

ARMED: ARgumentation Mining and reasoning about Enterprise architecture Decisions

Marc van Zee ^a

Dirk van der Linden^b

^a *University of Luxembourg, Luxembourg*

^b *University of Haifa, Israel*

Abstract

In this position paper we outline a research agenda for automatically capturing architectural decisions from documents of the Enterprise Architecture (EA) domain (e.g. minutes of meetings), and to develop formal methods to reason about extracted EA decisions and their underlying arguments.

1 Introduction

An Enterprise Architecture (EA) is used to model large enterprises in a holistic fashion by connecting their IT infrastructure and applications to the business processes they support. In turn this links them also to the products and services that are realized by those business processes [8, 15]. While a multitude of EA modelling languages have been developed in the past decades [7, 12, 24], the design decisions behind the resulting models are often left implicit. In the few frameworks that exist (c.f., [14, 34]), one of the biggest, and often avoided, issues is the *return of capturing effort*: Architects often have no time or incentive to explicitly document all of the decisions they make. This is a well-known problem in related areas such as Software Architecture (SA) as well [11].

The need for Enterprise Architecture is emphasized by the Commission Wijffels [35], who conduct a major study on how the willingness and stability of banks in the Netherlands can be improved. The motivation for this study is that “banks have been insufficiently willing and stable in the last years.” One of their main recommendation is for every bank to use an Enterprise Architecture (EA) as a central component. Failing to document EA related decisions can cause design integrity issues when architects want to maintain or change the current design. This means that due to a lacking insight of the rationale, new designs are constructed in an adhoc manner, without taking into consideration constraints implied by past design decisions [23]. Moreover the lessons learned in terms of anticipated and unanticipated consequences of design decisions are lost [18]. The need to explicitly document architectural decisions is further confirmed by a survey amongst software architects on recording design rationales [22]. In this survey, a large majority of architects (85,1%) agree with the importance of architectural rationalization.




Figure 1: Venn diagram of the related research areas

In this paper, we set out a research agenda to develop the ARMED framework (ARgumentation Mining and reasoning about Enterprise architecture Decisions). The framework is used to extract decisions from documents (e.g. minutes of meetings) in the EA domain using Natural Language Processing (NLP) techniques, and uses formal methods to reason about extracted enterprise architecture decisions and their underlying arguments. The ARMED framework lies at the intersection of three lines of research, namely enterprise architecture, formal reasoning, and argumentation mining (see Figure 1).

The rest of this paper is organized as follows: In Section 2 we give a cursory overview of relevant literature, discussing the three lines of research of Figure 1, in Section 3 we propose how to develop the ARMED framework by formulating our objectives, and in Section 4 we describe our methods and approach for each of the research objectives.

2 Literature Survey

We describe the state of the art in enterprise architecture, focussing on decision rationalisation, reasoning, focussing on argumentation and BDI theory, and argumentation mining.

2.1 Enterprise Architecture: Decision Rationalisation

Several EA modeling language have been developed over the past decades, of which we highlight three. In the ArchiMate project [12], the ArchiMate language for the representation of enterprise architectures was developed. ArchiMate is an open and independent modeling language to support the description, analysis and visualization of architecture (although such visualizations are not necessarily made with the involvement of their intended users [25]) within and across business domains in an unambiguous way, although the degree to which only a language itself can contribute to decreasing ambiguity between stakeholders has been investigated and questioned [26, 27]. The Unified Profile for DoDAF/MODAF (UPDM) [7] is another visual modeling standard that supports the development of architectures that comply with the USA Department of Defense Architecture Framework (DoDAF) and the UK Ministry of Defence Architecture Framework (MODAF). On a higher level, the Open Group standard TOGAF [24] was developed as a de facto global standard for Enterprise Architecture. The Open Group Architecture Forum, comprised of more than 200 enterprises, develops and maintains the TOGAF standard and publishes successive versions at regular intervals.

In the ArchiMate project, initial research into the capturing of architectural design decisions was conducted. Such work is needed to figure out what conceptual distinctions are important to accommodate in a language [28]. Following this preliminary work, we recently formalized an existing framework for EA decision making using graph and set theory, together with a number of integrity constraints that can be used as consistency checks on the formal framework [34]. In a similar attempt for decision rationalisation in enterprise architecture, we apply the ASPIC+ framework for structured argumentation to the Goal-oriented Requirements Language (GRL) [33, 30]. One of the biggest unresolved issues in each of the above mentioned frameworks is the return of capturing effort issue: Architects often have no time or incentive to explicitly document all of the decisions they make.

2.2 Reasoning

One of the outcomes of a recent survey that we conducted among a group of 35 enterprise architects is that quantitative data is desirable in decision making, but that this is not always available [29]. In fact, documents from the EA domain are more likely to use qualitative notions than quantitative ones. As a result, the large body of literature on decision theory based on probability distributions and utilities functions is not directly applicable. Therefore, the ARMED framework is based on a qualitative decision theory by combining a BDI (Belief-Desire-Intention) theory with formal argumentation, in order to represent decisions and their underlying arguments. We discuss these two areas in turn.

BDI Theory The logical framework of ARMED uses planning theories of beliefs, desires, and intentions, which have been developed specifically for partial plans of resource-bounded agents. This is motivated by one of the main findings of our survey, namely that plans that are being developed in EA decision making are usually partial and subject to change. We base our work on a recent proposal by

Yoav Shoham [19], who proposes to change the focus on intentions from its historical philosophical perspective to a computer science database perspective of revising plans in the context of beliefs. Shoham's framework is largely informal, but it has been formalized in a temporal logic by Icard *et al.* [9]. We have recently extended this work by linking revision of intentions to a semantical model [32, 31]. Shoham uses this logical framework as the basis of a start-up company called TimeFul¹, which develops a "smart calendar" application for smartphones. The app stores the users' decisions about the future and is able to reason about these decisions. Timeful has recently been acquired by Google for several million dollars, which confirms the relevance of this line of research. Although Shoham's framework is in principle suitable to formalize EA decisions on an abstract level, it is not possible to reason about the underlying arguments in this framework. As such, we propose to extend Shoham's framework with formal argumentation, such that it is possible to automatically translate EA decisions and their underlying arguments into a logical form, amendable for automatic reasoning.

Formal Argumentation Argumentation plays an important role in the communication of human and artificial agents and is a ubiquitous task in professional and everyday life. With the breakthrough of formal logic in the 20th century, semantical and proof theoretical questions became dominating issues in logic, especially following the influential paper of Dung [6] on argumentation semantics, providing a standard calculus for determining acceptable, conflict free sets of arguments from information that may be inconsistent. Argumentation has become an important area of study in artificial intelligence over the last fifteen years, especially in subfields such as nonmonotonic reasoning, multiple-source information systems, and modelling interactions between agents. Various aspects of decisions have been formalized in argumentation systems. For instance, Amgoud and Prade [2] propose an abstract argumentation framework for decision making based on decision principles that determine alternatives for decisions. Jureta *et al.* [10] propose a method to guide the justification of goal modeling choices, and include a detailed formal argumentation model that links to goal models. We recently proposed a similar framework for the Goal-oriented Requirements Language [30]. However, none of the above approaches attempts to parse argument from documents automatically using NLP techniques. The ARMED framework does so using argumentation mining, which we discuss in the next subsection. The forthcoming Handbook of Formal Argumentation² provides an overview of the field, ranging from abstract argumentation and non-numerical argumentation to algorithms and applications. ARMED is positioned in the last chapter "Natural language argumentation" of the last section (Natural language argumentation) of the handbook.

2.3 Argumentation Mining

The amount of data that is used in the decision making process in large enterprises is increasing rapidly. Big data systems store vast amounts of data. For many organizations the value of these systems can be summed up in one word: analytics. Big data systems allow organizations to analyse data in ways never been possible before. At the recent Landelijke Architecture Congress³ (the Dutch natural architecture conference), the main focus was on big data and how it can support architects in the decision making process. As such, big data and text mining are currently one of the most important challenges in society. The ARMED framework uses argumentation mining techniques to mine decisions and underlying arguments from EA documents.

Argumentation mining is a relatively new area in corpus-based discourse analysis involving the identification of argumentative structures within a document automatically, e.g., the premises, conclusion, and argumentation scheme of each argument, as well as argument-subargument and argument-counterargument relationships between different arguments that may occur in the document. Researchers have developed methods for argumentation mining in many different areas, such as legal documents [5, 3], on-line debates, product reviews, and user comments on proposed regulations [16]. An older strand of research (that uses the term 'argumentative structure' in a somewhat different sense than we do) has investigated the automatic classification of the sentences of a scientific article's abstract or full text in terms of their contribution of new knowledge to a field (e.g., [20]). Argumentation mining has ties to sentiment analysis as well (e.g., [21]). Currently, there are a few corpora with annotations for argu-

¹<http://www.timeful.com/>

²<http://formalargumentation.org/>

³<http://www.lacongres.nl/>

mentation mining research [5] although corpora with annotations for argument sub-components have recently become available (e.g., [16]).

3 Objectives

Our overall objective is to design a parser to extract data from documents of the EA domain (e.g. minutes of meetings), and to develop formal methods to reason about extracted enterprise architecture decisions and their underlying arguments. This breaks down into the following objectives:

1. Developing a logical system to formalize EA decisions and arguments by applying knowledge representation techniques on a corpus of EA documents, and validating this framework through a series of field studies.
2. Developing an argumentation system (for instance based on the ASPIC+ [17] framework) to reason about the dynamics of decisions and their underlying arguments using a theory of intentions.
3. Developing mining patterns for arguments and decisions, and implementing these patterns into a prototype able to parse the corpus of EA documents into a structured form.

The scope of the ARMED framework is the enterprise architecture decision making process as it is documented in EA documents. An “EA document” is in principle any document that an enterprise architect may use to make decisions. As such, they can be strategic documents, architectural documents, start architecture, ArchiMate models, design documents, minutes of meetings, or even emails. However, as a first step we focus on sentences in the minutes that are used in meetings. This includes for instance sentences of the form: “The board decided to invest in advertisement, because there was enough budget available.” or “The majority of members agreed on Lease-or-Buy Consultant motion to use Lease-or-Buy Consultants.” Preliminary analysis on such sentences shows that minutes contain decisions in an argumentative way (e.g., they give also reasons).

The motivation of the ARMED framework is that having the ability to reason about decisions and their underlying arguments in terms of a logic-based framework: (1) enables consistency checks of the underlying rationales, (2) enables a more precise capturing of design knowledge, and (3) enables impact/what-if analysis when confronted with changes to the decisions or underlying arguments.

4 Methods and Approach

The success criterion of ARMED is successfully developing the proof of concept system in Figure 2, where an arrow $A \rightarrow B$ means that A serves as input to B .

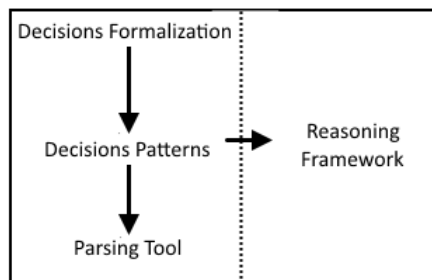


Figure 2: ARMED success criterion

- *Decisions Formalization*: A logical system (syntax and semantics) is developed to represent sentences in EA documents (objective 1).
- *Decisions/argument patterns*: A corpus of patterns for decisions and arguments are developed based on the formalisation.

- *Parsing Tool*: A parsing tool is developed that is able to label pieces of text from the EA documents by implementing the corpus patterns (objective 3).
- *Reasoning Framework*: A reasoning framework is developed that can reason about the dynamics of decisions and arguments (objective 2). The reasoning framework can perform conflict detection (consistency checks) and what-if analysis (describing the effect of changing a decision or an argument on the rest of the system).

Objective 1 We analyse a corpus of EA documents containing decisions and arguments that have been put forward for these decisions. For this, we have received initial input from three industrial partners from the Netherlands. These are large non-profit organisations, each consisting of over 2000 employees and six enterprise architects. As a starting point, we focus on analysing minutes of meetings from the EA domain. We apply knowledge representation techniques on this corpus, meaning that we characterize the sentences in the corpus, and then find a suitable formal (logical) language to model these characteristics. Suppose for instance a minute of a meeting contains the following sentence: “The board decided to invest in advertisement, because there was enough budget available”. The decision “invest in advertisement” is based on the argument “there was enough budget available”. However, there may be counter-arguments against this argument, or additional support in terms of evidence or previous decisions. We see the language that Amgoud et al. [1] developed recently as a useful starting point. Starting from a requirement by the NLP community to mine arguments as reasons about claims, they develop a logic for reasoning about reasons and claims. For instance, the above sentence is formalized as follows:

$$R(\text{enough_budget}) : C(\text{invest_in_advertisement}).$$

Both premise and conclusion can be negated, so it is for instance possible to have reasons for not claiming something, or not having a reason to claim something.

Objective 2 We recently conducted a survey among a group of 35 EA practitioners [29]. Two of the main findings are that EA decision making is a group process, and that due to the uncertainty of the domain, the plans that are being developed in EA are usually partial and subject to change. Therefore, we use the logic of intention from a database perspective by Shoham [19], and our recent extensions [32, 31] as our starting point. In order to develop a system for reasoning about conflicting intentions using formal argumentation, we see the ASPIC+ system [17] for structured argumentation as a promising candidate, because it has been developed by abstracting away from a logical theory. Therefore, we aim to instantiate the ASPIC+ system with the logic for intentions and we use the argumentation structure to characterize conflict resolution for intentions.

Objective 3 We implement a concrete pipeline from EA documents to a logical framework, which could be extended and used in commercial systems in the future to translate EA documents into a more structured, logical format. The TULE dependency parser will be used and extended in order to automatically build from text the logical formulae given in input to the ASPIC+ framework. The TULE dependency parser [13] is one of the open-source parsers with the best attested performances for Italian [4]. For English, performances are quite close to those of the Stanford parser, one of the most popular open source parsers. The great advantages of TULE is that it is rule-based, contrary to most other available parsers for Italian and English, such as the Stanford parser, that are statistical. A statistical parser needs an initial corpus of annotated trees on which it is possible to train the statistical model that the parser will use to annotate new text. Moreover, it is often the case that the statistical model, once trained on a bigger corpus, is able to properly parse new linguistic constructions but, at the same time, is no longer able to properly parse linguistic constructions that were previously handled. As a result, the performances of these parsers tend to oscillate around a certain maximum value of accuracy (usually, 90% of accuracy).

We extend the TULE parser with specific rules for dealing with minutes, thus enhancing the performance much beyond those provided by available statistical parsers. These parser are trained on “generic” corpora while, in order to produce reliable results in the present project, they should be trained on syntactically annotated corpora of minutes, which are currently unavailable. Furthermore, we design and develop a semantic role labeller specific for the output of the parser TULE. Semantic role labelling

is needed to identify the decisions and the arguments in text, from which we will build the input formulae given in input to the ASPIC+ framework. The solution that will be developed for the ARMED framework is again a rule-based syntax-semantic interface that takes as input the output of the TULE parser and builds the input formulae of the ASPIC+ framework. For instance, with respect to the sentence “The board decided to invest in advertisement, arguing there was enough budget available”, the TULE parser will identify “The board” as the syntactic subject of the main verb “decided”, “to invest in advertisement” as its syntactic complement, and “arguing there was enough budget available” as one of its adjuncts. Thus, the knowledge base of the syntax-semantic interface contains a rule associated to the verb “decide” saying that the subjects of that verb correspond to a formula A_0 and their complement to a formula A_1 that must given in input to the ASPIC+ framework. Furthermore, in case the verb includes an adjunct headed by the verb “to argue”, that adjunct correspond to a formula A_2 and the ASPIC+ framework is also fed with the logical dependency $A_2 \rightarrow A_0$, specifying that A_2 is an argument of the decision A_0 . Again, we expect that a rule-based semantic role labeller specifically built for dealing with minutes will have performances higher than statistical semantic role labeller available online.⁴

5 Conclusion

One of the biggest issues with decision capturing in Enterprise Architecture and related areas is the return of capturing effort. While almost everybody is convinced that it is useful and profitable to document decisions, it is often not done. Architects see no direct benefit of documenting a decisions, and there may be not time to do so either. In order to resolve this problem, we set out a research agenda to develop the ARMED framework. The ARMED framework consists of a parser that takes as input documents from the EA domains (e.g., minutes of meetings) and translates these into a logical form, focusing on decisions and arguments for these decisions. In order to reason about these entities, we propose to use a BDI logic to formalize EA decisions, and to extend it with argumentation to formalize underlying arguments. In order to extract decisions and arguments from text we use argumentation mining techniques. As such, the ARMED framework extends current knowledge on the crossroads of reasoning, logical methods, argumentation mining, and enterprise architecture. At the same time, by extending an existing parser, we expect that our research has practical uses as well.

In short, the ARMED framework is meant to lay the basis for a decision support system in Enterprise Architecture, one that can perform consistency checks on decisions and arguments, enables a more precise capturing of design knowledge, and enables impact/what-if analysis when confronted with changes to the decisions or underlying arguments.

6 Acknowledgements

We thank Leon van der Torre and Livio Robaldo for useful suggestions and discussions.

References

- [1] L. Amgoud, B. Besnard, and A. Hunter. Representing and reasoning about arguments mined from texts and dialogues. In *Proceedings of the 13th European Conference on Symbolic and Quantitative Approaches to Reasoning with Uncertainty (ECSQARU)*, 2015.
- [2] Leila Amgoud and Henri Prade. Using arguments for making and explaining decisions. *Artificial Intelligence*, 173(34):413 – 436, 2009.
- [3] Kevin D. Ashley and Vern R. Walker. Toward constructing evidence-based legal arguments using legal decision documents and machine learning. In *Proceedings of the Fourteenth International Conference on Artificial Intelligence and Law, ICAIL '13*, pages 176–180, New York, NY, USA, 2013. ACM.

⁴E.g. the one available at <http://barbar.cs.lth.se:8081/parse>

- [4] Cristina Bosco and Alessandro Mazzei. The evalita dependency parsing task: From 2007 to 2011. In Bernardo Magnini, Francesco Cutugno, Mauro Falcone, and Emanuele Pianta, editors, *EVALITA*, volume 7689 of *Lecture Notes in Computer Science*, pages 1–12. Springer, 2011.
- [5] Glenn Rowe Chris Reed, Raquel Mochales Palau and Marie-Francine Moens. Language resources for studying argument. In Nicoletta Calzolari (Conference Chair), Khalid Choukri, Bente Mae-gaard, Joseph Mariani, Jan Odijk, Stelios Piperidis, and Daniel Tapias, editors, *Proceedings of the Sixth International Conference on Language Resources and Evaluation (LREC'08)*, Marrakech, Morocco, may 2008. European Language Resources Association (ELRA). <http://www.lrec-conf.org/proceedings/lrec2008/>.
- [6] Phan Minh Dung. On the acceptability of arguments and its fundamental role in nonmonotonic reasoning, logic programming and n-person games. *Artif. Intell.*, 77(2):321–358, 1995.
- [7] Object Management Group. Unified profile for DoDAF and MODAF (UPDM) Technical report. <http://www.omg.org/spec/UPDM/2.0>, 2012.
- [8] Jan A P Hoogervorst. *Enterprise governance and enterprise engineering*. Enterprise Engineering Series. Springer, Dordrecht, 2009.
- [9] Thomas Icard, Eric Pacuit, and Yoav Shoham. Joint revision of belief and intention. *Proc. of the 12th International Conference on Knowledge Representation*, pages 572–574, 2010.
- [10] Ivan Jureta, John Mylopoulos, and Stéphane Faulkner. Analysis of multi-party agreement in requirements validation. In *Proceedings of the 2009 17th IEEE International Requirements Engineering Conference, RE, RE '09*, pages 57–66, Washington, DC, USA, 2009. IEEE Computer Society.
- [11] Philippe Kruchten, Rafael Capilla, and Juan Carlos Dueas. The decision view's role in software architecture practice. *Software, IEEE*, 26(2):36–42, 2009.
- [12] M. Lankhorst. *Enterprise architecture at work: modelling, communication, and analysis*. Springer, 2005.
- [13] L. Lesmo. The Rule-Based Parser of the NLP Group of the University of Torino. *Intelligenza Artificiale*, 2(4):46–57, June 2007.
- [14] Carlos Martinez. Enterprise architectures: A just-in-time approach for decision-making. <http://www2.mitre.org/public/eabok/pdf/Using-EAs-14-1210.pdf>, 2014. [Accessed: 5-September-2015].
- [15] M. Op 't Land, H.A. Proper, M. Waage, J. Cloo, and C. Steghuis. *Enterprise Architecture – Creating Value by Informed Governance*. Enterprise Engineering Series. Springer, Berlin, Germany, 2008.
- [16] Joonsuk Park and Claire Cardie. Assess: A tool for assessing the support structures of arguments in user comments. In *Computational Models of Argument - Proceedings of COMMA 2014, Atholl Palace Hotel, Scottish Highlands, UK, September 9-12, 2014*, pages 473–474, 2014.
- [17] Henry Prakken. An abstract framework for argumentation with structured arguments. *Argument and Computation*, 1(2):93–124, 2010.
- [18] HA Proper and M Op t Land. Lines in the water. In *Practice-Driven Research on Enterprise Transformation 2010 (PRET-2010)*, pages 193–216. Springer Berlin Heidelberg, 2010.
- [19] Yoav Shoham. Logical theories of intention and the database perspective. *Journal of Philosophical Logic*, 2009.
- [20] Ekaterina Shutova and Simone Teufel. Metaphor corpus annotated for source - target domain mappings. In *Proceedings of the International Conference on Language Resources and Evaluation, LREC 2010, 17-23 May 2010, Valletta, Malta*, 2010.

- [21] Swapna Somasundaran and Janyce Wiebe. Recognizing stances in ideological on-line debates. In *Proceedings of the NAACL HLT 2010 Workshop on Computational Approaches to Analysis and Generation of Emotion in Text*, CAAGET '10, pages 116–124, Stroudsburg, PA, USA, 2010. Association for Computational Linguistics.
- [22] Antony Tang, Muhammad Ali Babar, Ian Gorton, and Jun Han. A survey of architecture design rationale. *Journal of Systems and Software*, 79(12):1792 – 1804, 2006.
- [23] Antony Tang, Yan Jin, and Jun Han. A rationale-based architecture model for design traceability and reasoning. *J. Syst. Softw.*, 80(6):918–934, June 2007.
- [24] The Open Group. TOGAF version 9.1. the open group architecture framework. 2011.
- [25] Dirk van der Linden. An argument for more user-centric analysis of modeling languages' visual notation quality. In Anne Persson and Janis Stirna, editors, *Advanced Information Systems Engineering Workshops*, volume 215 of *Lecture Notes in Business Information Processing*, pages 114–120. Springer International Publishing, 2015.
- [26] Dirk van der Linden and S.J.B.A. Hoppenbrouwers. Challenges of identifying communities with shared semantics in enterprise modeling. In K. Sandkuhl, U. Seigerroth, and J. Stirna, editors, *The Practice of Enterprise Modeling*, volume 134 of *Lecture Notes in Business Information Processing*, pages 160–171. Springer, Berlin, Germany, 2012.
- [27] Dirk van der Linden, S.J.B.A. Hoppenbrouwers, A. Lartseva, and W. Molnar. Beyond terminologies: Using psychometrics to validate shared ontologies. *Applied Ontology*, 7(4):471–487, 2012.
- [28] Dirk van der Linden and H.A. Proper. On the accommodation of conceptual distinctions in conceptual modeling languages. In *Proceedings of Modellierung 2014 (MoK 2014)*, volume 225 of *Lecture Notes in Informatics*, 2014.
- [29] Dirk van der Linden and Marc van Zee. Insights from a Study on Decision Making in Enterprise Architecture. *To appear in: Short Paper Proceedings of the 8th IFIP WG 8.1 Working Conference on the Practice of Enterprise Modeling (PoEM)*, 2015.
- [30] Marc van Zee, Floris Bex, and Sepideh Ghanavati. Rationalization of goal models in grl using formal argumentation. In *Proceedings of RE: Next! track at the Requirements Engineering Conference 2015 (RE'15)*, August 2015.
- [31] Marc van Zee, Mehdi Dastani, Dragan Doder, and Leendert van der Torre. Agm revision of beliefs about action and time. In *Proceedings of the International Joint Conference on Artificial Intelligence*, July 2015.
- [32] Marc van Zee, Mehdi Dastani, Dragan Doder, and Leendert van der Torre. Consistency conditions for beliefs and intentions. In *Proceedings of the Twelfth International Symposium on Logical Formalizations of Commonsense Reasoning*, March 2015.
- [33] Marc van Zee and Sepideh Ghanavati. Capturing evidence and rationales with requirements engineering and argumentation-based techniques. In *Proceedings of the 26th Benelux Conference on Artificial Intelligence (BNAIC2014)*, November 2014.
- [34] Marc van Zee, Georgios Plataniotis, Diana Marosin, and Dirk van der Linden. Formalizing enterprise architecture decision models using integrity constraints. In *16th IEEE Conference on Business Informatics (CBI)*, April 2014.
- [35] Commission Wijffels. Naar een dienstbaar en stabiel bankwezen. <http://www.rijksoverheid.nl/bestanden/documenten-en-publicaties/kamerstukken/2013/06/28/rapport-naar-een-dienstbaar-en-stabiel-bankwezen/rapport-naar-een-dienstbaar-en-stabiel-bankwezen.pdf>, 2013. Accessed: 15-April-2015.