Improving interpretations of trust claims

Marc Sel

Information Security Group,
Royal Holloway, University of London
Egham, Surrey, TW20 0EX, United Kingdom
Marc.Sel.2013@live.rhul.ac.uk
marc@marcsel.eu

Abstract. This paper presents an approach to semantic modelling of large-scale trust ecosystems to improve the interpretations of trust claims. The problem of interpreting trust claims is described and relevant types of reasoning are analysed. A model based on $\mathcal{SROIQ}$ and OWL Description Logic is proposed. The novel elements are the creation of classes and properties on the basis of legal and regulatory sources that extend existing vocabularies (W3C, Dublin Core), and the use of these classes and properties to create assertions that represent information harvested from on-line information sources. The resulting model allows automated classification via a reasoner, as well as queries that support use cases from various actors. A general approach is presented, as well as results from a prototype implementation based on the European eIDAS and US FICAM trust ecosystems.

1 Introduction

1.1 Context

The digital society will continue to increase its reliance on electronic transactions. Such transactions are conducted between Service Providers and Service Consumers, possibly with the use of intermediaries. Relying on the outcome of a transaction performed via an ICT system, or making a selection which system to use in the first place, forces the user to take a trust decision. While the notion of trust is in widespread use, its meaning varies. For a basic treatment, refer to Gambetta et al. [4], Marsh [9] or Cofta [2].

1.2 Motivation

The motivation for the research described below stems from two observations. First, understanding what a specific trust claim actually means, what it is based on, as well as why it should be considered valid is still hard, and there is often room for different interpretations. This article promotes the view that one should not 'trust' but rather take an informed decision on the basis of evidence and reasoning. Second, various actors publish reasonably independent information on other other actors in the same ecosystem. For example regulators, central banks, and business registers provide contextual information that can contribute to verifying a claim. Today's trust models typically include such contextual information only in a limited way. More extensive usage of such information under formal semantics could potentially strengthen the verification of claims because it adds information typically from beyond the control of the actor whose claim is validated. It is to the benefit of honest parties that reliance on a transaction is based on a trust model with semantics and evidence understandable and agreeable to all.

1.3 Research contributions

This paper researches the type of reasoning that would allow a limitation of interpretation of trust claims. The problem of interpreting trust claims is described. A novel trust modelling approach is proposed, based on a Trust Claim Interpretation model that answers queries resulting from execution of a trust policy validation algorithm. Novel elements are the creation of classes and properties on the basis of legal and regulatory sources that extend existing vocabularies (W3C, Dublin Core), and the use of these classes and properties to create assertions that represent information harvested from on-line information sources. An implementation based on $\mathcal{SROIQ}$ and OWL Description Logic is presented.

1.4 Paper outline

The preceding section set the context and motivation, and provided an introduction to the research contribution. Section 2 describes the various types of trust statements addressed, and what existing work has been done in the area. Section 3 describes a new approach to trust modelling, including a novel trust modelling architecture. Section 4 discusses a prototype implementation, based on the choice of $\mathcal{SROIQ}$ and OWL DL. In section 5 strengths and weaknesses are analysed, as well as areas for further research.

2 Trust, trust modelling and prior art

2.1 Trust and trust modelling

A key part of the development of the electronic society is the introduction of an economy based on electronic transactions and trust. Transactions are conducted between Service Providers and Service Consumers. Trust can be provided by a range of possible mechanisms including, but not limited to, cryptographic protocols and legal or contractual liability. The meaning of 'trust' varies according to the circumstances, and the perspective of the trustee (who is trusted) and the trustor (who is trusting). Trust in cryptographic protocols, often relying on Trusted Third Parties, supports many Internet or closed user group transactions.

Qualifying information or a service in an electronic form as 'trusted' is non-trivial. Many different actors, mechanisms and artefacts collaborate to perform electronic trust transactions. In [12] an informal domain model was introduced. An introduction to eIDAS and FICAM trust models is provided in [13].

2.2 Prior art

With regard to trust, much research has been conducted to represent real world information and use it as the basis for decisions. A trust calculus for PKI and identity management is proposed by Huang and Nicol, [7]. Measuring and computing trust using subjective logic has been studied by Josang [8]. Hartig defined a trust-aware extension to SPARQL [5]. The Web Of Trust (WOT) project[1] defined an basic RDF vocabulary to facilitate the use of Public Key Cryptography. Shekarpour and Katebi reviewed trust calculation and models of trust rating, and proposed algorithms for propagation and aggregation of trust [14]. A formal notion of trust to enable reasoning about security properties is proposed by Fuchs, Gürgens and Rudolph [3].

[1] http://xmlns.com/wot/0.1/

3 Trust modelling: a new approach

3.1 Outline

A model for reduction of interpretations of trust claims is proposed, combining mathematical modelling with harvesting artefacts that include contextual information, followed by reasoning according to a well specified logic. In figure 1 the real world is represented by a globe from where two abstractions are derived. The first abstraction is composed of the actors in the left box. The actors' transactions rely on one or more trust models. To validate a particular reliance, a trust policy validation algorithm attempts to satisfy assumptions by issuing trust claim interpretation requests. The second abstraction is the trust claim interpretation (TCI) model, responding to these requests with responses. For this purpose, the TCI model contains a query engine as well as a knowledge base. The knowledge base contains assertions imported from the real world, and its contents is maintained consistent by a reasoning engine.

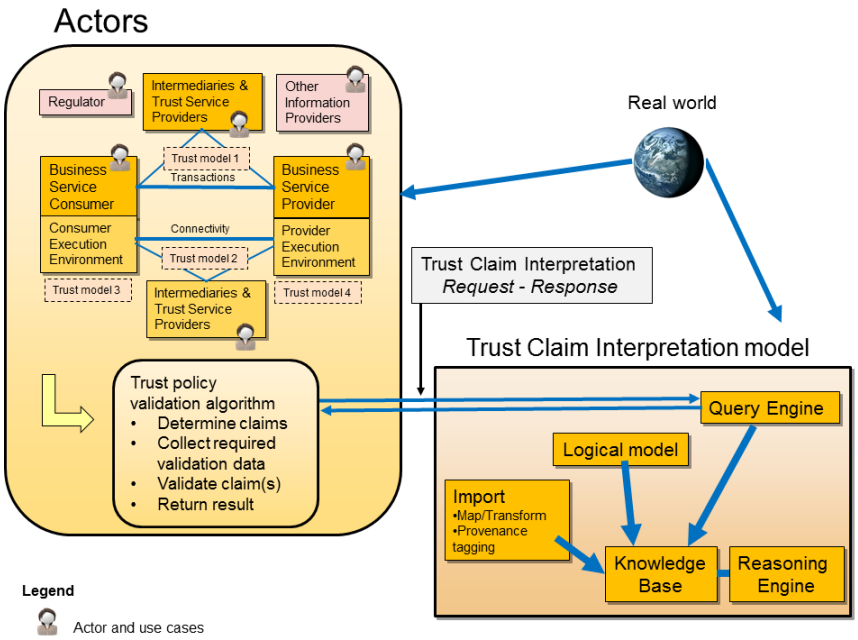

Fig. 1. Overview.

3.2 Actors and their use cases

Various actors are involved in the model. A Business Service Producer offers electronic business services, which are consumed by Business Service Consumers. Transactions and connectivity between producers and consumers can be protected by services of Trust Service Providers (TSP). The term TSP refers to a broad category of Service Providers including Identity Providers, Certification Authorities, Signature and Validation Service Providers, Time Stamping Authority services, registered electronic delivery services and trusted electronic archiving services. Other information providers offer additional information. They may be independent from the entities

they provide information about in varying degrees. A regulator can impose conditions on entities that provide services.

3.3 Trust models, trust policy and validation

A trust model combines safeguards such as cryptographic protocols, a policy, operational procedures and legal or contractual liability. There might be different trust models deployed at the level of the transactions, the connectivity and the integrity of the operational environment. A trust model is formalised in a trust policy, composed of a set of assumptions which contain a set of claims that need to be satisfied. Reliance can be validated according to a trust policy. The applicable set of assumptions is function of the actor's use case and the protected artefact. When verification of all claims yields positive evidence that they are satisfied, the assumption is considered satisfied. When all assumptions are satisfied, the trust policy is considered satisfied for that artefact. A generic trust policy is proposed:

- Claimant assumptions, through which a claimant does claim an identity:
  - Claimant functional assumptions that claim the identity for the claimant, specifying the level of assurance required for the authentication thereof, and all supporting evidence related to cryptographic meta data, algorithms and transformation devices.
  - Claimant cryptographic assumptions that address the actual cryptographic validation.
- Claim assumptions, through which a claimant does claim a trust related functionality which is different from identity:
  - Claim functional assumptions that address the functionality of the claim (message authentication, electronic signature validation, multi-party signatures validation, etc), the commitment assumed by the claimant, the level of assurance on timing evidences, and if applicable, the type of legal effect sought by the claim,
  - Claim cryptographic assumptions that address the actual cryptographic validation.

The following trust policy validation algorithm is proposed:

1. Determine assumptions and claims in function of protected artefact.
2. Collect supporting validation data. This includes certificates and information on supporting validation data.
3. Validate claimant and claim assumptions. In this step, all applicable assumptions and claims are validated, for which a request/response model is proposed.
4. Return result. The algorithm ends by returning the result to the actor.

3.4 The trust claim interpretation (TCI) model

The TCI model contains a representation of the real world, derived from normative knowledge and factual assertions. The normative knowledge is derived from authoritative sources such as legislation, regulation, standards and related information. To promote interoperability, the terminology to describe the normative knowledge should be based on existing terms and vocabularies, extended where necessary. Factual assertions are derived from on-line information including published meta data. The import functionality will map and if necessary transform the input artefacts for inclusion in the knowledge base. The model contains a mechanism to maintain data consistency, and a query engine to respond to trust claim queries invoked from the use cases.

An instantiation of the approach should have a solid basis, particularly for its semantic aspects. Reasoning should be deterministic, decidable and computable in a reasonable time. There should be support for the different use cases and their claims. It should be possible to

integrate contextual information in varying formats in a relatively easy way, to include publicly available information. The instantiation of the approach in this article is limited to the TCI model. Instantiating the trust policy validation algorithm and the corresponding trust policies and their assumptions is identified as a further research area. The first requirement indicates the need for a mathematical basis, with a focus on logic. Boolean logic, reputation scoring, subjective logic and description logic were compared, and the latter [1] was selected as the basis for the prototype implementation.

4 A prototype implementation

The logical model was defined in the logic $\mathcal{SROIQ}$ and implemented in OWL DL[2]. For a treatment of OWL DL refer to [6] and [11]. OWL DL was used because it is the syntactic fragment of OWL that abides the syntactic restriction that OWL axioms can be read as $\mathcal{SROIQ}$ axioms for which the structural restrictions are satisfied. This means that once $\mathcal{SROIQ}$ constructors and axioms are identified, these are described in DL classes and properties. Protégé[3] was used as programming environment. A Knowledge Base is a combination of T-, R- and A-boxes. The T- and R-boxes resulted from the modelling. The A-boxes resulted from manual imports. The reasoning capability is provided by the Hermit reasoner [10], built into Protégé. The query engine consists of the DL and SPARQL query interfaces of Protégé. The normative terminology is based on the EU eIDAS and the US FICAM definitions, and the individuals are based on evidence captured from on-line sources. To promote interoperability of the model, existing ontologies were used where possible. The implementation approach is now described.

4.1 Four step implementation approach

Identification of concepts and classes In the first step, $\mathcal{SROIQ}$ concepts were identified from the eIDAS and FICAM literature and modelled as OWL DL classes. The first concept that emerged was an anchorpoint that oversees supervision and publishes metadata. Supervisors and trust anchors may have a legal basis in a particular jurisdiction, or may be based on less stringent concepts such as a membership agreement. The second set of relations that emerged from this analysis were those between a service provider and consumer, making use of trust services. Such a TSP is overseen by a supervisor. The supervisor can point to the TSP's meta-data from within his own meta-data. This allows services consumers that invoke TSP services to validate against official meta-data. A third set of relations emerged around registers and assurance assessors. Since TSPs provide trust services against renumeration, they are typically officially registered organisations that pay taxes. Assurance assessors review that TSPs meet the requirements imposed on them, and report on this to the relevant supervisory authority. The analysis resulted in eleven concepts, listed in table 1. The relations between them are not included in this table, but are modelled in the OWL DL model. They implement the description above.

Reuse of existing vocabularies In the second step, as the prototype model aims to be inter-operable with existing definitions, vocabularies were evaluated for potential reuse or extension. The W3C list of ontologies [4] was used as a starting point. The DCMI's *dcterms* vocabulary was

[2] http://www.w3.org/2012/pdf/REC-owl2-direct-semantics-20121211.pdf
[3] http://protege.stanford.edu
[4] http://www.w3.org/wiki/Lists_of_ontologies/

found to be the most relevant standard, complemented by the W3C's *Organization*, and *Registered Organization* vocabularies. The first column of table 1 lists the $\mathcal{SROIQ}$ concept name. The second column provides a description. The third column indicates the basis for the semantic class. For further refining the class definition, three alternatives are possible. Either an existing vocabulary offers a relevant class that can directly be used, an existing vocabulary offers a class that can refined by subclassing it, or no relevant classes from existing vocabularies could be identified. In the latter case, a new class needs to be defined. In the current prototype, this latter alternative was not used. Whatever alternative is used, it yields the T-boxes.

$\mathcal{SROIQ}$ **concepts**		
$\mathcal{SROIQ}$ **conceptname**	**Description**	**Semantic implementation**
Jurisdiction	The extent or range of judicial, law enforcement, or other authority	Direct use of *dcterms: Jurisdiction*
LegalBasis	Legislation that provides authority	New subclass of *dcterms: BibliographicResource*
TrustMetaData-MR	Published meta-data about trust in machine readable format	New subclass of *dcterms: BibliographicResource*
TrustMetaData-HR	Published meta-data about trust in human readable format	New subclass of *dcterms: BibliographicResource*
TrustService	Service offering certificates, identity, authentication, time stamping, registered electronic delivery	New subclass of SKOS *skos:concept*
TrustAnchor	Formal organisation, mandated within some jurisdiction	New subclass of *org:FormalOrganization*
TrustSupervisor	Formal organisation, mandated within some jurisdiction	New subclass of *org:FormalOrganization*
TrustServiceProvider	Registered organisation providing trust services	New subclass of *regorg:RegisteredOrganization*
Register	Organisation that registers and makes available official information about other organisations	New subclass of *org:FormalOrganization*
TrustServiceAssuranceAssessor	Organisation that assesses the assurance level of a TSP service	New subclass of *org:FormalOrganization*
ContextualEvidenceProvider	Organisation that provides contextual evidence about an organisation or service	New subclass of *org:Organization*

Table 1. $\mathcal{SROIQ}$ concepts and their semantic interpretation

Roles and properties In the third step, $\mathcal{SROIQ}$ roles were defined and implemented as OWL DL properties. For classes based on existing vocabularies, existing object properties were reused as roles where possible, as well as existing data properties to capture relevant attributes. Otherwise, new definitions were created. This yields the R-boxes.

Individuals In the fourth step, individuals were created for the different classes of the model, yielding the A-boxes. In the current prototype this has been done manually. However it has been shown [12] that this can be automated using e.g. XSLT transformations.

4.2 Illustration of the four steps

The implementation of the *Jurisdiction* class illustrates the direct reuse of existing terminology. It is derived from the eIDAS and FICAM literature there is a need for such a concept, since claims will only be valid within a certain jurisdiction. The DCMI's class *dcterms:jurisdiction* is used directly in the TCI model. Then existing object properties such as *dcterms:coverage* are analysed. To capture the relation between a formal organisation and a jurisdiction, the new object property *hasJurisdiction* is introduced, with domain 'FormalOrganization' and range 'Jurisdiction'. To conclude, two individuals were created, EU Jurisdiction and US Jurisdiction.

The implementation of the *TrustServiceProvider* class illustrates the reuse of existing terminology by subclassing. It is derived there is a need for such a concept, since TSPs are used by both providers and consumers of business services. Both the European and the US regulations define a TSP. The W3C's *org* vocabulary is identified to contain the class *org:FormalOrganization*, and the *regorg* vocabulary contains the class *regorg:RegisteredOrganization*. *TrustAnchor* is subclassed of the latter. Then *isSupervisedOrCertifiedBy* and *providesTS* are created as additional roles. The data property *regorg:legalName* is reused. To conclude, TSP individuals are created.

4.3 Generating responses to requests

Once the KB contains T-, R- and A-boxes, and has been classified, queries can be answered. The present prototype implements elements of the validation of claimant functional assumptions for TSPs, TrustSupervisors and TrustAnchors. Generating responses to requests that result from invoking a trust validation policy is specified as illustration. In this case, claimant functional assumptions need to be validated that address the involved TSP and the TrustAnchor. Assumptions on TSP existence can be verified by the DL query *'TrustServiceProvider and registration some and providesTS some'*. This query yields the set of TSP individuals with these properties. Assumptions on TSP meta data and qualifications can be verified by the DL query *'TrustServiceProvider and isSupervisedOrCertifiedBy some and (publishesTMD-HR some or publishesTMD-MR some)'*. Assumptions on the legal basis of a trust supervisor can be verified by *'TrustSupervisor and hasLegalBasis some'*. The response to this query allows the distinction between a trust supervisor operating established on a legal basis (e.g. a national trust supervisor of one of the European countries) and a trust supervisor operating according to a less formal Membership agreement (e.g. the Kantara Initiative).

5 Strengths, weaknesses and further research

Analysing the proposed approach leads to the identificaton of the following strenghts. As the logical model is based on legislation and standards rather than on technical vocabulary only, it allows an interpretation that spans these two domains. As it builds on existing vocabularies from W3C and Dublin Core it allows interoperability, since rather than reinventing the wheel it starts from a terminology that has been created through large scale consensus. As it has a formal logic basis, composed of $\mathcal{SROIQ}$ and OWL DL, it ensure that interpretations conform to the logical definitions. It introduces transparency by allowing invocation of the explanation of the DL classification and inferences. It allows also the inclusion of contextual assertions from sources that are reasonably independent from the actor providing the trust claim.

The current concepts do not address securing the provenance of the various assertions in the knowledge base, as well as their timeliness. Also the formalisation of the degree of independence of providers of contextual assertions from the actors providing the claims is not addressed. The prototype implementation is limited to the TCI model and does not implement the trust policy

validation algorithm. The TCI request-response mechanism is currently only simulated by the query interface and does not support an http-like request-response protocol. It is further limited by the fact that individual assertions need to be entered manually.

Areas for further research include instantiating the trust policy validation algorithm, and its deployment in trust policies. Further areas include addressing the weaknesses identified in the preceding section. Securing the provenance and timeliness of the various assertions in the knowledge base, as well as the degree of independence of providers of contextual assertions related to the actors should be more formally addressed. Specialisations towards trust for IdPs and authentication, and towards trust for other TSPs can also be envisaged.

References

1. Franz Baader. *The description logic handbook: theory, implementation, and applications*. Cambridge university press, 2003.
2. Piotr Cofta. *Trust, Complexity and Control: Confidence in a Convergent World*. Wiley, 2007.
3. Andreas Fuchs, Sigrid Gürgens, and Carsten Rudolph. A formal notion of trust–enabling reasoning about security properties. In *Trust Management IV*, pages 200–215. Springer, 2010.
4. D. Gambetta. Can we trust trust? In D. Gambetta, editor, *Trust: Making and Breaking Cooperative Relations*, pages 213–237. Basil Blackwell, Oxford, 1988.
5. Olaf Hartig. Querying Trust in RDF data with tSPARQL. In Lora Aroyo, Paolo Traverso, Fabio Ciravegna, Philipp Cimiano, Tom Heath, Eero Hyvönen, Riichiro Mizoguchi, Eyal Oren, Marta Sabou, and Elena Simperl, editors, *The Semantic Web: Research and Applications*, volume 5554 of *Lecture Notes in Computer Science*, pages 5–20. Springer Berlin Heidelberg, 2009.
6. Pascal Hitzler, Markus Krötzsch, and Sebastian Rudolph. *Foundations of Semantic Web Technologies*. Chapman & Hall/CRC, 2009.
7. Jingwei Huang and David Nicol. A Calculus of Trust and its Application to PKI and identity management. ACM, 2009.
8. Audun Jøsang, Ross Hayward, and Simon Pope. Trust network analysis with subjective logic. In *Proceedings of the 29th Australasian Computer Science Conference - Volume 48*, ACSC '06, pages 85–94, Darlinghurst, Australia, Australia, 2006. Australian Computer Society, Inc.
9. Stephen Paul Marsh. Formalising trust as a computational concept. 1994. PhD thesis.
10. Boris Motik, Rob Shearer, and Ian Horrocks. Optimized Reasoning in Description Logics using Hypertableaux. In Frank Pfenning, editor, *Proc. of the 21st Conference on Automated Deduction (CADE-21)*, volume 4603 of *LNAI*, pages 67–83, Bremen, Germany, July 17–20 2007. Springer.
11. Sebastian Rudolph. Foundations of description logics. In *Reasoning Web. Semantic Technologies for the Web of Data*, pages 76–136. Springer, 2011.
12. Marc Sel. Using the semantic web to generate trust indicators. In Sachar Paulus, Norbert Pohlman, and Helmut Reimer, editors, *Securing business processes*, pages 106–119. Vieweg+Tuebner, Springer Science+Business Media, 2014.
13. Marc Sel. A comparison of trust models. In Sachar Paulus, Norbert Pohlman, and Helmut Reimer, editors, *Securing business processes*, pages 206–215. Vieweg+Tuebner, Springer Science+Business Media, 2015.
14. Saeedeh Shekarpour and S. D. Katebi. Modeling and evaluation of trust with an extension in semantic web. *Web Semant.*, 8(1):26–36, March 2010.