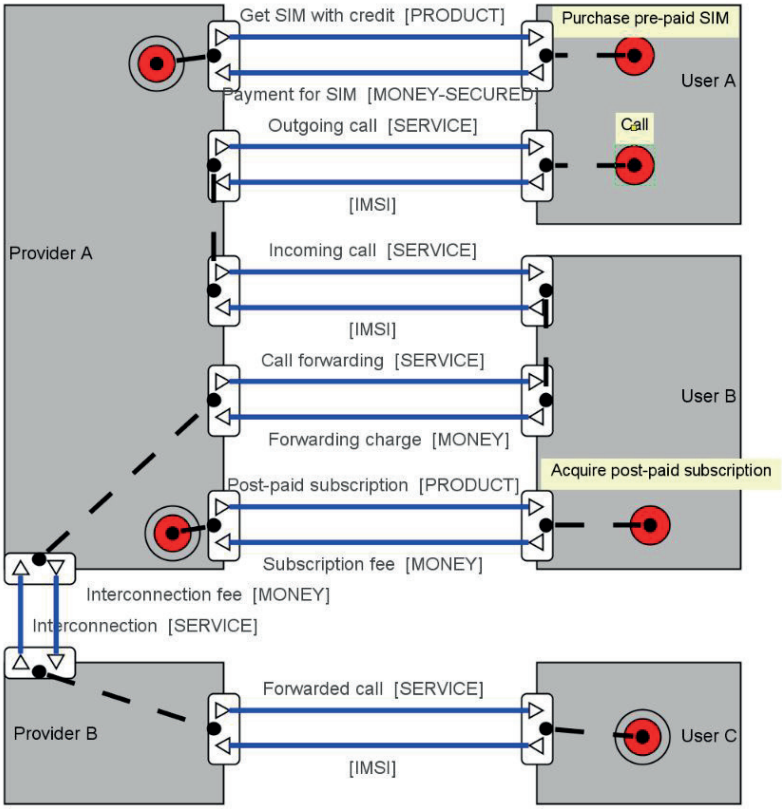

Figure 7.2: Value model of call forwarding to other provider

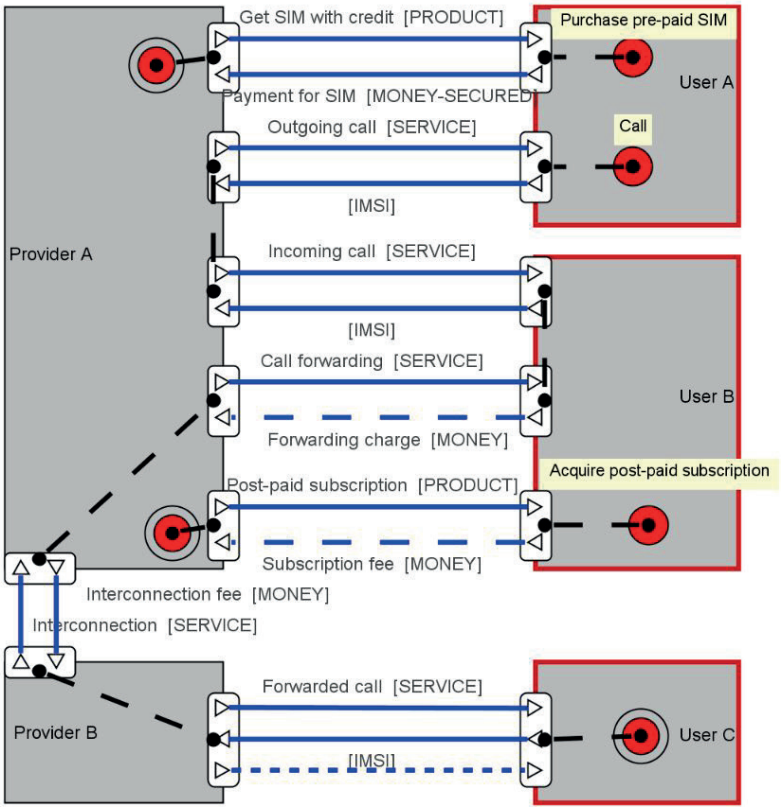

Figure 7.3: e^{3}fraud *model of fraudulent call forwarding to other provider*

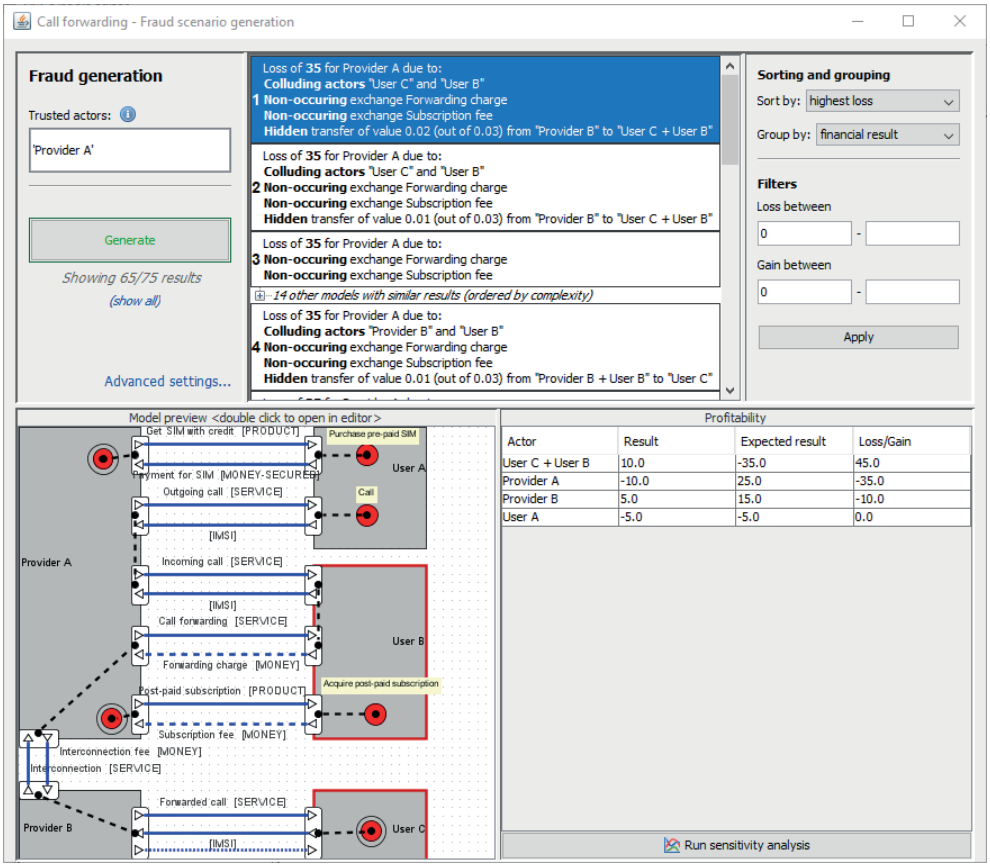

Figure 7.4: Screen-shot of the e^{3}fraud *tool's output for the value model in Fig. 7.2.*

very large number of calls to User B. The calls are usually charged at a preferential rate, as the two users are on the same network. User B, as the forwarding party will have to pay for the call outside the network. He will receive a very large bill at the end of the month, which he does not pay. In reality, User A commonly pays User B a small amount of money to start a post-paid subscription using fake credentials, so that User B cannot be traced by Provider A. Finally, User C, the end-recipient of the call would have a Revenue Sharing agreement with Provider B, by which he receives a payout for every call that comes through. This is common with 0900 numbers but also available with some "budget" subscriptions. For this model, we used the following realistic values: 5 euro for the pre-paid SIM including 500 minutes of free calls within Provider A's network, 10 Euro for the post-paid SIM, 5 Eurocents per minute for the forwarding charge and 3 Eurocents per minute for the interconnection fee.

By running the model in Figurescenario2 through the *e^{3}fraud* tool using default settings, we obtain the fraud scenario described above and visualized in Fig. 7.3 as the highest ranked non-ideal scenario. Fig. 7.4 shows a screen-shot of the actual output of *e^{3}tool*. The left part of the screen describes the highest ranked risk as:

Loss of **35** for Provider A due to:
Colluding actors "User B" and "User A" and "User C"
Non-occurring exchange Forwarding charge
Non-occurring exchange Monthly Payment
Hidden transfer of value 0.02 (out of 0.03) from "Provider B" to "User B + User C + User A"

The bottom-right of the screen provides more details: 25 Euro is the "ideal" profit for the ToA in the ideal case (i.e. the model provided as input to the fraud generation module). The assumption is that this value is what the ToA would have expected to obtain given its own estimates. However, in this scenario, the ToA will actually incur 10 Euros worth of costs, adding up to a relative loss of 35 Euro. The reasons for this is provided in the textual description above, but also shown in the visual preview shown in the bottom-left of the screen: three actors (Actor A and Actors C) are colluding, User B will not pay his bill this month (consisting of a subscription fee and several forwarding charges, and Provider B is passing two thirds of his revenue per call to the three colluding users.

The colluding users are able to obtain a profit of 5 euros, even after deducting the costs of the prepaid SIM. This is not much, so we might want to see whether this fraud can be scaled up. One way to do so is by making more calls. By clicking the "Run sensitivity analysis" button below the profitability table, the user is able to generate projections of how the non-ideal scenario might scale. If we choose to use the occurrences of "Call" as a parameter, we obtain a chart such as the one in Fig. 7.5 The x-axis represents the number of calls by User A and the y-axis represents the corresponding profit and loss for the actors. The steepest upward line is "User B + User A + User C" (the fraudsters). The next steepest upward line is "Provider B" and the downward line is Provider A (the ToA). The graph, therefore, shows that this particular scenario allows the two colluding actors to scale up the fraud. It also shows that, for the provider, this fraud has the potential to cause loses which get significantly worse with the number of minutes called. But this ignores the fact that calling more requires purchasing more SIM cards and the fraud generation module does not allow performing sensitivity analyses on two parameters. To overcome this, we can double-click the preview in the bottom-right of the fraud generation modules' window to open it in the editor. Then, we can tweak the occurrence rates of the number of SIMs or the number of calls. We can also explore the effect of countermeasures such as fair use policies, by limiting the number of calls.

Other damaging variations of this non-ideal scenario may be identified by scrolling further down the list. But some of the scenarios generated may be unfeasible or even impossible, but this has to be decided per scenario. For instance, it makes no (financial) sense for Actor A to initiate the calls in the first place if he is in on the fraud. Furthermore, other highly ranked risks imply a User colluding with Provider B, thus also obtaining Provider B's legitimate interconnection income. While collusion between user and providers is not impossible, in this case, it is extremely unlikely. Grouping the results per collusion would help in this case to eliminate unrealistic collusions from the analysis.

These results can help design fraud detection thresholds, identify the transaction that require (further) procedural or technical controls or even trigger a re-design of the service in order to eliminate or mitigate business value risks [191].

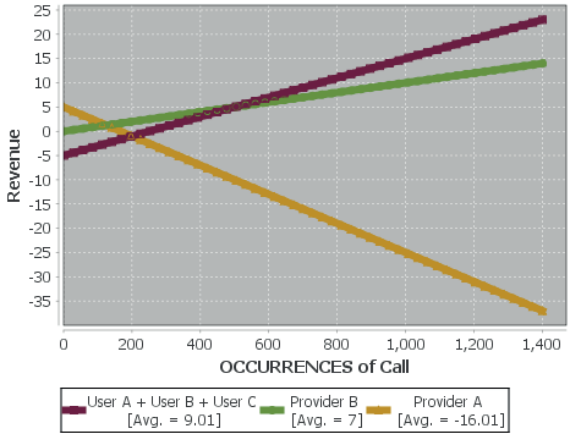

Figure 7.5: Sensitivity analysis of the non-ideal model of Fig. 7.3

7.4 Conclusions and future work

This *e³fraud* approach described in this chapter provides a novel, constructive, semi-automated method for conducting quantitative risk assessments of value models. The approach relies on a small set of misuse patterns commonly seen in telecom fraud, that so far has been sufficient to generate known fraud scenarios. The results highlight the potential of automating the identification, quantification, and ranking of business risks associated with one or more service offerings.

As noted earlier, telecom providers already have statistical means of estimating fraud [195]. But these means assume the future is like the past. Predictive models based on large data sets of past service deliveries do not necessarily indicate what the risks of new, innovative e-service provision arrangements are. *e³fraud* provides a supplementary approach that analyzes the architecture of a service provisioning network and identifies risks that follow from the structure of the network, by actually constructing the fraud mechanisms. This allows marketeers and managers to identify the source of these risks and take preventive measures before the risks materialize.

So far the approach has only been applied to cases with known fraud scenarios. Understandably, telecom providers are reluctant to disclose all of the fraud scenarios known to them, and so it will be very hard for us to know whether fraud scenarios generated by the tool were already known to them or not.

I have shown feasibility in other cases [14]. However, more real-world cases are of course needed to confirm generalizability. To test whether the tool is able to find fraud scenarios not known to *us,* I plan to collect new business models and try to identify unexpected scenarios. To further test the generalizability of the ideas of this chapter, I intend to analyze fraud in other kinds of e-services, outside the telecom sector. If this too results in the identification of unknown fraud possibilities, this would confirm the power of the three basic fraud operations used by the proposed approach. Alternatively, one can use the new scenarios to distill a more complete set of fraud patterns and heuristics. Implementing them into the tool's generation

and ranking modules in a customizable way would further the flexibility and applicability of the approach. For instance, it might be the case that in certain scenarios, higher gain does not necessarily imply higher likelihood. Therefore, it is worth exploring alternative heuristics in future research.

Finally, a larger search space also raises issues of resource exhaustion; care must be given to trimming the search space and streamlining code. One way of limiting the search space seems to be differentiating between clients and providers. Then, only clients can be assumed to collude or attempt to bypass payments. Another way of managing the potentially very large lists of results is filtering. Therefore, one of the main topic for improvement in the future is integrating a filter functionality in the tool. Examples filters are: removing sub-ideal models that do not cause a loss, removing sub-ideal models that are not profitable for any of the actors and only showing the most profitable or costly sub-ideal model per collusion type. Furthermore, integrating a model editor into the tool might further increase its usability: firstly, users would need not go through the export/import process for every instance of a model and secondly, users would be able to visualize the fraud directly on the value model.

8

Value-driven identification of sustainability risks using coordination models

Based on *Value-driven risk analysis of coordination models* [16]

Coordination processes are business processes that involve independent profit-and-loss responsible business actors who collectively provide to a customer. Coordination processes are meant to be profitable for the business actors that execute them. However, because business actors are independent, there is also an increased risk of fraud. To compute profitability as well as quantify the risk of fraud, we need to attach value models to coordination process models.

In this chapter, I propose guidelines for deriving a value model from any coordination process model. Next, I show how this approach can be used to identify possibilities of fraud offered by a coordination process, as well as quantify the financial impact of known fraudulent processes. Finally, I discuss additional applications, such as identifying commercially superfluous tasks, or missing tasks needed to achieve a financially sustainable process.

8.1 Introduction

Today, *electronic commercial* services, are an important source of revenue for many businesses. For instance, consider companies such as Netflix, Spotify, or Paypal. Most e-services share two common attributes: (1) they are paid, usually by a customer and (2) they are provided by a complex network of enterprises. As a result, these services are open to opportunities to commit fraud. For example, a fraudulent actor may use the telephone subscription of someone else to place expensive phone calls.

Although fraud is often performed by misusing a business or coordination processes, its impact is actually on the business *value* level. Therefore, we need an instrument to analyze and express its financial effects for all actors involved. In line with related work on value-based fraud analysis [14, 15], we use an e^3*value* model [68] for this purpose. Because a value model represents what actors exchange with each other in terms of *economically valuable* objects (such as products, services or information), it is fundamentally different from a process model. Abstracting away from operational details, e^3*value* models only show *what* is offered, and not *how*.

Unfortunately, for many commercial services, information contained in a value model only exists in the mind of stakeholders, but an explicitly stated model is lacking. Coordination process models, however, often *are* available or can be harvested from existing coordination and orchestration systems [197]. While several approaches can be useful for designing a process model based on given value model, such as the ones described in [173, 184, 198–201], to the best of the authors' knowledge no previous work exists looks at an inverse technique.

This chapter puts forth a new set of guidelines by which an available *BPMN* coordination process model can be used to derive a corresponding e^3*value* model (Sect. 8.3). With the resulting value model, we can use existing tools to identify and prioritize fraud and misuse scenarios (see Sect. 8.4.1), as well as estimate the impact in terms of lost value, and potential gain in terms of misplaced value for the actors involved (see Sect. 8.4.2). Risk assessment could be performed directly on the coordination model. But linking the process layer to the value layer provides a structure approach to quantifying procedural risks. As an added bonus, it can help with rationalizing coordination models in terms of the economic value each activity produces (see Sect. 8.5).

8.2 Related work

Value models and coordination models have different goals and thus represent different types of information. At the same time, they are also related because they express different aspects of the same artifact, namely a set of enterprises and customers aiming to make a profit or to increase their economic value.

When designing a new e-business network, the designer starts with the development of a value model, often as a result of a series of business development workshops. The primary goal is to arrive at a shared understanding amongst the participating enterprises about what they offer each other of economic value, *without* considering *how* these value propositions are executed in terms of operational business processes. This allows identification of potentially profitable e-business models from a management point of view. If a profitable e-business network has been designed, the next step is to assess operational feasibility, which includes

assessing and mitigating risks of fraud. This requires a coordination process model. Schuster et al discuss the design of UMM models from *e*^{3}*value*/REA models [199, 200]. The design of a process model from a value model can be also based on a consideration of trust issues [14, 173] or on the distinction between ownership and possession of value objects [184, 201].

Relating value- and process models can also be viewed as a formal consistency checking problem [202]. Such approaches assume overlap between value models and process models, which need to be kept consistent. For example, in *e*^{3}*value* there is the notion of 'actor', which corresponds well with the idea of 'resource pools' in BPMN. Considering the relation between value and process model as a consistency checking problem is useful because it exposes hints, which can be used to derive process models based on a given value model.

In the remainder of this chapter I consider the case where businesses are already cooperating, but they want to assess the business value of this cooperation, for example in order to assess if the cooperation is still profitable, to assess the economic necessity of all parts of the coordination process, or to assess the potential for fraud, as we do in this chapter. In the real world, many business processes and coordination processes evolve without regular consideration of the underlying value model, and it has been observed earlier that identification of the value proposition of a business process is a key concern of practitioners [203].

In the next section, I show how to derive a value model from an existing process model, in Sect. 8.4 I show how to use the resulting pair of models in fraud analysis and in Sect. 8.5 I discuss further applications.

8.3 From coordination process model to value model

As the value model represents different information about a value web than a coordination process model does, deriving a value model from a process model cannot be fully automated: information needs to be added to as well as deleted from a process model. Moreover, to add this information, value design decisions need to be made, such as which step in a process actually adds value for which actor, how much value is added, and which dependencies exist among economic transactions. These informal decisions - underlined in Fig. 8.2 - cannot be automated, and which of these decisions need to be made differs per process model and depends on the intended value model. The rest of this section elaborates on the derivation process proposed in Fig. 8.2 and gives guidelines for these decisions.

I take the simple process of setting up a new home Internet connection which requires network credentials as a *running example*. This applies to some telephony connections and/or ADSL connections where each user is authenticated to the provider via a username and password that are not linked to the equipment used to enable logical access to the provider's network (e.g. modem).

The normal (i.e. ideal, from the perspective of the provider) process by which a new subscriber requests and receives access to the network is shown in Fig. 8.1. When a customer places an order for a new Internet connection, it triggers the generation of new access credentials. While the user pays for the first month of service, the credentials are sent to him by mail. A technician is scheduled a week or two later to install the necessary equipment (usually a modem). Once the equipment is installed, the credentials can be used to obtain access to the Internet. Note that, for simplicity and didactical reasons, I assume the provider does not wait for the payment to be received before proceeding with setting up the connection. Since

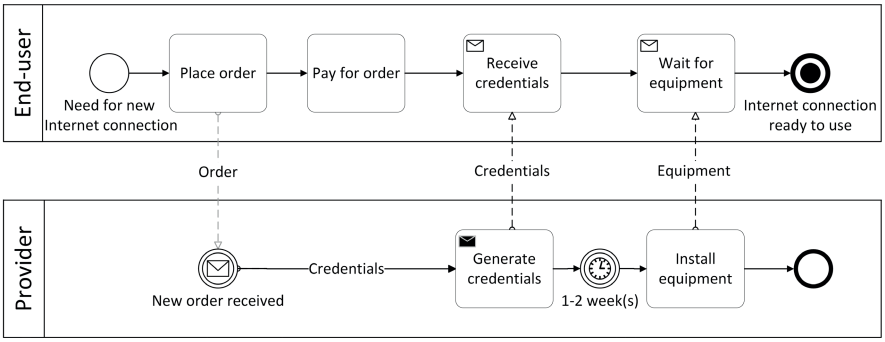

Figure 8.1: Ideal coordination model for setting up a new home Internet connection

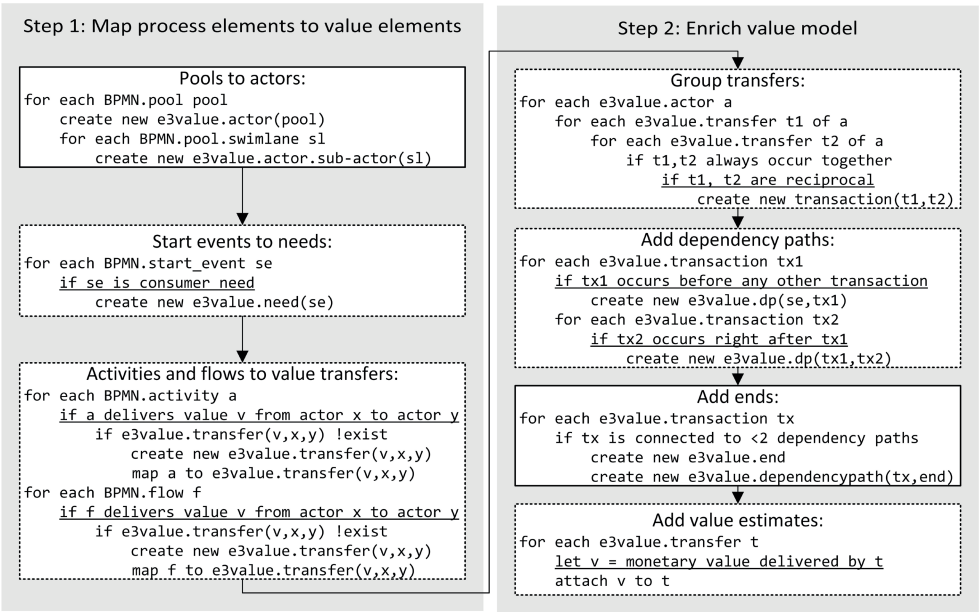

Figure 8.2: Proposed derivation approach: solid boxes can be fully automated; dotted boxes require human decisions (underlined).

modelling physical objects (such as money) as a message is only necessary if their arrival acts as a message event or is expected as input for some activity [3, Section 5.2] I omit modelling the payment as a cross-border message flow.

8.3.1 Mapping process elements to value elements

Several BPMN concepts do have corresponding concepts in e^3*value*. Elements such as BPMN pools, swimlanes, start points and flows share semantic similarities e^3*value* actors, sub-actors, needs and value transfers. I propose the following mapping:

8.3.1.1 Pools to actors:

Instantiate every BPMN pool as an *e³value* actor and every BPMN swimlane as an *e³value* sub-actor.

Running example: BPMN Swimlanes *End-user* and *Provider* are instantiated as *e³value* Actors with the same name.

8.3.1.2 Start events to needs

Select the BPMN Start Event(s) that correspond(s) to consumer need(s). Instantiate as corresponding *e³value* need(s), located at the same actor as the selected start event.

Running example: BPMN Start Event *Need for new Internet connection* becomes an *e³value* Need associated with *End-user*.

8.3.1.3 Activities and/or flows to Value Transfers:

Per activity and per flow, state if they deliver value and to which actor. Then, create a corresponding transfer in the value model.

Guideline: A BPMN activity maps to an e3value activity if and only if the BPMN activity results in a potential profit. In many situations, this is not the case; many BPMN activities are generating costs. Therefore, the mapping from BPMN activities to e3value activities is non-trivial. To find such a mapping, the modeler should ask himself: which BPMN activities and flows relate logically, such that together, they create a profit. Three situations are possible:

- Single activity/flow maps to a transfer
- Multiple activities/flows map to a transfer
- Activity/flow does not map to any transfer. This happens if the flow does not contribute to a transfer of economic value.

Running example:

- BPMN activity *Pay for order* provides (monetary) value to the *Provider*. Therefore, it is mapped to an *e³value* transfer of type *MONEY* which I name *Payment*.
- BPMN activities *Generate credentials*, *Receive credentials* and *Install equipment*, as well as the flows connecting them to each other and to the other components of the process model together provide (service) value to the *End-user*, in the form of Internet access. Therefore, they are grouped into an *e³value* transfer of type *SERVICE* which I name *Internet access*.
- BPMN activity *Place order* and the corresponding message flow only serve as a coordination mechanism and do not provide any value to any of the actors. Therefore, they do not have a corresponding transfer in the value model.

The value model after this first step is shown in Fig. 8.3a.

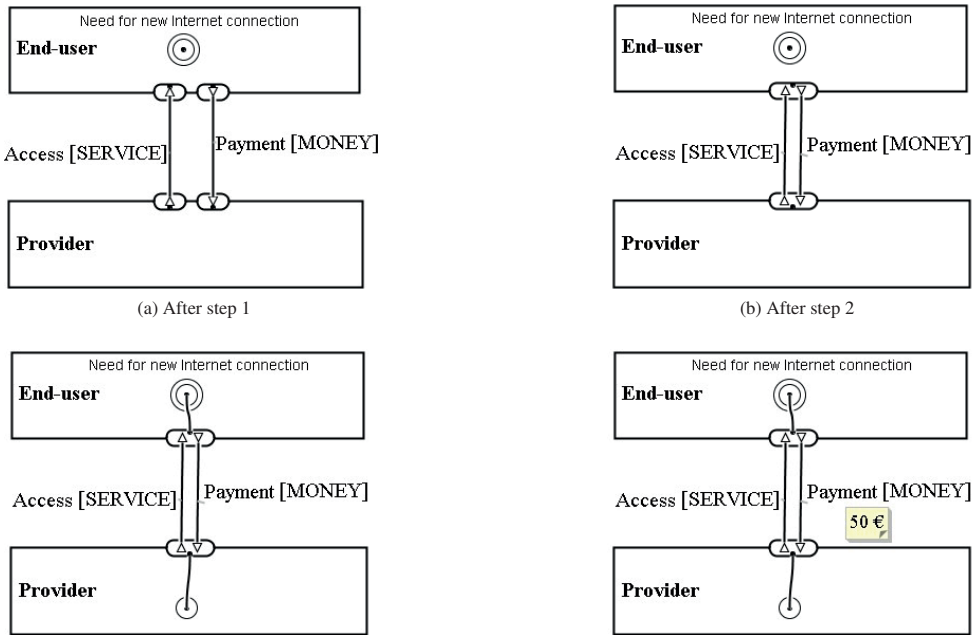

(a) After step 1

(b) After step 2

(c) After step 3

(d) After step 2

Figure 8.3: Evolution of the derived value model for setting up a new home Internet connection

8.3.2 Enriching the value model

8.3.2.1 Group transfers:

Per actor, reciprocal transfers which always occur together should be grouped as part of a single *e³value* interface.

Guideline: For each actor, two transfers he is engaged in are reciprocal (and therefore part of the same interface) if that actor considers that the outgoing transfer provides adequate compensation for what he offers [204]. Note that this does not have to be a on-to-one mapping: an interface may contain any number of incoming and outgoing transfers. While BPMN does not contain sufficient information to decide when two transfers are reciprocal, the execution semantics of BPMN can help ruling out transfers which are not. Specifically, exclusive gateways, event gateways and multiple start events give birth to *alternative paths* [205]. Depending on the conditions of the split or which start event is triggered, activities belonging to one of the alternative paths might not be executed. Conversely, a process model with a single start node and no OR or XOR splits will always terminate, and all activities will be executed [206]. Two *e³value* transfers between the same two actors, can be grouped if and only if *all* BPMN activities that were mapped to these transfers in the previous step are part of the same path. Conversely, if any two activities required for the realization of any of the two transfers are located on alternative paths, then the two transfers should not be part of the same interface.

Running example: The two transfers (*Payment* and *Access*) are reciprocal and part of the same path and can therefore be grouped.

8.3.2.2 Add dependency paths:

Following the sequence and message flow of the original BPMN model from the start, add corresponding dependency paths between the elements of the value model.

*Guideline:*Since *e³value* models lack procedural information such as timing, the goal of this step is not to accurately represent the order in which the transactions take place but rather the causal dependencies between these transaction. Therefore, care must be again given to *alternatives paths*. As a guideline, map parallel gateways to AND nodes and exclusive gateways and event-based gateways to OR nodes. *e³value* OR node are annotated with *fractions*. These fractions should reflect the relative likelihood of the condition-events (in case of event-based gateways) or of the conditions (in case of exclusive gateways).

Running example: we just need to connect the Need to the only transaction.

8.3.2.3 Add ends:

Add *e³value* ends as needed to any transactions without a connection to a dependency path.

Running example: we are left with one transaction which has no outgoing dependency paths so we add an end-point and connect it to this transaction.

8.3.2.4 Add value estimates:

Quantify the value being generated or transfered by the activities in the process model and attach these values to the corresponding transfers.

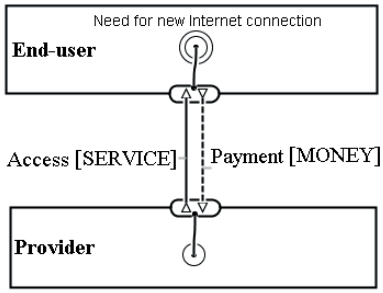

Figure 8.4: Highest ranked sub-ideal model generated by the e3fraud *tool from the model in Fig. 8.3d*

Guideline:e^{3}value provides two ways of attaching monetary values to transfers: if the object being transferred has the same value for both the actors involved, then this value is attached to the transfer itself; otherwise, each individual valuation is attached to the corresponding end of the transfer.

Running example: We add the value of the payment to the "Payment [MONEY]" transfer. We may also add the valuation by any or both of the actors of the "Access [SERVICE]" to the model.

The final value model is shown in Fig. 8.3d.

8.4 Applications to fraud analysis

Once we derive a value model from an *ideal* coordination model, we can leverage previous work on value models in order to run various value-based analyses on it, such as fraud assessment using e3fraud [15]. Or, if we started out with a coordination model which includes fraudulent activities, we can do impact estimation using e3value [68].

8.4.1 Fraud assessment of an ideal coordination process

In this section, I apply the *e^{3}fraud* methodology for automated identification and prioritization of business risks in e-service networks [15] to a derived value model and discuss the implications of the results on the initial process model. The associated *e^{3}fraud* tool[1] can generate possible fraudulent variations, in terms of (1) transactions that might not occur as agreed, (2) transactions that were not expected to occur and (3) collusion, where two or more actors thought to be independent act together against the interests of the provider. It also orders these sub-ideal scenarios based on potential gain for the fraudster, impact for the service provider, or both.

For instance, if we load the derived value model of the simple running example (Fig. 8.3d) into the *e^{3}fraud* tool, breaking transactionality by bypassing the only payment is – quite obviously – identified as the most damaging scenarios. Figure 8.4 shows the corresponding

[1] https://github.com/danionita/e3fraud

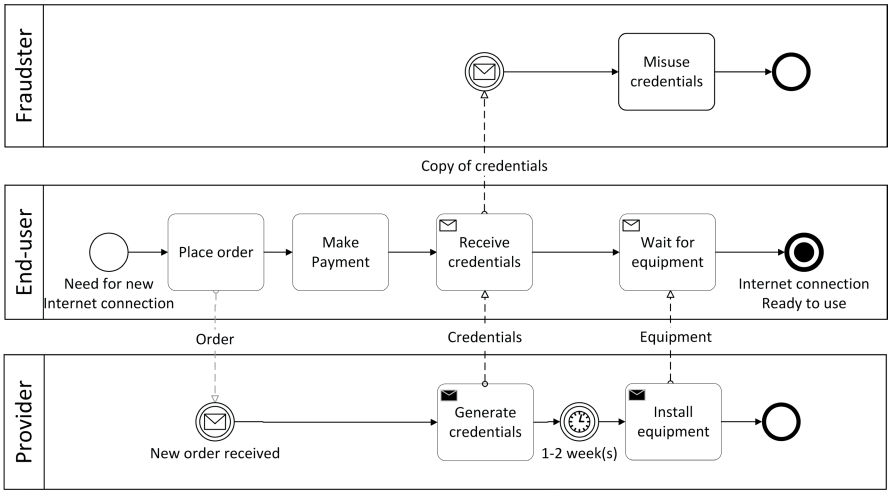

Figure 8.5: Manually created sub-ideal process model of setting up a new home Internet connection

e³fraud model, as generated by the *e³fraud* tool. The *Payment* transfer is marked as dashed to highlight the fact that it does not occur.

Leveraging the decisions made during the derivation process, one can now extend the *e³fraud* analysis by mapping a fraudulent scenario back to the original process model, adding mitigations to this process model, and assessing the impact of those mitigations on the profitability as well as the fraud risks of the value model. This too is a partly automated and partly manual process, and could be a topic for further investigation.

8.4.2 Impact estimation of a sub-ideal coordination process

The above approach allows us to find fraud using a process model and a corresponding value model of the ideal, non-fraudulent way of doing business. In many economic sectors, there are however known process models of fraud. For these process models, the approach can help estimating the economic impact of the fraud by constructing the corresponding value model of the fraud.

For instance, a known vulnerability of the process of setting up a new Internet connection – as described in Sect. 8.3 – involves exploiting the time delay between receiving the credentials and the physical installation of the equipment by a technician. A BPMN model of this fraudulent process is shown in Fig. 8.5.

By applying the proposed model transformation steps, we end up with a corresponding value model of the fraud, as shown in Fig. 8.6. We can now evaluate this value model using the established *e³value* profitability analysis ([68]) to estimate the profit made by the fraudster as well as the associated costs for the Internet provider. Furthermore, we can apply extensions of *e³value* aimed at analyzing sub-ideal value models – such as *e³control* [173] – in order to help implement preventive measures.

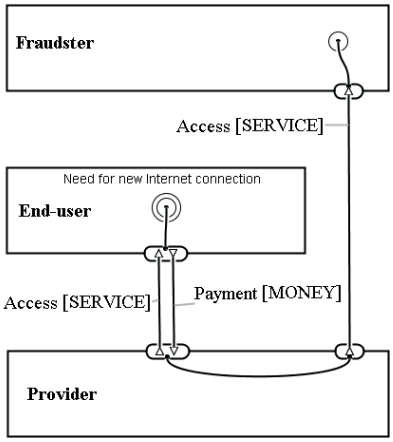

Figure 8.6: Value model derived from the model in Fig. 8.5

8.5 Case study: the roaming service

In the previous sections I used a simple, didactic example to introduce the proposed derivation approach and and demonstrate how it can be used to assess the potential for fraud in an existing non-fraudulent process as well as to estimate the financial impact of a fraudulent process. Next, I test the approach on a realistic business model obtained from a telecom service provider in order to discuss alternative applications which were not visible in the first example. Specifically, I investigate if we can leverage the process-to-value mapping created as part of the derivation process to identify potential risks related to the commercial feasibility of a coordination process. With this purpose in mind, I obtained and analyzed a coordination model of the process of calling from abroad, also known as roaming. This is a telephony service which involves multiple providers (both the home and the visited provider need to collaborate to provide the service) and several payments (between providers and between providers and the user).

The ideal process by which roaming services are provided and charged is shown in Fig. 8.7. The process is triggered when the subscriber receives or initiated a call. Calling is a looping activity that triggers a technical sub-process (mobile subscriber identification, network routing, and so on). When a call is ended, a record of that call is saved. At fixed intervals, call records are billed and these bills are sent to the respective home providers. In turn, the home provider performs a corresponding payment and adds these costs to the subscriber's monthly bill.

We derive a value model from the process model shown in Fig. 8.7 above by applying the transformation steps described in Sect. 8.3.

1. Map process elements to value elements:

    1.1. Pools to actors: BPMN swimlanes *Subscriber (roamer)*, *Visited TSP* and *Home TSP* are instantiated as e^3*value* Actors.

    1.2. Start events to needs: BPMN start event becomes an e^3*value* Need.

    1.3. Activities and flows to value transfers:

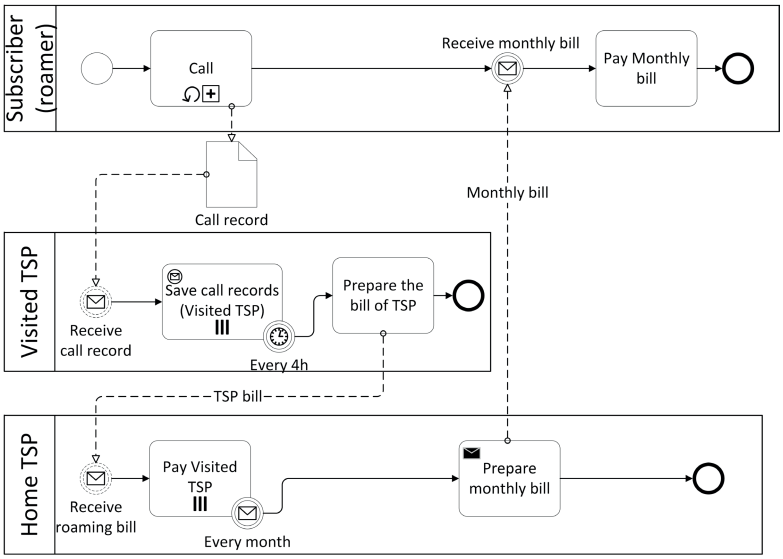

Figure 8.7: Ideal process model - roaming service

- BPMN activity *Call* provides (service) value to the *Subscriber / Roamer*. Therefore, it is mapped to an e^3*value* transfer of type *SERVICE* which we also name *Call*.
- BPMN activity *Prepare the bill of TSP*, as well as the event *Receive the roaming bill* and the message flow connecting them together provide value to the *Home TSP* (he now knows that his subscriber performed roaming calls and can charge accordingly). Therefore, they are grouped into an e^3*value* transfer of type *Proof of roaming* which we name *Bill to TSP*.
- BPMN activities *Prepare monthly bill*, the event *Receive monthly bill* and the message flow connecting them together provide value to the *Subscriber Roamer* (he now knows how much he has to pay for his roaming usage). Therefore, they are grouped into an e^3*value* transfer of type *Proof of roaming* which we name *Bill to customer*.
- BPMN activity *Pay monthly bill* provides (monetary) value to *Home TSP*. Therefore, it is mapped to an e^3*value* transfer of type *MONEY* which we name *Customer payment*
- BPMN activity *Pay the roaming costs* provides (monetary) value to *Visiting TSP*. Therefore, it is mapped to an e^3*value* transfer of type *MONEY* which we name *TSP payment*.
- BPMN activity *Save call records* only serves as a logging mechanism and do not provide any immediate value to any of the actors. Therefore, these activities do not have a corresponding transfer in the value model.

2. Enrich value model

2.1. Group transfers into transactions:
   - The two transfers *TSP payment* and *Bill to TSP* are reciprocal and therefore grouped into a single transaction.
   - The two transfers *Customer payment* and *Bill to customer* are reciprocal and therefore grouped into a single transaction.
   - The transfer *Call* has no reciprocal transfer. This unusual result is discussed in Sect. 8.5.1 below.

2.2. Add dependency paths: Using the message flows of the BPMN model as a guide, we add dependency paths accordingly.

2.3. Add end(s): We add an e^3*value* End-point to the transaction without a dependency path.

2.4. Add value estimates: We add the value of the payment to the "Customer Payment [MONEY]" and "TSP Payment [MONEY]" transfers. We may also add the valuation by any or both of the actors of the "Call [SERVICE]" or of the two bills to the model.

The resulting value model is shown in Fig. 8.8. Note that the transfer *Call* has no reciprocal transfer. This unusual result is discussed next.

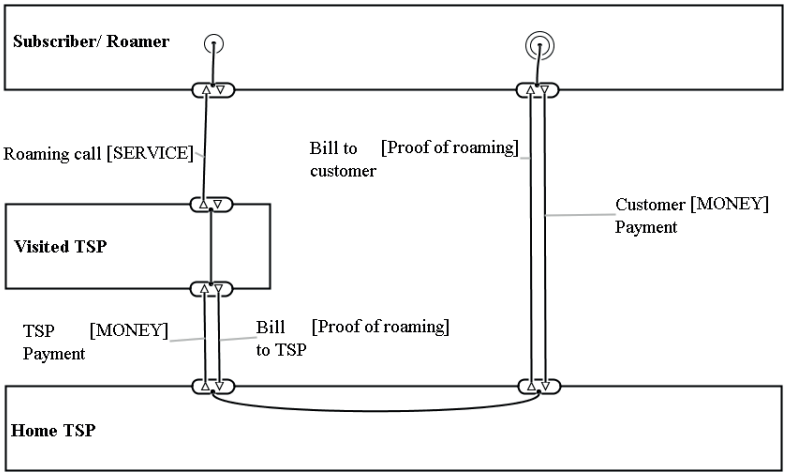

Figure 8.8: Ideal value model - roaming service

8.5.1 Non-reciprocal transfers

In the value model shown in Fig. 8.8, obtained by applying the proposed derivation method to the process model shown in Fig. 8.7, there is a transaction consisting of a single, non-reciprocal transfer: the *Call[SERVICE]*.

A non-reciprocal transfer implies that something realizing the reciprocal value transfer is missing from the initial process model. Such duality problems in value models created with the proposed derivation approach can have one of two causes: either (1) there is a problem

in the process or (2) there is a problem in the model of the process. The former is indicative of a financially unfeasible process, which means the process is either altruistic or fraudulent. The latter means that the activities or flows realizing the reciprocal transfer are intentionally left out (i.e. are out of scope of the model) or unintentionally omitted (as a result of improper modelling or poor data quality). Deciding which cause applies in a certain case is important as it might trigger a re-design of the process or an update of the model. Checking the process model against consistency rules, such as realizability [207], local enforceability [208] and desynchronizability [209] might help identify which of the three situations described above we are in and why, but this is subject of further research.

In the example of Fig. 8.8, the process model is incomplete: it lacks information with regard to what is provided by the customer whenever he wishes to make a call, namely an identifier (commonly referred to as an IMSI[2] in telecom) which acts as proof that the subscriber has the right to perform roaming calls. This right was obtained when the SIM card was first purchased or is included in the monthly subscription fee.

The fact that reciprocal value-producing tasks missing from the process model will result in incomplete or incorrect value models when transformed using the proposed approach suggests that one could also use the approach proposed in this chapter to rationalize and validate coordination models and processes in terms of their financial sustainability.

8.5.2 Superfluous activities

Another result of the derivation of a value model from the process model in Fig. 8.7, is that one of the tasks - namely, *Save call records* was not identified as being part of a value transfer. Similarly, in Sect. 8.3.1 we did not map the *Place order activity* of Fig. 8.1 to a value transfer. This indicates that these tasks do not have a commercial purpose. Therefore, from a commercial point of view, the process can be streamlined by eliminating the task. But perhaps from another point of view, e.g. auditing, the task may still have to be included. Whichever the case, but observations such as these could provide a starting point for process optimization activities.

8.6 Conclusions and future work

The derivation approach described in this chapter can be used to construct a value model from a multi-actor BPMN process model, which in turn allows profitability analysis of the original process model, the generation of fraudulent variations of the process model, and the analysis of the financial effects of changes (fraudulent or not) in the process model. By starting with a fraudulent instead of an ideal process model, it can also be used to estimate the impact of fraud and of fraud-mitigating measures.

The derivation approach proposed is feasible: it was applied by two authors independently to two case studies and was found to produce very similar results. Of course, further real-world validation is needed to get a better idea as to how difficult and error-prone the derivation process is.

[2]International Mobile Subscriber Identity, used to identify the user of a cellular network and is a unique identification associated with all cellular networks. [210]

I envision supporting the process by means of a software tool. Such a tool would implement the algorithmic part of Fig. 8.2, and guide the user through the non-automatable decisions he/she has to make. Another related topic which merits further investigation is whether a similar tool could use these decisions to maintain dynamic consistency between the two models, thereby supporting a wider range of applications, such as sensitivity analyses.

I believe that associating value models to coordination process models empowers the organization by promoting an understanding of the value creation activities inside the process and allowing usage of value analysis tools, such as income/cost estimations and fraud assessment.

Part V

Conclusions

9

Overarching conclusions and future work

In this dissertation, I have explored a multitude of ways in which conceptual models can strengthen, enhance and augment the risk assessment of socio-technical systems. We looked at how tangible models can increase the efficiency and effectiveness of collaborative Target of Assessment modelling tasks, at how argumentation models can help maintain an overview of risks and risk mitigation decisions both during and after a risk assessment, and at how value models can be used to quantify as well as identify risks.

In this final chapter, I summarize the lessons learned and use them to argue about the applicability of tangible modelling, argumentation modelling, and value modelling to information security risk assessment. I then discuss the complementarity of the proposed approaches. Finally, I return to my original research questions, discuss to which extent I've answered them, and suggest avenues for future work.

9.1 Conclusions

9.1.1 Tangible modelling and its role in risk assessment

Overall, the three experiments provide evidence that iconicity not only improves understandability but also modelling speed and model quality and that tangibility promotes collaboration by facilitating uniform participation of all group members. The experiments also suggest that tangibility magnifies the positive effects of iconicity as well as the negative effects of abstractness on understandability, modelling speed, and model quality.

This has implications for risk assessment which rely on accurate models of the Target of Assessment, such as those performed by (semi-)automated risk analysis tools. Consultants, security experts, and auditors could also find value in applying tangible modelling principles to their interaction with customers. Even generic risk assessment frameworks could integrate collaborative tangible modelling as part of context establishment phase. Finally, the results of the modelling experiments are also relevant to researchers looking into collaborative and participative enterprise modelling, as well as to the broader field of conceptual modelling.

9.1.2 Argumentation modelling and its role in risk assessment

Good security is invisible and perfect security, impossible. Security arguments can show that a system is secure to some extent by providing structured, but human-readable explanations as to which risks were considered and how they were mitigated. This is important for a variety of reasons, ranging from certification to compliance, and from awareness to assurance.

Unsurprisingly, most argumentation modeling tools employ a simplified version of Toulmin's argument structure for conceptualizing security arguments but vary in terms of either the granularity by which they decompose the argument or in the way they represent inter-argument structures. However, very few tools exist which address the specifics of security argumentation, and their audience is mostly academic.

Indeed, confronting the argumentation tools discussed in Chapter 4 with practical security arguments shows that in order to be usable, security argumentation techniques need to be simple and reduce themselves to the essential information that needs to be present in order to argue about (in-)security of a system or software: the links between mitigations, risks, and system components or modules. As these links can be of type "many-to-many", graphs are a natural fit for representing these links.

9.1.3 Value modelling and its role in risk assessment

Value models provide an established way of modelling the co-creation of value by independent profit-loss responsible entities. But co-creating value may bring about new vulnerabilities stemming from the potentially falsifiable assumptions that have to be made about the behavior of third-parties and even customers. Chapters 6 to 8 provide concrete examples of how existing value-driven analyses can be extended to help understand and quantify the effects of such vulnerabilities. As shown in those chapters, value models can serve as a useful tool not only for estimating the profitability of a (new) e-business idea but – with the proposed extensions – they can also be used for quantitative risk and sensitivity analysis.

In their recent review of value models, Weigand and Jeewanie propose strengthening their connection with management research (Weigand, Jeewanie, 2009) so as to leverage their systemic perspective on how to do business (Zott, Amit, Massa, 2011), with a focus on value creation, delivery and capture (Lambert, 2010). The conceptual addition described in Chapter 6, the methodological addition of Chapter 7 and derivation approach of Chapter 8 moves e^3*value* in the direction of a (security) decision support framework, and this constitutes a strengthening of the systematic perspective on how to do business. Partial automation of the value model-driven risk assessment process can speed up the analysis process, while visualizations such as charts and graphs enrich the analysis by improving understandability of its results. Being model-driven approaches also facilitate re-use, as well as being constructionist in nature.

9.1.4 Complementarity

The techniques described in this dissertation are supplementary rather than complementary, in that they each focus on a different aspect of risk assessment and can be used independently of each other or in combination with other risk analysis techniques. The tangible modelling approach of Part II focuses on obtaining a clearer picture of the Target of Assessment. The argumentation-based techniques of Part III aim to improve the traceability of risk identification, analysis, and mitigation. The value-driven methods of Part IV are geared towards quantifying risks, and provides a specialized means for identifying fraud and sustainability risks.

The insights on collaborative modelling of Part II could be used to construct a model to serve as a reference for one of the argumentation-based risk assessment techniques of Part III, but also for other risk assessment techniques. The risk quantification and fraud identification methods and tools of Part IV could inform and support the decisions made as part of a risk argumentation session such as the ones proposed in Part III. While within each Part, Chapters share a conceptual model and are often iterations of each other, there is no shared conceptual or procedural relationship across Parts II, III, and IV. This is because each part started out as a solution to a different problem: getting to grips with a formal ToA modeling language, providing structure to risk assessment brainstorming meetings, and quickly estimating the risk of fraud, respectively.

9.2 Answers to research questions

I started my doctoral research with the goal of *improving information security risk assessment in a model-driven way without unnecessary quantification*. To this end, I've proposed several modelling, analysis and estimation techniques all of which rely on explicit conceptual models and are aimed at streamlining, structuring or otherwise improving the risk assessment of modern socio-technical information systems without requiring hard-to-obtain data such as likelihood and impact estimations. None of the approaches have been used "in the wild", and some remain at a proof-of-concept level. Nevertheless, the insights gathered during their development and validation helps shed light on the three research questions I posed at the start:

RQ1 *How can the effort and resources required to perform an IS risk assessment be reduced?'*
In Part II, I looked at how different ways to model and represent a socio-technical system for the purpose of risk assessment play a role in improving the process of defining the Target of Assessment. I uncovered evidence that the iconicity of the modelling

elements is correlated with the efficiency and efficacy of the modelling task, and that tangibility amplifies these effects and increases collaboration when modelling as a group. However, the most time- and resource-consuming part of a Risk Assessment is often risk identification and analysis [32]. To mitigate this, one can perform a qualitative risk assessment - such as described in Part III - or an automated risk assessment. Tangible and iconic modelling may assist in obtaining a Target of Assessment model which is correct and complete enough to be fed into such an automated tool.

RQ2 *How can the defensibility, understandability and re-usability of risk mitigation decisions be improved?* In order to revisit risk mitigation decisions, the reasoning behind this decision should to be available. In Part III, I showed that argumentation modelling can be used to encode the arguments put forth during a risk assessment. Based on my experiments it appears that for the purpose of defending mitigation decisions, as well as informing similar decisions in the future, it is necessary to store at least two arguments: a risk argument and a mitigation argument. The risk argument - called an "attacker" argument in Part III - explains why the risk is pertinent to the given system. The mitigation argument - called a "defender argument" in Part III - explains how the proposed countermeasure addresses one or more risks. Storing only these two arguments keeps the argumentation model at a human-readable level of complexity. A tree structure appears to provide good understandability of such a model. In addition, relating these arguments to a model of the Target of Assessment provides a foundation for more formal analyses as well as for navigating (e.g. by filtering) the argumentation model.

RQ3 *How can IS risk assessments be better integrated with established enterprise processes?* In order to integrate better with existing enterprise processes, IS risk assessment can (1) take existing enterprise models as input and (2) produce results which can be mapped back to enterprise models or fed back into enterprise processes. In Part IV, I showed how fraud and sustainability risk assessment can be performed based solely on existing business process models and business value models. In addition, the results produced are designed to inform the fraud detection and prevention processes. In Part III, I introduced two argumentation-based risk assessment methods which are capable of modelling risk arguments and mitigation decisions an mapping them to any kind of enterprise system model. Qualitative, brainstorming-based risk assessment methods are often employed in practice. Argument-modelling provides some structure to these meetings but is flexible enough to be applied within virtually any established enterprise process.

9.3 Future work

9.3.1 Tangible modelling

Two major limitations of our tangible modelling experiments were the low sample size and the unavailability of experts. Furthermore, there is reason to believe that confounding factors, mostly related to the personality of the participants, have skewed the results. It is likely that some of the positive effects are more pronounced in certain demographics and are only significant with certain types of socio-technical modelling languages. Future work should

therefore attempt to reproduce the experiments with larger, more diverse samples in an attempt to replicate and refine the effects I observed. It is worth starting from modelling languages whose concepts are easy to map to iconic graphical representation.

9.3.2 Argumentation modelling

In the words of Buckingham-Shum [211], diagramming tools differ not only in the type of information they are able to represent but especially in regard to the trade-off they make between expressiveness and usability. This is true also for argumentation graphs, which can explode in size when all known risks and relevant mitigations pertaining to a real system are added. Therefore, ensuring scalability is critical to maintaining reasonable usability. Only some of the tools available provide ways of navigating the graph, for example by searching, filtering or collapsing parts of the argumentation structure. We believe this topic has to be better investigated before security argumentation modeling becomes usable in practice. To further enhance scalability, automation and re-usability are also relevant topics not only in security argumentation, but security in general. Future work could look therefore into ways by which the argumentation graph can be filled in semi-automatically, for instance by recognizing patterns, linking to knowledge bases or parsing the output of vulnerability scanners.

9.3.3 Value modelling

The proposed value-driven risk analysis approaches are easily extensible. Future work might reveal fraud heuristics other than the three defined in *e^3fraud* (collusion, non-occurrence and hidden transactions), which can be added to the conceptual meta-model used by the automated fraud generation engine. Model patterns may be developed and used as templates, thereby increasing the usability and efficiency of the approach. Finally, integration with enterprise systems can feed values into the value models, ensuring up-to-date valuations as well as reducing the amount of error prone manual work.

References

[1] Mass Soldal Lund, Bjørnar Solhaug, and Ketil Stølen. *Model-driven risk analysis: the CORAS approach*. Springer Science & Business Media, 2010.

[2] Wolter Pieters, Dina Hadziosmanovic, Aleksandr Lenin, Lorena Montoya, and Jan Willemson. *TREsPASS: plug-and-play attacker profiles for security risk analysis*. IEEE Security & Privacy poster abstracts, 2014.

[3] Object Management Group (OMG). *BPMN 2.0 by Example*. Technical Report dtc/2010-06-02, Object Management Group (OMG), June 2010.

[4] Net Losses. *Estimating the global cost of cybercrime*. McAfee, Centre for Strategic & International Studies, 2014.

[5] Mathias Weske. *Business process management: concepts, languages, architectures*. Springer Science & Business Media, 2012.

[6] LLC. Social-Engineer. *The Social Engineering Infographic*. `https://www.social-engineer.org/social-engineering/social-engineering-infographic/`, April 2014.

[7] Business Continuity Institute (BCI). *Cyber Resilience Report*, 2017.

[8] Gurpreet Dhillon and James Backhouse. *Current Directions in Is Security Research: Towards Socio-organizational Perspectives*. Information Systems Journal, 11(2):127–153, 2001.

[9] John Organ and Larry Stapleton. *Information systems risk through a socio-technical lens: future directions in systems risk research*. IFAC Proceedings Volumes, 45(10):138–143, 2012.

[10] Konstantinia Charitoudi and Andrew Blyth. *A socio-technical approach to cyber risk management and impact assessment*. Journal of Information Security, 4(01):33, 2013.

[11] Tyler W Moore, Christian W Probst, Kai Rannenberg, and Michel van Eeten. *Assessing ICT Security Risks in Socio-Technical Systems (Dagstuhl Seminar 16461)*. In Dagstuhl Reports, volume 6. Schloss Dagstuhl-Leibniz-Zentrum fuer Informatik, 2017.

[12] Dimitris Raptis, Theodosis Dimitrakos, Bjørn Axel Gran, and Ketil Stølen. *The coras approach for model-based risk management applied to e-commerce domain*. In Borka Jerman-Blazic and Tomaz Klobucar, editors, Communications and Multimedia Security, volume 228 of *IFIP Conference Proceedings*, pages 169–181. Kluwer, 2002.

[13] Dan Ionita, Roel Wieringa, Jaap Gordijn, and Ahmed Yesuf. *Quantitative, Value-driven Risk Analysis of e-Services*. Journal of Information Systems, 2017. Submitted, pending acceptance.

[14] Dan Ionita, Roel J Wieringa, Lars Wolos, Jaap Gordijn, and Wolter Pieters. *Using value models for business risk analysis in e-service networks*. In IFIP Working Conference on The Practice of Enterprise Modeling, pages 239–253. Springer International Publishing, 2015.

[15] Dan Ionita, Roel J Wieringa, and Jaap Gordijn. *Automated identification and prioritization of business risks in e-service networks*. In International Conference on Exploring Services Science, pages 547–560. Springer International Publishing, Springer International Publishing, May 2016.

[16] Dan Ionita, Jaap Gordijn, Ahmed Seid Yesuf, and Roel Wieringa. *Value-driven risk analysis of coordination models*. In IFIP Working Conference on The Practice of Enterprise Modeling, pages 102–116. Springer International Publishing, 2016.

[17] Alexandr Vasenev, Dan Ionita, Tomasso Zoppi, Andrea Ceccarelli, and Roelf J. Wieringa. *Towards security requirements: Iconicity as a feature of an informal modeling language*, 2017.

[18] Alexandr Vasenev, Lorena Montoya, Andrea Ceccarelli, Anhtuan Le, and Dan Ionita. *Threat navigator: grouping and ranking malicious external threats to current and future urban smart grids*. In Smart Grid Inspired Future Technologies: First International Conference, SmartGIFT 2016, Liverpool, UK, May 19-20, 2016, Revised Selected Papers, pages 184–192. Springer International Publishing, 2017.

[19] Dan Ionita and Roel Wieringa. *Web-based Collaborative Security Requirements Elicitation.* In REFSQ Workshops, volume 1564. CEUR-WS, March 2016.

[20] Dan Ionita, Roel Wieringa, Jan-Willem Bullee, and Alexandr Vasenev. *Tangible modelling to elicit domain knowledge: an experiment and focus group*. In P. Johannesson, M. Li Lee , S. W. Liddle , A. L. Opdahl, and López , editors, International Conference on Conceptual Modeling, volume 9381 of *Lecture Notes in Computer Science*, pages 558–565, Berlin, October 2015. Springer International Publishing, Springer. LNCS 9381.

[21] Alexandr Vasenev, Lorena Montoya, and Dan Ionita. *Outlining an 'Evaluation Continuum': Structuring Evaluation Methodologies for Infrastructure-Related Decision Making Tools*, pages 243–249. Springer International Publishing, Cham, 2017.

[22] Henry Prakken, Dan Ionita, and Roel Wieringa. *Risk Assessment As an Argumentation Game.* In João Leite, Tran Cao Son, Paolo Torroni, Leon van der Torre, and Stefan Woltran, editors, International Workshop on Computational Logic in Multi-Agent Systems, volume 8143 of *Lecture Notes in Computer Science*, pages 357–373. Springer Berlin Heidelberg, Springer, 2013.

[23] Dan Ionita, Jan-Willem Bullee, and Roel J. Wieringa. *Argumentation-Based Security Requirements Elicitation: The Next Round*. In Proceedings of the 2014 IEEE 1st International Workshop on Evolving Security and Privacy Requirements Engineering (ESPRE), pages 7–12. Springer, Aug 2014.

[24] D. Ionita, R. Kegel, A. Baltuta, and R. Wieringa. *ArgueSecure: Out-of-the-Box Security Risk Assessment*. In 2016 IEEE 24th International Requirements Engineering Conference Workshops (REW), pages 74–79, September 2016.

[25] Dan Ionita, Julia Kaidalova, Alexandr Vasenev, and Roel Wieringa. *A study on tangible participative enterprise modelling*. In International Conference on Conceptual Modeling, pages 139–148. Springer International Publishing, To be published, 2016.

[26] Dan Ionita, Margaret Ford, Alexandr Vasenev, and Roel Wieringa. *Graphical modeling of Security Arguments: Current state and future directions*. In International Workshop on Graphical Models for Security. Springer, 2017.

[27] Dan Ionita, Denieve Nazareth, Alexandr Vasenev, Frank van der Velde, and Roel Wieringa. *The role of tangibility and iconicity in collaborative modelling tasks*. In Conceptual Modeling: Research in Progress. CEUR, 2017.

[28] Dan Ionita. *Context-sensitive Information security Risk identification and evaluation techniques*. In Requirements Engineering Conference (RE), 2014 IEEE 22nd International, pages 485–488. IEEE, 2014.

[29] D Ionita, PH Hartel, W Pieters, and R Wieringa. *Current established risk assessment methodologies and tools, July 2013*. Technical report, University of Twente, 2013.

[30] Dan Ionita, SK Koenen, and RJ Wieringa. *Modelling telecom fraud with e3value*. Technical report, University of Twente, 2014.

[31] D. Ionita, R.J. Wieringa, J.H. Bullee, and A. Vasenev. *Investigating the usability and utility of tangible modelling of socio-technical architectures*. Technical report TR-CTIT-15-03, University of Twente, June 2015.

[32] Dan Ionita. *Current Established Risk Assessment Methodologies and Tools*. Master's thesis, University of Twente, July 2013.

[33] Atle Refsdal, Bjørnar Solhaug, and Ketil Stølen. *Cyber-risk management*. In Cyber-Risk Management, pages 33–47. Springer, 2015.

[34] Donald A Norman. *Some observations on mental models*. Mental models, 7(112):7–14, 1983.

[35] Jan-Willem Bullee, L. Montoya, Marianne Junger, and Pieter H. Hartel. *Telephone-based social engineering attacks: An experiment testing the success and time decay of an intervention*, pages 107–114. Cryptology and Information Security Series. IOS Press, 1 2016. Foreground = 100Type of audience = scientific community, industry; Size of audience = 150; Countries addressed = international;.

[36] Sarah Granger. *Social engineering fundamentals, part I: hacker tactics*. Security Focus, December, 18, 2001.

[37] Ira S Winkler and Brian Dealy. *Information Security Technology? Don't Rely on It. A Case Study in Social Engineering.* In USENIX Security Symposium, volume 5, pages 1–1, 1995.

[38] Thomas R Peltier. *Social engineering: Concepts and solutions*. Information Systems Security, 15(5):13–21, 2006.

[39] Jan-Willem H Bullée, Lorena Montoya, Wolter Pieters, Marianne Junger, and Pieter H Hartel. *Regression Nodes: Extending attack trees with data from social sciences*. In Socio-Technical Aspects in Security and Trust (STAST), 2015 Workshop on, pages 17–23. IEEE, 2015.

[40] Michael Nidd, Marieta Georgieva Ivanova, Christian W. Probst, and Axel Tanner. *Tool-based risk assessment of cloud infrastructures as socio-technical systems*. In Ryan Ko and Raymond Choo, editors, The cloud security ecosystem, pages 495–517. Elsevier, Syngress, Amsterdam, June 2015.

[41] Lotfi ben Othmane, Rohit Ranchal, Ruchith Fernando, Bharat Bhargava, and Eric Bodden. *Incorporating attacker capabilities in risk estimation and mitigation*. Computers & Security, 51:41–61, 2015.

[42] Zahra Mohaghegh and Ali Mosleh. *Incorporating organizational factors into probabilistic risk assessment of complex socio-technical systems: Principles and theoretical foundations*. Safety Science, 47(8):1139–1158, 2009.

[43] M.S. Lund, B. Solhaug, and K. St2len. *Model-driven risk analysis*. Springer Berlin Heidelberg, 2011.

[44] J.O. Aagedal, F. den Braber, T. Dimitrakos, B.A. Gran, D. Raptis, and K. Stolen. *Model-based risk assessment to improve enterprise security*. In Enterprise Distributed Object Computing Conference, 2002. EDOC '02. Proceedings. Sixth International, pages 51–62, 2002.

[45] Henk Jonkers, Marc Lankhorst, Rene van Buuren, Stijn Hoppenbrowers, Marcello Bosangue, and Leendert van der Torre. *Concepts for Modelling Enterprise Architectures*. International Journal of Cooperative Information Systems, 13(03):257–287, 2004.

[46] J. Stirna and A. Persson. *Ten Years Plus with EKD: Reflections from Using an Enterprise Modeling Method in Practice*. In Proceedings of the 11th International Workshop on Exploring Modeling Methods in Systems Analysis and Design (EMMSAD07). Springer, 2007.

[47] Mark S Fox and Michael Gruninger. *Enterprise modeling*. AI magazine, 19(3):109, 1998.

[48] Frank Innerhofer-Oberperfler and Ruth Breu. *Using an Enterprise Architecture for IT Risk Management.* In ISSA, pages 1–12, 2006.

[49] John A Zachman. *A framework for information systems architecture*. IBM systems journal, 26(3):276–292, 1987.

[50] Ronda R Henning et al. *Use of the Zachman architecture for security engineering*. Available at:¡ csrc. nist. gov/nissc/1996/papers/NISSC96/paper044/baltppr. pdf¿[Accessed 1 May 2011], 1996.

[51] L Lori DeLooze. *Applying security to an enterprise using the Zachman framework*. SANS Institute, 2001.

[52] Levent Ertaul and Raadika Sudarsanam. *Security Planning Using Zachman Framework for Enterprisesy*, 2005.

[53] German BSI. *BSI Standards 100-1, 100-2, 100-3, 100-4*. `https://www.bsi.bund.de/EN/Publications/BSIStandards/standards.html`, 2013.

[54] Bomil Suh and Ingoo Han. *The IS risk analysis based on a business model*. Information & Management, 41(2):149–158, 2003.

[55] Kurt Sandkuhl, Janis Stirna, Anne Persson, and Matthias Wißotzki. *Enterprise Modeling: Tackling Business Challenges with the 4EM Method*. Springer, Heidelberg, 2014.

[56] J. S. Park, B. Montrose, and J. N. Froscher. *Tools for Information Security Assurance Arguments*. In Proceedings of the DARPA Information Survivability Conference (DISCEX '01), volume 1, pages 287–296, 2001.

[57] R. E. Bloomfield, S. Guerra, A. Miller, M. Masera, and C. B. Weinstock. *International Working Group on Assurance Cases (for Security)*. IEEE Security Privacy, 4(3):66–68, May 2006.

[58] Charles B Haley, Robin Laney, Jonathan D Moffett, and Bashar Nuseibeh. *Arguing Satisfaction of Security Requirements*. Integrating security and software engineering: Advances and future visions, pages 16–43, 2006.

[59] Patchin Curtis and Mark Carey. *Risk assessment in practice*. Technical report, Deloitte, 2012.

[60] Leyla Bilge and Tudor Dumitras. *Before we knew it: an empirical study of zero-day attacks in the real world*. In Proceedings of the 2012 ACM conference on Computer and communications security, pages 833–844. ACM, 2012.

[61] Thomas Dufresne and James Martin. *Process modeling for e-business*. Î n: Information Systems Department, George Mason University, L. Kerschberg (ed), 2003.

[62] Matjaz B. Juric. *Business Process Execution Language for Web Services BPEL and BPEL4WS 2Nd Edition*. Packt Publishing, 2006.

[63] Chun Ouyang, Wil M. P. van der Aalst, Marlon Dumas, and Arthur. *Translating BPMN to BPEL*. Technical Report BPM-06-02, BPM Center, 2006.

[64] Stephen A. White. *Using BPMN to Model a BPEL Process*. BPTrends, 3:1–18, 01 2005.

[65] Chun Ouyang, Marlon Dumas, Arthur H.M. ter Hofstede, and Wil M.P. van der Aalst. *Pattern-based translation of BPMN process models to BPEL web services*. International Journal of Web Services Research (JWSR), 5(1):42–62, 2007.

[66] Matthias Weidlich, Gero Decker, Alexander Großkopf, and Mathias Weske. *On the Move to Meaningful Internet Systems: OTM 2008: OTM 2008 Confederated International Conferences, CoopIS, DOA, GADA, IS, and ODBASE 2008, Monterrey, Mexico, November 9-14, 2008, Proceedings, Part I*, chapter BPEL to BPMN: The Myth of a Straight-Forward Mapping, pages 265–282. Springer Berlin Heidelberg, 2008.

[67] Jan C. Recker and Jan Mendling. *On the Translation between BPMN and BPEL: Conceptual Mismatch between Process Modeling Languages*. In Thibaud Latour and Michael Petit, editors, 18th International Conference on Advanced Information Systems Engineering. Proceedings of Workshops and Doctoral Consortiums, pages 521–532. Namur University Press, 2006.

[68] Jaap Gordijn and Hans Akkermans. *Designing and Evaluating E-Business Models*. IEEE Intelligent Systems, 16(4):11–17, July 2001.

[69] Hans Weigand. *The e3value Ontology for Value Networks: Current State and Future Directions*. Journal of Information Systems, 30(2):113–133, 2016.

[70] Birger Andersson, Maria Bergholtz, Ananda Edirisuriya, Tharaka Ilayperuma, Paul Johannesson, Jaap Gordijn, Bertrand Grégoire, Michael Schmitt, Eric Dubois, Sven Abels, et al. *Towards a reference ontology for business models*. In International Conference on Conceptual Modeling, pages 482–496. Springer, 2006.

[71] Reza Samavi, Eric Yu, and Thodoros Topaloglou. *Strategic reasoning about business models: a conceptual modeling approach*. Information Systems and e-Business Management, 7(2):171–198, 2009.

[72] Alexander Osterwalder and Yves Pigneur. *An eBusiness model ontology for modeling eBusiness*. BLED 2002 Proceedings, page 2, 2002.

[73] Robert F Lusch, Stephen L Vargo, and Gregor Wessels. *Toward a conceptual foundation for service science: Contributions from service-dominant logic*. IBM systems journal, 47(1):5–14, 2008.

[74] Hans Weigand. *Value encounters–modeling and analyzing co-creation of value*. In Conference on e-Business, e-Services and e-Society, pages 51–64. Springer, 2009.

[75] Verna Allee. *A value network approach for modeling and measuring intangibles*. Transparent Enterprise, Madrid. Available at http://www. vernaallee. c om, 2002.

[76] Jaap Gordijn, Hans Akkermans, and Hans Van Vliet. *Business Modelling is not Process Modelling*. In Conceptual Modeling for E-Business and the Web, ECOMO 2000, volume 1921 of *LNCS*, pages 40–51. Springer, 2000.

[77] Stephen E. Toulmin. *The Uses of Argument*. Cambridge University Press, 1958.

[78] S.E. Toulmin, R.D. Rieke, and A. Janik. *An Introduction to Reasoning*. Macmillan, 1979.

[79] Jintae Lee and Kum-Yew Lai. *What's in design rationale?* Human–Computer Interaction, 6(3-4):251–280, 1991.

[80] Allan Maclean, Richard M. Young, and Thomas P. Moran. *Design Rationale: The Argument behind the Artefact*. In Proceedings of the Computer Human Interaction conference (CHI), 1989.

[81] Simon J. Buckingham Shum, Allan MacLean, Victoria M. E. Bellotti, and Nick V. Hammond. *Graphical Argumentation and Design Cognition*. 12(3):267–300, 1997.

[82] John Mylopoulos, Alex Borgida, Matthias Jarke, and Manolis Koubarakis. *Telos: Representing knowledge about information systems*. ACM Transactions on Information Systems (TOIS), 8(4):325–362, 1990.

[83] Gerhard Fischer, Andreas C. Lemke, Raymond McCall, and Anders I. Morch. *Making Argumentation Serve Design*. Hum.-Comput. Interact., 6(3):393–419, September 1991.

[84] Lukasz Cyra and Janusz Górski. *Support for Argument Structures Review and Assessment*. Reliability Engineering & System Safety, 96(1):26–37, 2011. Special Issue on Safecomp 2008.

[85] *Defence Standard 00-56 Issue 4 (part 1): Safety Management Requirements for Defence Systems*, July 2007.

[86] Timothy Patrick Kelly. *Arguing Safety: A Systematic Approach to Managing Safety Cases*. University of York, 1999.

[87] Luke Emmet. *Using Claims, Arguments and Evidence: A Pragmatic View–and tool support in ASCE*. www.adelard.com.

[88] Tim Kelly and Rob Weaver. *The Goal Structuring Notation – a Safety Argument Notation*. In Proeedings of Dependable Systems and Networks 2004 Workshop on Assurance Cases, 2004.

[89] *Adelard Safety Case Development (ASCAD) Manual*. London, UK, 2010.

[90] Kimberly M Markham, Joel J Mintzes, and M Gail Jones. *The concept map as a research and evaluation tool: Further evidence of validity*. Journal of research in science teaching, 31(1):91–101, 1994.

[91] Joeran Beel and Stefan Langer. *An exploratory analysis of mind maps*. In Proceedings of the 11th ACM symposium on Document engineering, pages 81–84. ACM, 2011.

[92] Martin J Eppler. *A comparison between concept maps, mind maps, conceptual diagrams, and visual metaphors as complementary tools for knowledge construction and sharing*. Information visualization, 5(3):202–210, 2006.

[93] J. Górski, A. Jarzębowicz, R. Leszczyna, J. Miler, and M. Olszewski. *Trust Case: Justifying Trust in an It Solution*. Reliability Engineering and System Safety, 89(1):33–47, 2005.

[94] Guttorm Sindre and Andreas L. Opdahl. *Eliciting Security Requirements with Misuse Cases*. Requirements Engineering, 10(1):34–44, 2005.

[95] John Rushby. *The interpretation and evaluation of assurance cases*. SRI International, Menlo Park, CA, USA, 2015.

[96] Charles B Haley, Jonathan D Moffett, Robin Laney, and Bashar Nuseibeh. *Arguing Security: Validating Security Requirements Using Structured Argumentation*. In Proceedings of Third Symposium on Requirements Engineering for Information Security (SREIS'05) held in conjunction with the 13th International Requirements Engineering Conference (RE'05), 2005.

[97] Jeff Rowe, Karl Levitt, Simon Parsons, Elizabeth Sklar, Andrew Applebaum, and Sharmin Jalal. *Argumentation Logic to Assist in Security Administration*. In Proceedings of the 2012 New Security Paradigms Workshop, NSPW '12, pages 43–52, New York, NY, USA, 2012. ACM.

[98] Charles Haley, Robin Laney, Jonathan Moffett, and Bashar Nuseibeh. *Security Requirements Engineering: A Framework for Representation and Analysis*. IEEE Trans. Softw. Eng., 34(1):133–153, January 2008.

[99] Virginia N. L. Franqueira, Thein Than Tun, Yijun Yu, Roel Wieringa, and Bashar Nuseibeh. *Risk and Argument: A Risk-based Argumentation Method for Practical Security.* In RE, pages 239–248. IEEE, 2011.

[100] Yijun Yu, Virginia N. L. Franqueira, Thein Than Tun, Roel Wieringa, and Bashar Nuseibeh. *Automated Analysis of Security Requirements through Risk-based Argumentation*. Journal of Systems and Software, 106:102–116, 2015.

[101] Henry Prakken. *An Abstract Framework for Argumentation with Structured Arguments*. Argument & Computation, 1:93–124, 2010.

[102] L. Aversano, G. Canfora, A. De Lucia, and S. Stefanucci. *Understanding SQL through iconic interfaces*. In Computer Software and Applications Conference, 2002. COMPSAC 2002. Proceedings. 26th Annual International, pages 703–708, 2002.

[103] Saskia Bakker, Debby Vorstenbosch, Elise van den Hoven, Gerard Hollemans, and Tom Bergman. *Tangible Interaction in Tabletop Games: Studying Iconic and Symbolic Play Pieces*. In Proceedings of the International Conference on Advances in Computer Entertainment Technology, ACE '07, pages 163–170, New York, NY, USA, 2007. ACM.

[104] William K. Horton. *The ICON Book: Visual Symbols for Computer Systems and Documentation*. John Wiley & Sons, Inc., New York, NY, USA, 1st edition, 1994.

[105] A. Grosskopf, J. Edelman, and M. Weske. *Tangible Business Process Modeling - Methodology and Experiment Design*. In 1st International Workshop on Empirical Research in Business Process Management (ER-BPM'09), pages 53–64, Ulm, Germany, 2009. Springer.

[106] C. Heath, L. Coles-Kemp, and P. Hall. *Logical Lego? Co-constructed perspectives on service design*. In Proceedings of NordDesign 2014, page 416. Aalto Design Factory, 2014.

[107] Marc Rettig. *Prototyping for Tiny Fingers*. Commun. ACM, 37(4):21–27, April 1994.

[108] Albert Fleischmann, Werner Schmidt, and Christian Stary. *Tangible or Not Tangible – A Comparative Study of Interaction Types for Process Modeling Support*. In Proceedings of the 16th International Conference on Human-Computer Interaction (HCI), Part II: Advanced Interaction Modalities and Techniques, pages 544–555. Springer, 2014.

[109] Neil J Smelser, Paul B Baltes, et al. *Analytical induction: methodology*. In International encyclopedia of the social & behavioral sciences, volume 11. Elsevier Amsterdam, 2001.

[110] R.J. Wieringa. *Design Science Methodology for Information Systems and Software Engineering*. Springer, 2014.

[111] Thomas Lloyd Short. *Peirce's theory of signs*. Cambridge University Press, 2007.

[112] Daniel Moody. *The "Physics" of Notations: Toward a Scientific Basis for Constructing Visual Notations in Software Engineering*. IEEE Transactions on Software Engineering, 35(6):756–779, November 2009.

[113] Adrian Frutiger. *Signs and symbols: their design and meaning*. Van Nostrand Reinhold Company, 1989.

[114] G. W. Fitzmaurice, H. Ishii, and William A. S. Buxton. *Bricks: Laying the Foundations for Graspable User Interfaces*. In Proceedings of the SIGCHI Conference on Human Factors in Computing Systems, CHI '95, pages 442–449, New York, NY, USA, 1995. ACM Press/Addison-Wesley Publishing Co.

[115] O. Zuckerman and Ayelet Gal-Oz. *To TUI or not to TUI: Evaluating performance and preference in tangible vs. graphical user interfaces*. International Journal of Human-Computer Studies, 71(7–8):803 – 820, 2013.

[116] John Sweller. *Cognitive Load During Problem Solving: Effects on Learning*. Cognitive Science, 12(2):257–285, 1988.

[117] G. Miller. *The Magical Number Seven, Plus or Minus Two: Some Limits on Our Capacity for Processing Information*. The Psychological Review, 63:81–97, 1956.

[118] David Hecht, Miriam Reiner, and Avi Karni. *Enhancement of response times to bi- and tri-modal sensory stimuli during active movements*. Experimental Brain Research, 185(4):655–665, 2008.

[119] I. Vessey and D. Galletta. *Cognitive Fit: An Empirical Study of Information Acquisition*. Information Systems Research, 2(1):63–84, 1991.

[120] Marija Bjeković, Henderik A Proper, and Jean-Sébastien Sottet. *Embracing pragmatics*. In International Conference on Conceptual Modeling, pages 431–444. Springer, 2014.

[121] Ilona Wilmont, Sytse Hengeveld, Erik Barendsen, and Stijn Hoppenbrouwers. *Cognitive Mechanisms of Conceptual Modelling*, pages 74–87. Springer Berlin Heidelberg, Berlin, Heidelberg, 2013.

[122] M.J. Kim and M.L. Maher. *The impact of tangible user interfaces on designers' spatial cognition*. Design Studies, 29:222—253, May 2008.

[123] Michael S. Horn, Erin Treacy Solovey, R. Jordan Crouser, and Robert J.K. Jacob. *Comparing the Use of Tangible and Graphical Programming Languages for Informal Science Education*. In Proceedings of the SIGCHI Conference on Human Factors in Computing Systems, CHI '09, pages 975–984, New York, NY, USA, 2009. ACM.

[124] Vikram Parmar, Gert Groeneveld, Ashis Jalote-Parmar, and David Keyson. *Tangible User Interface for Increasing Social Interaction Among Rural Women*. In Proceedings of the 3rd International Conference on Tangible and Embedded Interaction, TEI '09, pages 139–145, New York, NY, USA, 2009. ACM.

[125] Jonathan Edelman, Alexander Grosskopf, Mathias Weske, and L. Leifer. *Tangible business process modeling: a new approach*. In Proceedings of the 17th international conference on engineering design, ICED, volume 9, pages 153–168, August 2009.

[126] Alexander Lübbe-Grosskopf. *Tangible Business Process Modeling: Design and Evaluation of a Process Model Elicitation Technique*. PhD thesis, Universität Potsdam, 2011.

[127] The TREsPASS Project, D1.3.1. *Initial prototype of the socio-technical security model*, 2013. Deliverable D1.3.1.

[128] J. Nielsen. *Usability 101: Introduction to Usability*. Jakob Nielsen's Alertbox, 2003.

[129] J. R. Lewis. *Psychometric Evaluation of an After-scenario Questionnaire for Computer Usability Studies: The ASQ*. SIGCHI Bull., 23(1):78–81, January 1991.

[130] E. Vriezekolk, S. Etalle, and R. Wieringa. *Experimental Validation of a Risk Assessment Method*. In 21st International Working Conference on Requirements Engineering: Foundations for Software Quality (REFSQ 15). Springer, 2015.

[131] J.A. Garde and M.C. van der Voort. *The Procedure Usability Game: A Participatory Game for Development of Complex Medical Procedures & Products*. In Proceedings of the CIRP IPS2 Conference 2009, 2009.

[132] J. Barjis. *Collaborative, Participative and Interactive Enterprise Modeling*. In Enterprise Information Systems, volume 24 of *Lecture Notes in Business Information Processing*, pages 651–662. Springer, 2009.

[133] F3-Consortium. *F3 Reference Manual*. Technical Report MSU-CSE-00-2, SISU, Stockholm, 1994. ESPRIT III Project 6612.

[134] Doris Weitlaner, Annemarie Guettinger, and Markus Kohlbacher. *Intuitive Comprehensibility of Process Models*. In Proceedings of the 5th International Conference on Running Processes (S-BPM ONE 2013). Springer, 2013.

[135] O. I. Lindland, G. Sindre, and A. Solvberg. *Understanding quality in conceptual modeling*. IEEE Software, 11(2):42–49, March 1994.

[136] Bhuvan Unhelkar. *The Quality Strategy for UML*, pages 1–26. John Wiley & Sons, Inc., 2005.

[137] P Robinson. *Task complexity, task difficulty, and task production: exploring interactions in a componential framework*. Applied Linguistics, 22(1):27–57, 2001.

[138] Michiel Renger, Gwendolyn L. Kolfschoten, and Gert-Jan De Vreede. *Challenges in collaborative modelling: a literature review and research agenda*. International Journal of Simulation and Process Modelling, 4(3-4):248–263, 2008.

[139] Agnès Front, Dominique Rieu, and Marco Oswaldo Santorum. *A participative end-user modeling approach for business process requirements*. In Proceedings of the 15th Working Conference on Business Process Modeling, Development and Support). Springer, 2014.

[140] Alexandr Vasenev, Lorena Montoya, and Andrea Ceccarelli. *A Hazus-based method for assessing robustness of electricity supply to critical smart grid consumers during flood events*. In International Conference on Availability, Reliability and Security (ARES), page 6. IEEE, 2016.

[141] Kasper Hornbæk. *Current Practice in Measuring Usability: Challenges to Usability Studies and Research*. Int. J. Hum.-Comput. Stud., 64(2):79–102, February 2006.

[142] Brigid Barron. *When smart groups fail*. The journal of the learning sciences, 12(3):307–359, 2003.

[143] Robert J McQueen, Karen Rayner, and Ned Kock. *Contribution by participants in face-to-face business meetings: Implications for collaborative technology*. Journal of Systems and Information Technology, 3(1):15–34, 1999.

[144] P. Bordia. *Face-to-Face Versus Computer-Mediated CommunlGation:A Synthesis of the Experimental Literature*. The Journal of Business Communication, 34:99–120, 1997.

[145] H. J. M. Tabachneck-Schijf, J. H. Verpoorten, R. L. W. van de Weg, and R. J. Wieringa. *The influence of conceptual user models on the creation and interpretation of diagrams representing reactive systems*. In Current Research in Information Sciences and Technologies. Multidisciplinary approaches to global information systems, Spain, pages 452–456, Badajoz, Spain, 2006. Open Institute of Knowledge.

[146] L.W. Barasalou, W.K. Simmons, A.K. Barbey, and C.D. Wilson. *Grounding conceptual knowledge in modality-specific systems*. Trends in Cognitive Science, 7(2):84–91, February 2003.

[147] V. van der Velde. *Communication, concepts and grounding*. Neural Networks, 62:112–117, 2015.

[148] Stijn Hoppenbrouwers and Ilona Wilmont. *Focused conceptualisation: framing questioning and answering in model-oriented dialogue games*. In IFIP Working Conference on The Practice of Enterprise Modeling, pages 190–204. Springer, 2010.

[149] W.S. Robinson. *The logical structure of analytic induction*. American Sociological Review, 16(6):812–818, December 1951.

[150] F. Znaniecki. *The Method of Sociology*. Octagon Books, 1934.

[151] Joshua Gold. *Data breaches and computer hacking: liability & insurance issues*. American Bar Association's Government Law Committee Newsletter, Fall, 2011.

[152] Travis D Breaux and David L Baumer. *Legally "reasonable" security requirements: A 10-year FTC retrospective*. Computers & Security, 30(4):178–193, 2011.

[153] European Parliament & Council. *Regulation (EU) 2016/679 of the European Parliament and of the Council of 27 April 2016 on the protection of natural persons with regard to the processing of personal data and on the free movement of such data, and repealing Directive 95/46/EC (General Data Protection Regulation)*. Official Journal of the European Union, L119/59:1–88, May 2016.

[154] Donald G Firesmith. *Analyzing and specifying reusable security requirements*. Technical report, DTIC Document, 2003.

[155] Yijun Yu, Thein Than Tun, Alessandra Tedeschi, Virginia N. L. Franqueira, and Bashar Nuseibeh. *OpenArgue: Supporting argumentation to evolve secure software systems.* In RE, pages 351–352. IEEE, August 2011.

[156] Phan Minh Dung. *On the Acceptability of Arguments and Its Fundamental Role in Nonmonotonic Reasoning, Logic Programming and N-person Games*. Artificial Intelligence, 77:321–357, 1995.

[157] ENISA. *ENISA: Cloud Computing Security Risk Assessment*. Technical report, ENISA, May 2009.

[158] Cloud Security Alliance. *Top Threats to Cloud Computing V1.0*. Technical report, Cloud Security Alliance, March 2010.

[159] A.H. Dutoit, R. McCall, I. Mistrik, and B. Paech. *Rationale Management in Software Engineering*. Springer, 2007.

[160] E. Lumsdaine and M. Lumsdaine. *Creative Problem Solving*. IEEE Potentials, 13(5):4–9, December 1994.

[161] T. Raz and E. Michael. *Use and Benefits of Tools for Project Risk Management*. International Journal of Project Management, 19(1):9–17, 2001.

[162] Terry Lyons and Martin Skitmore. *Project Risk Management in the Queensland Engineering Construction Industry: A Survey*. International Journal of Project Management, 22(1):51–61, 2004.

[163] C. Solis and N. Ali. *Distributed Requirements Elicitation Using a Spatial Hypertext Wiki*. In 2010 5th IEEE International Conference on Global Software Engineering, pages 237–246, August 2010.

[164] Betty H. C. Cheng and Joanne M. Atlee. *Research Directions in Requirements Engineering*. In 2007 Future of Software Engineering, FOSE '07, pages 285–303, Washington, DC, USA, 2007. IEEE Computer Society.

[165] Norbert Seyff, Paul Grunbacher, Neil Maiden, and Amit Tosar. *Requirements Engineering Tools Go Mobile*. In Proceedings of the 26th International Conference on Software Engineering, ICSE '04, pages 713–714, Washington, DC, USA, 2004. IEEE Computer Society.

[166] K. Mohan and B. Ramesh. *Ontology-based support for variability management in product and service families.* In Proceedings of the 36th Hawaii International Conference on System Sciences, Hawaii, pages 9–18, Jan 2003.

[167] Des Freedman. *The Phone Hacking Scandal: Implications for Regulation*. Television & New Media, 13(1):17–20, 2012.

[168] Jaap Gordijn and Hans Akkermans. *Value based requirements engineering: Exploring innovative e-commerce idea*. Requirements Engineering Journal, 8(2):114–134, 2003.

[169] Ross Anderson. *Why information security is hard-an economic perspective*. In Computer security applications conference, 2001. acsac 2001. proceedings 17th annual, pages 358–365. IEEE, 2001.

[170] R. Normann and R. Ramírez. *Designing Interactive Strategy - From Value Chain to Value Constellation*. John Wiley & Sons Inc., Chichester, UK, 1994.

[171] The TREsPASS Project, D7.3.1. *Results from Case Study B*, 2014. Deliverable D7.3.1.

[172] Vera Kartseva, Jaap Gordijn, and Yao-Hua Tan. *Toward a modeling tool for designing control mechanisms for network organizations*. International Journal of Electronic Commerce, 10(2):58–84, 2005.

[173] Vera Kartseva, Jaap Gordijn, and Yao-Hua Tan. *Designing Value-Based Inter-organizational Controls Using Patterns*, volume 14 of *LNBIP*, pages 276–301. Springer Berlin Heidelberg, 2009.

[174] Vera Kartseva. *Designing Controls for Network Organization: A Value-Based Approach*. PhD thesis, Vrije Universiteit Amsterdam, 2008.

[175] Dan Baker. *International Revenue Share Fraud: Are We Winning the Battle Against Telecom Pirates?* Black Swan Telecom Journal, November 2012.

[176] Wikipedia. *Business Telephone System — Wikipedia, The Free Encyclopedia*, 2014.

[177] SMARTVOX. *How secure is your Asterisk PBX?* `http://kb.smartvox.co.uk/asterisk/secure-asterisk-pbx-part-1/` accessed Nov 2014, 2014.

[178] D. Richard. Kuhn, National Institute of Standards, and Technology (U.S.). *PBX vulnerability analysis [microform] : finding holes in your PBX before someone else does / D. Richard Kuhn*. U.S. Dept. of Commerce, Technology Administration, National Institute of Standards and Technology ; For sale by the Supt. of Docs., U.S. G.P.O Gaithersburg, Md. : [Washington, D.C, 2001.

[179] Terry Regan. *PBX Security in the VoIP environment*. `http://www.spitfire.co.uk/pdf/05_PBX_Security_in_the_VoIP_environment-white_paper_140313_2.pdf` accessed Nov 2014, March 2013.

[180] Prince Mayurank Singh. *Integrating business value in enterprise architecture modeling and analysis*, August 2013.

[181] Wil Janssen, Rene van Buuren, and Jaap Gordijn. *Business Case Modelling for E-Services*. In D. R. Vogel, P. Walden, J. Gricar, and G. Lenart, editors, Proceedings of the 18th BLED conference (e-Integration in Action), pages cdrom,, Maribor, SL, 2005. University of Maribor.

[182] Jaap Gordijn and Hans Van Vliet. *On the Interaction between Business Models and Software Architecture in Electronic Commerce*. In Addendum to the proceedings of the 7th European Software Engineering Conference/Foundations of Software Engineering / ESEC 1999, 1999.

[183] Roel Wieringa and Jaap Gordijn. *Value-oriented design of correct service coordination protocols*. In Proceedings of the 20th ACM Symposium on Applied Computing, pages 1320–1327. ACM Press, 2005.

[184] Vincent Pijpers and Jaap Gordijn. *Bridging business value models and process models in aviation value webs via possession rights*. In 40th Annual Hawaii International Conference on System Sciences, HICSS 2007. IEEE, 2007.

[185] Jaap Gordijn and Roel Wieringa. *A value-oriented approach to e-business process design*. In Proceedings of the 15th International Conference, CAiSE 2003, volume 2681 of *LNCS*, pages 390–403. Springer Verlag, 2003.

[186] Bradley Reaves, Ethan Shernan, Adam Bates, Henry Carter, and Patrick Traynor. *Boxed Out: Blocking Cellular Interconnect Bypass Fraud at the Network Edge*. In 24th USENIX Security Symposium (USENIX Security 15), pages 833–848, Washington, D.C., August 2015. USENIX Association.

[187] Javier Carbo, Jesus Garcia, and JoseM. Molina. *Trust and Reputation in E-services: Concepts, Models and Applications*. In E-Service Intelligence, volume 37 of *Studies in Computational Intelligence*, pages 327–345. Springer Berlin Heidelberg, 2007.

[188] Yao-Hua Tan, Wout Hofman, Jaap Gordijn, and Joris Hulstijn. *A Framework for the Design of Service Systems*. In Service Systems Implementation, Service Science: Research and Innovations in the Service Economy, pages 51–74. Springer US, 2011.

[189] Inam Soomro and Naved Ahmed. *Towards Security Risk-Oriented Misuse Cases*. In Business Process Management Workshops, volume 132 of *Lecture Notes in Business Information Processing*, pages 689–700. Springer Berlin Heidelberg, 2013.

[190] Eric S. K. Yu. *Models for Supporting the Redesign of Organizational Work*. In Proceedings of Conference on Organizational Computing Systems, COCS '95, pages 226–236, New York, NY, USA, 1995. ACM.

[191] MichaelH. Cahill, Diane Lambert, JoséC. Pinheiro, and DonX. Sun. *Detecting Fraud in the Real World*. In Handbook of Massive Data Sets, volume 4 of *Massive Computing*, pages 911–929. Springer US, 2002.

[192] Markus Ruch and Stefan Sackmann. *Customer-Specific Transaction Risk Management in E-Commerce*. In Value Creation in E-Business Management, volume 36 of *Lecture Notes in Business Information Processing*, pages 68–79. Springer Berlin Heidelberg, 2009.

[193] Dritsoula L. and Musacchio J. *A game of clicks: Economic incentives to fight click fraud in ad networks*. In Performance Evaluation Review, volume 41, pages 12–15, 2014.

[194] Wolter Pieters, SE Banescu, and Simona Posea. *System Abuse by Service Composition: Analysis and Prevention*. In CESUN 2012: 3rd International Engineering Systems Symposium, Delft University of Technology, The Netherlands, 18-20 June 2012, 2012.

[195] David J. Hand Richard J. Bolton. *Statistical Fraud Detection: A Review*. Statistical Science, 17(3):235–249, 2002.

[196] Sean Ross. *How does revenue sharing work in practice?* Investopedia, 2015. [Online; accessed 12-December-2015].

[197] W. M. P. van der Aalst. *Discovering Coordination Patterns using Process Mining*. In L Bocchi and P. Ciancarini, editors, Proceedings of the First International Workshop on Petri Nets and Coordination, PNC04, pages 49–64, CNR PISA, Italy, 2004. STAR (Servizio Tipografico Area della Ricerca).

[198] Hassan Fatemi, Marten van Sinderen, and Roel Wieringa. *E3value to BPMN Model Transformation*, pages 333–340. Springer Berlin Heidelberg, Berlin, Heidelberg, 2011.

[199] Rainer Schuster, Thomas Motal, Christian Huemer, and Hannes Werthner. *From Economic Drivers to B2B Process Models: A Mapping from REA to UMM*. In Witold Abramowicz and Robert Tolksdorf, editors, Business Information Systems, 13th International Conference, BIS 2010, Berlin, Germany, May 3-5, 2010. Proceedings, volume 47 of *Lecture Notes in Business Information Processing*, pages 119–131. Springer, 2010.

[200] Rainer Schuster and Thomas Motal. *From e3-value to REA: Modeling Multi-party E-business Collaborations*. In Proceedings of the 2009 IEEE Conference on Commerce and Enterprise Computing, CEC '09, pages 202–208, Washington, DC, USA, 2009. IEEE Computer Society.

[201] H. Weigand, P. Johannesson, B. Andersson, M. Bergholtz, A. Edirisuriya, and T. Ilayperuma. *Value modeling and the transformation from value model to process mode*. In G. Doumeingts, J. Muller, G. Morel, and B. Vallespir, editors, Enterprise Interoperability: New Challenges and Approaches, pages 1–10, London, UK, 2007. Springer-Verlag.

[202] L. Bodenstaff. *Managing Dependency Relations in Inter-Organizational Models*. PhD thesis, University of Twente, June 2010.

[203] Marta Indulska, Jan C. Recker, Michael Rosemann, and Peter Green. *Business process modeling : current issues and future challenges*. Lecture Notes in Computer Science, 5565:501–514, June 2009.

[204] Jaap Gordijn. *E-business value modelling using the e3-value ontology*, chapter 5, pages 98–127. Elsevies Butterworth-Heinemann, Oxford, UK, 2004. Preprint available.

[205] S. Onoda, Y. Ikkai, T. Kobayashi, and N. Komoda. *Definition of deadlock patterns for business processes workflow models*. In HICSS 1999: Proceedings of the Thirty-second Annual Hawaii International Conference on System Sciencess, Washington, DC, USA, volume 5, pages 55–65, Jan 1999.

[206] Ahmed Awad and Frank Puhlmann. *Business Information Systems: 11th International Conference, BIS 2008, Innsbruck, Austria, May 5-7, 2008. Proceedings*, chapter Structural Detection of Deadlocks in Business Process Models, pages 239–250. Springer Berlin Heidelberg, 2008.

[207] Xiang Fu, Tevfik Bultan, and Jianwen Su. *Conversation Protocols: A Formalism for Specification and Verification of Reactive Electronic Services*, pages 188–200. Springer Berlin Heidelberg, 2003.

[208] Gero Decker and Mathias Weske. *Local Enforceability in Interaction Petri Nets*. In Gustavo Alonso, Peter Dadam, and Michael Rosemann, editors, Business Process Management: 5th International Conference, BPM 2007, Brisbane, Australia, September 24-28, 2007. Proceedings, pages 305–319. Springer Berlin Heidelberg, 2007.

[209] Gero Decker, Alistair Barros, Frank Michael Kraft, and Niels Lohmann. *Non-desynchronizable Service Choreographies*. In Athman Bouguettaya, Ingolf Krueger, and Tiziana Margaria, editors, Service-Oriented Computing – ICSOC 2008: 6th International Conference, Sydney, Australia, December 1-5, 2008. Proceedings, pages 331–346. Springer Berlin Heidelberg, 2008.

[210] John Scourias. *Overview of the Global System for Mobile Communications*. Technical report, Simon Fraser University, 1995.

[211] Simon Buckingham Shum. *The Roots of Computer Supported Argument Visualization*, pages 3–24. Springer London, London, 2003.

Other titles in the SIKS dissertation series since 2009

2011 01 Botond Cseke (RUN), Variational Algorithms for Bayesian Inference in Latent Gaussian Models
02 Nick Tinnemeier (UU), Organizing Agent Organizations. Syntax and Operational Semantics of an Organization-Oriented Programming Language
03 Jan Martijn van der Werf (TUE), Compositional Design and Verification of Component-Based Information Systems
04 Hado van Hasselt (UU), Insights in Reinforcement Learning; Formal analysis and empirical evaluation of temporal-difference
05 Bas van der Raadt (VU), Enterprise Architecture Coming of Age - Increasing the Performance of an Emerging Discipline.
06 Yiwen Wang (TUE), Semantically-Enhanced Recommendations in Cultural Heritage
07 Yujia Cao (UT), Multimodal Information Presentation for High Load Human Computer Interaction
08 Nieske Vergunst (UU), BDI-based Generation of Robust Task-Oriented Dialogues
09 Tim de Jong (OU), Contextualised Mobile Media for Learning
10 Bart Bogaert (UvT), Cloud Content Contention
11 Dhaval Vyas (UT), Designing for Awareness: An Experience-focused HCI Perspective
12 Carmen Bratosin (TUE), Grid Architecture for Distributed Process Mining
13 Xiaoyu Mao (UvT), Airport under Control. Multiagent Scheduling for Airport Ground Handling
14 Milan Lovric (EUR), Behavioral Finance and Agent-Based Artificial Markets
15 Marijn Koolen (UvA), The Meaning of Structure: the Value of Link Evidence for Information Retrieval
16 Maarten Schadd (UM), Selective Search in Games of Different Complexity
17 Jiyin He (UVA), Exploring Topic Structure: Coherence, Diversity and Relatedness
18 Mark Ponsen (UM), Strategic Decision-Making in complex games
19 Ellen Rusman (OU), The Mind's Eye on Personal Profiles
20 Qing Gu (VU), Guiding service-oriented software engineering - A view-based approach
21 Linda Terlouw (TUD), Modularization and Specification of Service-Oriented Systems
22 Junte Zhang (UVA), System Evaluation of Archival Description and Access
23 Wouter Weerkamp (UVA), Finding People and their Utterances in Social Media
24 Herwin van Welbergen (UT), Behavior Generation for Interpersonal Coordination with Virtual Humans On Specifying, Scheduling and Realizing Multimodal Virtual Human Behavior
25 Syed Waqar ul Qounain Jaffry (VU), Analysis and Validation of Models for Trust Dynamics
26 Matthijs Aart Pontier (VU), Virtual Agents for Human Communication - Emotion Regulation and Involvement-Distance Trade-Offs in Embodied Conversational Agents and Robots
27 Aniel Bhulai (VU), Dynamic website optimization through autonomous management of design patterns
28 Rianne Kaptein (UVA), Effective Focused Retrieval by Exploiting Query Context and Document Structure
29 Faisal Kamiran (TUE), Discrimination-aware Classification
30 Egon van den Broek (UT), Affective Signal Processing (ASP): Unraveling the mystery of emotions
31 Ludo Waltman (EUR), Computational and Game-Theoretic Approaches for Modeling Bounded Rationality
32 Nees-Jan van Eck (EUR), Methodological Advances in Bibliometric Mapping of Science
33 Tom van der Weide (UU), Arguing to Motivate Decisions
34 Paolo Turrini (UU), Strategic Reasoning in Interdependence: Logical and Game-theoretical Investigations

35 Maaike Harbers (UU), Explaining Agent Behavior in Virtual Training
36 Erik van der Spek (UU), Experiments in serious game design: a cognitive approach
37 Adriana Burlutiu (RUN), Machine Learning for Pairwise Data, Applications for Preference Learning and Supervised Network Inference
38 Nyree Lemmens (UM), Bee-inspired Distributed Optimization
39 Joost Westra (UU), Organizing Adaptation using Agents in Serious Games
40 Viktor Clerc (VU), Architectural Knowledge Management in Global Software Development
41 Luan Ibraimi (UT), Cryptographically Enforced Distributed Data Access Control
42 Michal Sindlar (UU), Explaining Behavior through Mental State Attribution
43 Henk van der Schuur (UU), Process Improvement through Software Operation Knowledge
44 Boris Reuderink (UT), Robust Brain-Computer Interfaces
45 Herman Stehouwer (UvT), Statistical Language Models for Alternative Sequence Selection
46 Beibei Hu (TUD), Towards Contextualized Information Delivery: A Rule-based Architecture for the Domain of Mobile Police Work
47 Azizi Bin Ab Aziz (VU), Exploring Computational Models for Intelligent Support of Persons with Depression
48 Mark Ter Maat (UT), Response Selection and Turn-taking for a Sensitive Artificial Listening Agent
49 Andreea Niculescu (UT), Conversational interfaces for task-oriented spoken dialogues: design aspects influencing interaction quality

2012 01 Terry Kakeeto (UvT), Relationship Marketing for SMEs in Uganda
02 Muhammad Umair (VU), Adaptivity, emotion, and Rationality in Human and Ambient Agent Models
03 Adam Vanya (VU), Supporting Architecture Evolution by Mining Software Repositories
04 Jurriaan Souer (UU), Development of Content Management System-based Web Applications
05 Marijn Plomp (UU), Maturing Interorganisational Information Systems
06 Wolfgang Reinhardt (OU), Awareness Support for Knowledge Workers in Research Networks
07 Rianne van Lambalgen (VU), When the Going Gets Tough: Exploring Agent-based Models of Human Performance under Demanding Conditions
08 Gerben de Vries (UVA), Kernel Methods for Vessel Trajectories
09 Ricardo Neisse (UT), Trust and Privacy Management Support for Context-Aware Service Platforms
10 David Smits (TUE), Towards a Generic Distributed Adaptive Hypermedia Environment
11 J.C.B. Rantham Prabhakara (TUE), Process Mining in the Large: Preprocessing, Discovery, and Diagnostics
12 Kees van der Sluijs (TUE), Model Driven Design and Data Integration in Semantic Web Information Systems
13 Suleman Shahid (UvT), Fun and Face: Exploring non-verbal expressions of emotion during playful interactions
14 Evgeny Knutov (TUE), Generic Adaptation Framework for Unifying Adaptive Web-based Systems
15 Natalie van der Wal (VU), Social Agents. Agent-Based Modelling of Integrated Internal and Social Dynamics of Cognitive and Affective Processes.
16 Fiemke Both (VU), Helping people by understanding them - Ambient Agents supporting task execution and depression treatment
17 Amal Elgammal (UvT), Towards a Comprehensive Framework for Business Process Compliance
18 Eltjo Poort (VU), Improving Solution Architecting Practices
19 Helen Schonenberg (TUE), What's Next? Operational Support for Business Process Execution
20 Ali Bahramisharif (RUN), Covert Visual Spatial Attention, a Robust Paradigm for Brain-Computer Interfacing
21 Roberto Cornacchia (TUD), Querying Sparse Matrices for Information Retrieval
22 Thijs Vis (UvT), Intelligence, politie en veiligheidsdienst: verenigbare grootheden?
23 Christian Muehl (UT), Toward Affective Brain-Computer Interfaces: Exploring the Neurophysiology of Affect during Human Media Interaction
24 Laurens van der Werff (UT), Evaluation of Noisy Transcripts for Spoken Document Retrieval
25 Silja Eckartz (UT), Managing the Business Case Development in Inter-Organizational IT Projects: A Methodology and its Application
26 Emile de Maat (UVA), Making Sense of Legal Text
27 Hayrettin Gurkok (UT), Mind the Sheep! User Experience Evaluation & Brain-Computer Interface Games

28 Nancy Pascall (UvT), Engendering Technology Empowering Women
29 Almer Tigelaar (UT), Peer-to-Peer Information Retrieval
30 Alina Pommeranz (TUD), Designing Human-Centered Systems for Reflective Decision Making
31 Emily Bagarukayo (RUN), A Learning by Construction Approach for Higher Order Cognitive Skills Improvement, Building Capacity and Infrastructure
32 Wietske Visser (TUD), Qualitative multi-criteria preference representation and reasoning
33 Rory Sie (OUN), Coalitions in Cooperation Networks (COCOON)
34 Pavol Jancura (RUN), Evolutionary analysis in PPI networks and applications
35 Evert Haasdijk (VU), Never Too Old To Learn – On-line Evolution of Controllers in Swarm- and Modular Robotics
36 Denis Ssebugwawo (RUN), Analysis and Evaluation of Collaborative Modeling Processes
37 Agnes Nakakawa (RUN), A Collaboration Process for Enterprise Architecture Creation
38 Selmar Smit (VU), Parameter Tuning and Scientific Testing in Evolutionary Algorithms
39 Hassan Fatemi (UT), Risk-aware design of value and coordination networks
40 Agus Gunawan (UvT), Information Access for SMEs in Indonesia
41 Sebastian Kelle (OU), Game Design Patterns for Learning
42 Dominique Verpoorten (OU), Reflection Amplifiers in self-regulated Learning
43 Withdrawn
44 Anna Tordai (VU), On Combining Alignment Techniques
45 Benedikt Kratz (UvT), A Model and Language for Business-aware Transactions
46 Simon Carter (UVA), Exploration and Exploitation of Multilingual Data for Statistical Machine Translation
47 Manos Tsagkias (UVA), Mining Social Media: Tracking Content and Predicting Behavior
48 Jorn Bakker (TUE), Handling Abrupt Changes in Evolving Time-series Data
49 Michael Kaisers (UM), Learning against Learning - Evolutionary dynamics of reinforcement learning algorithms in strategic interactions
50 Steven van Kervel (TUD), Ontologogy driven Enterprise Information Systems Engineering
51 Jeroen de Jong (TUD), Heuristics in Dynamic Sceduling; a practical framework with a case study in elevator dispatching

2013 01 Viorel Milea (EUR), News Analytics for Financial Decision Support
02 Erietta Liarou (CWI), MonetDB/DataCell: Leveraging the Column-store Database Technology for Efficient and Scalable Stream Processing
03 Szymon Klarman (VU), Reasoning with Contexts in Description Logics
04 Chetan Yadati (TUD), Coordinating autonomous planning and scheduling
05 Dulce Pumareja (UT), Groupware Requirements Evolutions Patterns
06 Romulo Goncalves (CWI), The Data Cyclotron: Juggling Data and Queries for a Data Warehouse Audience
07 Giel van Lankveld (UvT), Quantifying Individual Player Differences
08 Robbert-Jan Merk (VU), Making enemies: cognitive modeling for opponent agents in fighter pilot simulators
09 Fabio Gori (RUN), Metagenomic Data Analysis: Computational Methods and Applications
10 Jeewanie Jayasinghe Arachchige (UvT), A Unified Modeling Framework for Service Design.
11 Evangelos Pournaras (TUD), Multi-level Reconfigurable Self-organization in Overlay Services
12 Marian Razavian (VU), Knowledge-driven Migration to Services
13 Mohammad Safiri (UT), Service Tailoring: User-centric creation of integrated IT-based homecare services to support independent living of elderly
14 Jafar Tanha (UVA), Ensemble Approaches to Semi-Supervised Learning Learning
15 Daniel Hennes (UM), Multiagent Learning - Dynamic Games and Applications
16 Eric Kok (UU), Exploring the practical benefits of argumentation in multi-agent deliberation
17 Koen Kok (VU), The PowerMatcher: Smart Coordination for the Smart Electricity Grid
18 Jeroen Janssens (UvT), Outlier Selection and One-Class Classification
19 Renze Steenhuizen (TUD), Coordinated Multi-Agent Planning and Scheduling
20 Katja Hofmann (UvA), Fast and Reliable Online Learning to Rank for Information Retrieval
21 Sander Wubben (UvT), Text-to-text generation by monolingual machine translation
22 Tom Claassen (RUN), Causal Discovery and Logic

23 Patricio de Alencar Silva (UvT), Value Activity Monitoring
24 Haitham Bou Ammar (UM), Automated Transfer in Reinforcement Learning
25 Agnieszka Anna Latoszek-Berendsen (UM), Intention-based Decision Support. A new way of representing and implementing clinical guidelines in a Decision Support System
26 Alireza Zarghami (UT), Architectural Support for Dynamic Homecare Service Provisioning
27 Mohammad Huq (UT), Inference-based Framework Managing Data Provenance
28 Frans van der Sluis (UT), When Complexity becomes Interesting: An Inquiry into the Information eXperience
29 Iwan de Kok (UT), Listening Heads
30 Joyce Nakatumba (TUE), Resource-Aware Business Process Management: Analysis and Support
31 Dinh Khoa Nguyen (UvT), Blueprint Model and Language for Engineering Cloud Applications
32 Kamakshi Rajagopal (OUN), Networking For Learning; The role of Networking in a Lifelong Learner's Professional Development
33 Qi Gao (TUD), User Modeling and Personalization in the Microblogging Sphere
34 Kien Tjin-Kam-Jet (UT), Distributed Deep Web Search
35 Abdallah El Ali (UvA), Minimal Mobile Human Computer Interaction
36 Than Lam Hoang (TUe), Pattern Mining in Data Streams
37 Dirk Börner (OUN), Ambient Learning Displays
38 Eelco den Heijer (VU), Autonomous Evolutionary Art
39 Joop de Jong (TUD), A Method for Enterprise Ontology based Design of Enterprise Information Systems
40 Pim Nijssen (UM), Monte-Carlo Tree Search for Multi-Player Games
41 Jochem Liem (UVA), Supporting the Conceptual Modelling of Dynamic Systems: A Knowledge Engineering Perspective on Qualitative Reasoning
42 Léon Planken (TUD), Algorithms for Simple Temporal Reasoning
43 Marc Bron (UVA), Exploration and Contextualization through Interaction and Concepts

2014 01 Nicola Barile (UU), Studies in Learning Monotone Models from Data
02 Fiona Tuliyano (RUN), Combining System Dynamics with a Domain Modeling Method
03 Sergio Raul Duarte Torres (UT), Information Retrieval for Children: Search Behavior and Solutions
04 Hanna Jochmann-Mannak (UT), Websites for children: search strategies and interface design - Three studies on children's search performance and evaluation
05 Jurriaan van Reijsen (UU), Knowledge Perspectives on Advancing Dynamic Capability
06 Damian Tamburri (VU), Supporting Networked Software Development
07 Arya Adriansyah (TUE), Aligning Observed and Modeled Behavior
08 Samur Araujo (TUD), Data Integration over Distributed and Heterogeneous Data Endpoints
09 Philip Jackson (UvT), Toward Human-Level Artificial Intelligence: Representation and Computation of Meaning in Natural Language
10 Ivan Salvador Razo Zapata (VU), Service Value Networks
11 Janneke van der Zwaan (TUD), An Empathic Virtual Buddy for Social Support
12 Willem van Willigen (VU), Look Ma, No Hands: Aspects of Autonomous Vehicle Control
13 Arlette van Wissen (VU), Agent-Based Support for Behavior Change: Models and Applications in Health and Safety Domains
14 Yangyang Shi (TUD), Language Models With Meta-information
15 Natalya Mogles (VU), Agent-Based Analysis and Support of Human Functioning in Complex Socio-Technical Systems: Applications in Safety and Healthcare
16 Krystyna Milian (VU), Supporting trial recruitment and design by automatically interpreting eligibility criteria
17 Kathrin Dentler (VU), Computing healthcare quality indicators automatically: Secondary Use of Patient Data and Semantic Interoperability
18 Mattijs Ghijsen (UVA), Methods and Models for the Design and Study of Dynamic Agent Organizations
19 Vinicius Ramos (TUE), Adaptive Hypermedia Courses: Qualitative and Quantitative Evaluation and Tool Support
20 Mena Habib (UT), Named Entity Extraction and Disambiguation for Informal Text: The Missing Link
21 Kassidy Clark (TUD), Negotiation and Monitoring in Open Environments
22 Marieke Peeters (UU), Personalized Educational Games - Developing agent-supported scenario-based training

23 Eleftherios Sidirourgos (UvA/CWI), Space Efficient Indexes for the Big Data Era
24 Davide Ceolin (VU), Trusting Semi-structured Web Data
25 Martijn Lappenschaar (RUN), New network models for the analysis of disease interaction
26 Tim Baarslag (TUD), What to Bid and When to Stop
27 Rui Jorge Almeida (EUR), Conditional Density Models Integrating Fuzzy and Probabilistic Representations of Uncertainty
28 Anna Chmielowiec (VU), Decentralized k-Clique Matching
29 Jaap Kabbedijk (UU), Variability in Multi-Tenant Enterprise Software
30 Peter de Cock (UvT), Anticipating Criminal Behaviour
31 Leo van Moergestel (UU), Agent Technology in Agile Multiparallel Manufacturing and Product Support
32 Naser Ayat (UvA), On Entity Resolution in Probabilistic Data
33 Tesfa Tegegne (RUN), Service Discovery in eHealth
34 Christina Manteli (VU), The Effect of Governance in Global Software Development: Analyzing Transactive Memory Systems.
35 Joost van Ooijen (UU), Cognitive Agents in Virtual Worlds: A Middleware Design Approach
36 Joos Buijs (TUE), Flexible Evolutionary Algorithms for Mining Structured Process Models
37 Maral Dadvar (UT), Experts and Machines United Against Cyberbullying
38 Danny Plass-Oude Bos (UT), Making brain-computer interfaces better: improving usability through post-processing.
39 Jasmina Maric (UvT), Web Communities, Immigration, and Social Capital
40 Walter Omona (RUN), A Framework for Knowledge Management Using ICT in Higher Education
41 Frederic Hogenboom (EUR), Automated Detection of Financial Events in News Text
42 Carsten Eijckhof (CWI/TUD), Contextual Multidimensional Relevance Models
43 Kevin Vlaanderen (UU), Supporting Process Improvement using Method Increments
44 Paulien Meesters (UvT), Intelligent Blauw. Met als ondertitel: Intelligence-gestuurde politiezorg in gebiedsgebonden eenheden.
45 Birgit Schmitz (OUN), Mobile Games for Learning: A Pattern-Based Approach
46 Ke Tao (TUD), Social Web Data Analytics: Relevance, Redundancy, Diversity
47 Shangsong Liang (UVA), Fusion and Diversification in Information Retrieval

---

2015 01 Niels Netten (UvA), Machine Learning for Relevance of Information in Crisis Response
02 Faiza Bukhsh (UvT), Smart auditing: Innovative Compliance Checking in Customs Controls
03 Twan van Laarhoven (RUN), Machine learning for network data
04 Howard Spoelstra (OUN), Collaborations in Open Learning Environments
05 Christoph Bösch (UT), Cryptographically Enforced Search Pattern Hiding
06 Farideh Heidari (TUD), Business Process Quality Computation - Computing Non-Functional Requirements to Improve Business Processes
07 Maria-Hendrike Peetz (UvA), Time-Aware Online Reputation Analysis
08 Jie Jiang (TUD), Organizational Compliance: An agent-based model for designing and evaluating organizational interactions
09 Randy Klaassen (UT), HCI Perspectives on Behavior Change Support Systems
10 Henry Hermans (OUN), OpenU: design of an integrated system to support lifelong learning
11 Yongming Luo (TUE), Designing algorithms for big graph datasets: A study of computing bisimulation and joins
12 Julie M. Birkholz (VU), Modi Operandi of Social Network Dynamics: The Effect of Context on Scientific Collaboration Networks
13 Giuseppe Procaccianti (VU), Energy-Efficient Software
14 Bart van Straalen (UT), A cognitive approach to modeling bad news conversations
15 Klaas Andries de Graaf (VU), Ontology-based Software Architecture Documentation
16 Changyun Wei (UT), Cognitive Coordination for Cooperative Multi-Robot Teamwork
17 André van Cleeff (UT), Physical and Digital Security Mechanisms: Properties, Combinations and Trade-offs
18 Holger Pirk (CWI), Waste Not, Want Not! - Managing Relational Data in Asymmetric Memories
19 Bernardo Tabuenca (OUN), Ubiquitous Technology for Lifelong Learners
20 Lois Vanhée (UU), Using Culture and Values to Support Flexible Coordination

21 Sibren Fetter (OUN), Using Peer-Support to Expand and Stabilize Online Learning
22 Zhemin Zhu (UT), Co-occurrence Rate Networks
23 Luit Gazendam (VU), Cataloguer Support in Cultural Heritage
24 Richard Berendsen (UVA), Finding People, Papers, and Posts: Vertical Search Algorithms and Evaluation
25 Steven Woudenberg (UU), Bayesian Tools for Early Disease Detection
26 Alexander Hogenboom (EUR), Sentiment Analysis of Text Guided by Semantics and Structure
27 Sándor Héman (CWI), Updating compressed colomn stores
28 Janet Bagorogoza (TiU), Knowledge Management and High Performance; The Uganda Financial Institutions Model for HPO
29 Hendrik Baier (UM), Monte-Carlo Tree Search Enhancements for One-Player and Two-Player Domains
30 Kiavash Bahreini (OU), Real-time Multimodal Emotion Recognition in E-Learning
31 Yakup Koç (TUD), On the robustness of Power Grids
32 Jerome Gard (UL), Corporate Venture Management in SMEs
33 Frederik Schadd (TUD), Ontology Mapping with Auxiliary Resources
34 Victor de Graaf (UT), Gesocial Recommender Systems
35 Jungxao Xu (TUD), Affective Body Language of Humanoid Robots: Perception and Effects in Human Robot Interaction

2016 01 Syed Saiden Abbas (RUN), Recognition of Shapes by Humans and Machines
02 Michiel Christiaan Meulendijk (UU), Optimizing medication reviews through decision support: prescribing a better pill to swallow
03 Maya Sappelli (RUN), Knowledge Work in Context: User Centered Knowledge Worker Support
04 Laurens Rietveld (VU), Publishing and Consuming Linked Data
05 Evgeny Sherkhonov (UVA), Expanded Acyclic Queries: Containment and an Application in Explaining Missing Answers
06 Michel Wilson (TUD), Robust scheduling in an uncertain environment
07 Jeroen de Man (VU), Measuring and modeling negative emotions for virtual training
08 Matje van de Camp (TiU), A Link to the Past: Constructing Historical Social Networks from Unstructured Data
09 Archana Nottamkandath (VU), Trusting Crowdsourced Information on Cultural Artefacts
10 George Karafotias (VUA), Parameter Control for Evolutionary Algorithms
11 Anne Schuth (UVA), Search Engines that Learn from Their Users
12 Max Knobbout (UU), Logics for Modelling and Verifying Normative Multi-Agent Systems
13 Nana Baah Gyan (VU), The Web, Speech Technologies and Rural Development in West Africa - An ICT4D Approach
14 Ravi Khadka (UU), Revisiting Legacy Software System Modernization
15 Steffen Michels (RUN), Hybrid Probabilistic Logics - Theoretical Aspects, Algorithms and Experiments
16 Guangliang Li (UVA), Socially Intelligent Autonomous Agents that Learn from Human Reward
17 Berend Weel (VU), Towards Embodied Evolution of Robot Organisms
18 Albert Meroño Peñuela (VU), Refining Statistical Data on the Web
19 Julia Efremova (Tu/e), Mining Social Structures from Genealogical Data
20 Daan Odijk (UVA), Context & Semantics in News & Web Search
21 Alejandro Moreno Célleri (UT), From Traditional to Interactive Playspaces: Automatic Analysis of Player Behavior in the Interactive Tag Playground
22 Grace Lewis (VU), Software Architecture Strategies for Cyber-Foraging Systems
23 Fei Cai (UVA), Query Auto Completion in Information Retrieval
24 Brend Wanders (UT), Repurposing and Probabilistic Integration of Data; An Iterative and data model independent approach
25 Julia Kiseleva (TU/e), Using Contextual Information to Understand Searching and Browsing Behavior
26 Dilhan Thilakarathne (VU), In or Out of Control: Exploring Computational Models to Study the Role of Human Awareness and Control in Behavioural Choices, with Applications in Aviation and Energy Management Domains
27 Wen Li (TUD), Understanding Geo-spatial Information on Social Media
28 Mingxin Zhang (TUD), Large-scale Agent-based Social Simulation - A study on epidemic prediction and control
29 Nicolas Höning (TUD), Peak reduction in decentralised electricity systems - Markets and prices for flexible planning

30 Ruud Mattheij (UvT), The Eyes Have It
31 Mohammad Khelghati (UT), Deep web content monitoring
32 Eelco Vriezekolk (UT), Assessing Telecommunication Service Availability Risks for Crisis Organisations
33 Peter Bloem (UVA), Single Sample Statistics, exercises in learning from just one example
34 Dennis Schunselaar (TUE), Configurable Process Trees: Elicitation, Analysis, and Enactment
35 Zhaochun Ren (UVA), Monitoring Social Media: Summarization, Classification and Recommendation
36 Daphne Karreman (UT), Beyond R2D2: The design of nonverbal interaction behavior optimized for robot-specific morphologies
37 Giovanni Sileno (UvA), Aligning Law and Action - a conceptual and computational inquiry
38 Andrea Minuto (UT), Materials that Matter - Smart Materials meet Art & Interaction Design
39 Merijn Bruijnes (UT), Believable Suspect Agents; Response and Interpersonal Style Selection for an Artificial Suspect
40 Christian Detweiler (TUD), Accounting for Values in Design
41 Thomas King (TUD), Governing Governance: A Formal Framework for Analysing Institutional Design and Enactment Governance
42 Spyros Martzoukos (UVA), Combinatorial and Compositional Aspects of Bilingual Aligned Corpora
43 Saskia Koldijk (RUN), Context-Aware Support for Stress Self-Management: From Theory to Practice
44 Thibault Sellam (UVA), Automatic Assistants for Database Exploration
45 Bram van de Laar (UT), Experiencing Brain-Computer Interface Control
46 Jorge Gallego Perez (UT), Robots to Make you Happy
47 Christina Weber (UL), Real-time foresight - Preparedness for dynamic innovation networks
48 Tanja Buttler (TUD), Collecting Lessons Learned
49 Gleb Polevoy (TUD), Participation and Interaction in Projects. A Game-Theoretic Analysis
50 Yan Wang (UVT), The Bridge of Dreams: Towards a Method for Operational Performance Alignment in IT-enabled Service Supply Chains

2017 01 Jan-Jaap Oerlemans (UL), Investigating Cybercrime
02 Sjoerd Timmer (UU), Designing and Understanding Forensic Bayesian Networks using Argumentation
03 Daniël Harold Telgen (UU), Grid Manufacturing; A Cyber-Physical Approach with Autonomous Products and Reconfigurable Manufacturing Machines
04 Mrunal Gawade (CWI), Multi-core Parallelism in a Column-store
05 Mahdieh Shadi (UVA), Collaboration Behavior
06 Damir Vandic (EUR), Intelligent Information Systems for Web Product Search
07 Roel Bertens (UU), Insight in Information: from Abstract to Anomaly
08 Rob Konijn (VU) , Detecting Interesting Differences:Data Mining in Health Insurance Data using Outlier Detection and Subgroup Discovery
09 Dong Nguyen (UT), Text as Social and Cultural Data: A Computational Perspective on Variation in Text
10 Robby van Delden (UT), (Steering) Interactive Play Behavior
11 Florian Kunneman (RUN), Modelling patterns of time and emotion in Twitter #anticipointment
12 Sander Leemans (TUE), Robust Process Mining with Guarantees
13 Gijs Huisman (UT), Social Touch Technology - Extending the reach of social touch through haptic technology
14 Shoshannah Tekofsky (UvT), You Are Who You Play You Are: Modelling Player Traits from Video Game Behavior
15 Peter Berck (RUN), Memory-Based Text Correction
16 Aleksandr Chuklin (UVA), Understanding and Modeling Users of Modern Search Engines
17 Daniel Dimov (UL), Crowdsourced Online Dispute Resolution
18 Ridho Reinanda (UVA), Entity Associations for Search
19 Jeroen Vuurens (UT), Proximity of Terms, Texts and Semantic Vectors in Information Retrieval
20 Mohammadbashir Sedighi (TUD), Fostering Engagement in Knowledge Sharing: The Role of Perceived Benefits, Costs and Visibility
21 Jeroen Linssen (UT), Meta Matters in Interactive Storytelling and Serious Gaming (A Play on Worlds)
22 Sara Magliacane (VU), Logics for causal inference under uncertainty
23 David Graus (UVA), Entities of Interest — Discovery in Digital Traces
24 Chang Wang (TUD), Use of Affordances for Efficient Robot Learning

25 Veruska Zamborlini (VU), Knowledge Representation for Clinical Guidelines, with applications to Multimorbidity Analysis and Literature Search
26 Merel Jung (UT), Socially intelligent robots that understand and respond to human touch
27 Michiel Joosse (UT), Investigating Positioning and Gaze Behaviors of Social Robots: People's Preferences, Perceptions and Behaviors
28 John Klein (VU), Architecture Practices for Complex Contexts
29 Adel Alhuraibi (UvT), From IT-BusinessStrategic Alignment to Performance: A Moderated Mediation Model of Social Innovation, and Enterprise Governance of IT"
30 Wilma Latuny (UvT), The Power of Facial Expressions
31 Ben Ruijl (UL), Advances in computational methods for QFT calculations
32 Thaer Samar (RUN), Access to and Retrievability of Content in Web Archives
33 Brigit van Loggem (OU), Towards a Design Rationale for Software Documentation: A Model of Computer-Mediated Activity
34 Maren Scheffel (OU), The Evaluation Framework for Learning Analytics
35 Martine de Vos (VU), Interpreting natural science spreadsheets
36 Yuanhao Guo (UL), Shape Analysis for Phenotype Characterisation from High-throughput Imaging
37 Alejandro Montes Garcia (TUE), WiBAF: A Within Browser Adaptation Framework that Enables Control over Privacy
38 Alex Kayal (TUD), Normative Social Applications
39 Sara Ahmadi (RUN), Exploiting properties of the human auditory system and compressive sensing methods to increase noise robustness in ASR
40 Altaf Hussain Abro (VUA), Steer your Mind: Computational Exploration of Human Control in Relation to Emotions, Desires and Social Support For applications in human-aware support systems
41 Adnan Manzoor (VUA), Minding a Healthy Lifestyle: An Exploration of Mental Processes and a Smart Environment to Provide Support for a Healthy Lifestyle
42 Elena Sokolova (RUN), Causal discovery from mixed and missing data with applications on ADHD datasets
43 Maaike de Boer (RUN), Semantic Mapping in Video Retrieval
44 Garm Lucassen (UU), Understanding User Stories - Computational Linguistics in Agile Requirements Engineering
45 Bas Testerink (UU), Decentralized Runtime Norm Enforcement
46 Jan Schneider (OU), Sensor-based Learning Support
47 Jie Yang (TUD), Crowd Knowledge Creation Acceleration
48 Angel Suarez (OU), Collaborative inquiry-based learning

2018 01 Han van der Aa (VUA), Comparing and Aligning Process Representations
02 Felix Mannhardt (TUE), Multi-perspective Process Mining
03 Steven Bosems (UT), Causal Models For Well-Being: Knowledge Modeling, Model-Driven Development of Context-Aware Applications, and Behavior Prediction
04 Jordan Janeiro (TUD), Flexible Coordination Support for Diagnosis Teams in Data-Centric Engineering Tasks
05 Hugo Huurdeman (UVA), Supporting the Complex Dynamics of the Information Seeking Process

“Listen, Morty, I hate to break it to you but what people call “love” is just a chemical reaction that compels animals to breed. It hits hard, Morty, then it slowly fades, leaving you stranded in a failing marriage. I did it. Your parents are gonna do it. Break the cycle, Morty. Rise above. Focus on science.”

Rick Sanchez