NeOn: Lifecycle Support for Networked Ontologies

Integrated Project (IST-2005-027595)

Priority: IST-2004-2.4.7 — "Semantic-based knowledge and content systems"

D4.3.1 Review of Trust Models as a Criterion for Ontology Customization

Deliverable Co-ordinator: **Alexander Kubias**

Deliverable Co-ordinating Institution: **University of Koblenz-Landau (UKO-LD)**

Other Authors: **Marko Babic (UKARL), Holger Lewen (UKARL); Martin Dzbor (OU); Klaas Dellschaft (UKO-LD); Jose Manuel Gómez-Pérez (iSOCO)**

In this deliverable we survey the state of the art regarding trust, trust models and trust systems for Semantic Web applications. Furthermore, the treatment of trust mechanisms in NeOn is described by presenting appropriate use cases.

Document Identifier:	NEON/2007/D4.3.1/v1.0	Date due:	October 31, 2007
Class Deliverable:	NEON EU-IST-2005-027595	Submission date:	December 14, 2007
Project start date	March 1, 2006	Version:	v1.0
Project duration:	4 years	State:	Final
		Distribution:	Public

NeOn Consortium

This document is part of the NeOn research project funded by the IST Programme of the Commission of the European Communities by the grant number IST-2005-027595. The following partners are involved in the project:

Open University (OU) – Coordinator Knowledge Media Institute – KMi Berrill Building, Walton Hall Milton Keynes, MK7 6AA United Kingdom Contact person: Martin Dzbor, Enrico Motta E-mail address: {m.dzbor, e.motta}@open.ac.uk	**Universität Karlsruhe – TH (UKARL)** Institut für Angewandte Informatik und Formale Beschreibungsverfahren – AIFB D-76128 Karlsruhe Germany Contact person: Peter Haase E-mail address: pha@aifb.uni-karlsruhe.de
Universidad Politécnica de Madrid (UPM) Campus de Montegancedo 28660 Boadilla del Monte Spain Contact person: Asunción Gómez Pérez E-mail address: asun@fi.ump.es	**Software AG (SAG)** Uhlandstrasse 12 64297 Darmstadt Germany Contact person: Walter Waterfeld E-mail address: walter.waterfeld@softwareag.com
Intelligent Software Components S.A. (ISOCO) Calle de Pedro de Valdivia 10 28006 Madrid Spain Contact person: Jesús Contreras E-mail adress: jcontreras@isoco.com	**Institut 'Jožef Stefan' (JSI)** Jamova 39 SL–1000 Ljubljana Slovenia Contact person: Marko Grobelnik E-mail address: marko.grobelnik@ijs.si
Institut National de Recherche en Informatique et en Automatique (INRIA) ZIRST – 665 avenue de l'Europe Montbonnot Saint Martin 38334 Saint-Ismier France Contact person: Jérôme Euzenat	**University of Sheffield (USFD)** Dept. of Computer Science Regent Court 211 Portobello street S14DP Sheffield United Kingdom Contact person: Hamish Cunningham
Universität Kolenz-Landau (UKO-LD) Universitätsstrasse 1 56070 Koblenz Germany Contact person: Steffen Staab E-mail address: staab@uni-koblenz.de	**Consiglio Nazionale delle Ricerche (CNR)** Institute of cognitive sciences and technologies Via S. Marino della Battaglia 44 – 00185 Roma-Lazio Italy Contact person: Aldo Gangemi E-mail address: aldo.gangemi@istc.cnr.it
Ontoprise GmbH. (ONTO) Amalienbadstr. 36 (Raumfabrik 29) 76227 Karlsruhe Germany Contact person: Jürgen Angele E-mail address: angele@ontoprise.de	**Food and Agriculture Organization of the United Nations (FAO)** Viale delle Terme di Caracalla 00100 Rome Italy Contact person: Marta Iglesias E-mail address: marta.iglesias@fao.org
Atos Origin S.A. (ATOS) Calle de Albarracín, 25 28037 Madrid Spain Contact person: Tomás Pariente Lobo E-mail address: tomas.parientelobo@atosorigin.com	**Laboratorios KIN, S.A. (KIN)** C/Ciudad de Granada, 123 08018 Barcelona Spain Contact person: Antonio López E-mail address: alopez@kin.es

Work package participants

The following partners have taken an active part in the work leading to the elaboration of this document, even if they might not have directly contributed to the writing of this document or its parts:

- Intelligent Software Components S.A. (ISOCO)
- Universität Koblenz-Landau (UKO-LD)
- Universität Karlsruhe – TH (UKARL)
- The Open University (OU)

Change Log

Version	Date	Amended by	Changes
0.1	27-07-2007	Alexander Kubias	Outline and initial sections
0.2	30-07-2007	Alexander Kubias	Chapter about Trust Modeling extended
0.3	09-11-2007	Klaas Dellschaft	Extended the use cases, conclusion and the executive summary.
0.4	16-11-2007	Klaas Dellschaft	Prepared the deliverable for sending it to the QA.

Executive Summary

In this deliverable, we survey the state of the art of trust systems and trust models. One can identify two major research directions: reputation-based trust and credential-based trust. In reputation-based trust systems, a more subjective viewpoint is taken on the topic. Reputation-based systems are for example used in the context of internet shops like eBay or Amazon where users rate each other or provide reviews of products. The reviews and ratings may then e.g. be used for deciding whether a specific seller is trustworthy enough to buy products from him.

In contrast, in credential-based systems a more objective and stricter viewpoint is taken. They may for example be used for restricting the access to resources in a system so that these are only available to people with appropriate credentials. It is the task of the owner of a resource to define which other users are trustworthy enough to give them the necessary credentials for accessing a resource. This also includes mechanisms for clearly identifying and authenticating users. Thus, credential-based systems also include authentication mechanisms like public key infrastructures.

We will start this deliverable in section 1 with a general overview on the most important dimensions for classifying and distinguishing different trust systems. It is followed by a survey of existing reputation-based and credential-based systems in section 2 and 3, respectively. In section 4, we will summarize the differences between the two approaches.

Then, in section 5 and 6 we will relate the two previously identified categories of trust to the use cases and already ongoing work in NeOn. In the use cases, reputation-based systems may be used for providing personalized views on ontologies by e.g. hiding ontology elements that are below a certain trust value. In the same way, it may also be used for (semi-)automatically resolving inconsistencies in merged knowledge bases. Credential-based systems may be used in the context of exchanging e-invoices,e.g. for restricting access to financial data. In the semantic nomenclature use case of NeOn, a reputation-based trust system may be used for rating the sources of knowledge in a database of pharmaceutical products while a credential-based trust system may be used for restricting the access to private parts of the BOTplus ontology to members of the General Spanish Council of Pharmacists while other parts may be accessed by a wider public.

There is already ongoing work in NeOn where frameworks for handling reputation- and credential-based trust are developed. With regard to reputation-based trust there exists a proposal for Open Rating Systems that will be further developed in context of WP2 (see [SAd+07] for a first description) while the framework for treating access rights in ontologies (see [DKG+07]) is a credential-based trust system that will be further developed in context of WP4.

Contents

List of Figures

Chapter 1

Trust Modeling

Currently, there is a remarkable confusion caused by the variety of terms describing trust-related systems. Moreover, there is also a lack of coherence in the terminology, since authors often propose new systems from scratch.

To bring order into the field, this deliverable describes proposals and developments using a consistent terminology. With reference to the most relevant literature the required terms are explained and differentiated from each other (section 1.2). Based on this terminology, the most relevant properties as well as a rough classification of trust are introduced (sections 1.3 and 1.4). After that, a more fine-grained classification of trust systems is given regarding the implementation of trust in artificial agents. With reference to Semantic Web applications there exist two different major approaches for managing trust and trust systems. These approaches are either based on policies and or on reputation. Hence, the concept of trust relevant for the both approaches is considered in this work. Other approaches like social trust are not in the focus of most researchers. These less relevant trust concepts are not described in this deliverable, but they can be found in [Mar94].

In the following, some general aspects of trust modeling (e.g. metrics, model semantics etc.) which can be applied to both, reputation-based as well as policies-based trust systems, are explained before being defined in more detail in chapters 2 and 3. These joint aspects for both approaches of trust management are described in sections 1.1, 1.2, and 1.3. Subsequent to these descriptions a top-down classification of trust-related mechanisms is provided. Namely, starting with a classification of the general concept of trust, gradually, more fine-grained classifications of integral parts of a trust system are provided. Figure 1.1 shows the integral parts of such a multi-level trust system.

Thereby, some exemplary names of different systems, models, and metrics are given which all will be explained and discussed in chapters 2 and 3. At this point, the concrete examples of two different trust systems, namely a reputation-based and a credential-based system, each containing two different models which again are composed of different metrics (in case of the reputation-based system) or certificates (in case of the credential-based system) are given in order to point out the multilayer structure of such a trust system.

Having provided some general classifications for the concept of trust in section 1.4, a classification of relevant trust systems is provided in section 1.5. The multilayer set-up of a trust system is also used as a guideline for the subsequent structure of this chapter.

Since each trust system can be composed of different trust models, the next section 1.6 deals with characterization and classification of appropriate trust models. That followed, each of the models can integrate different metrics to evaluate trust values. Therefore, section 1.6.1 deals with characterization and classification of relevant metrics.

At the very end of this chapter, some aspects relevant for implementation of such a multi-level trust system, e.g. trust semantics, are discussed.

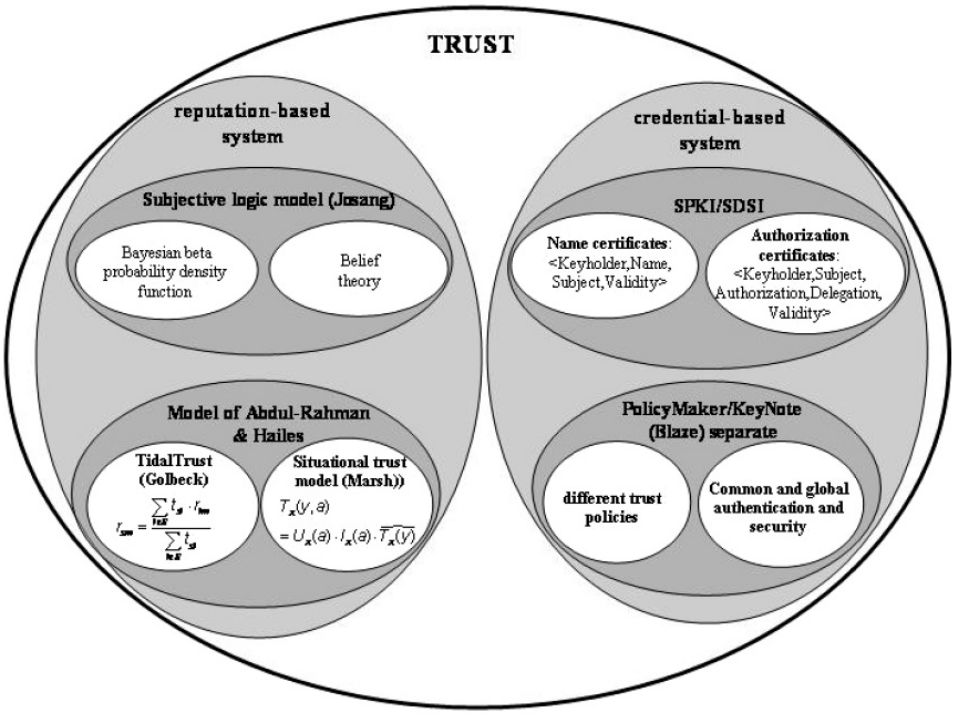

Figure 1.1: Multilayer Structure of Trust Systems

1.1 General Aspects

Trust is crucial as soon as risk, uncertainty, or interdependence exist [McK01, MDS95]. For instance, trust can be useful for deciding if a certain download from a server is risky or not. If the server is recommended by a friend who is trusted, there should be a lower risk that the downloaded file includes a virus. Trust can also be helpful in order to weaken uncertainty. Furthermore, trust can be used to dissolve interdependence.

Correspondingly, there is a large variety of trust related literature, ranging from specific applications to general models. However, this variety causes conceptual confusion on trust and makes comparing one trust study to another problematic. Hence, similar to this survey, there are several papers in social sciences summarizing an interpretation of existing research on trust.

One frequently cited work is [McK96] which integrates existing work in social sciences by proposing two kinds of trust topologies:

- classification of four qualities of trust (competence, benevolence, integrity, and predictability)
- definitions of six related trust types that form a model.

Alternatively cited work is [Gef02] which reduces the trust decision to three of these qualities, leaving out predictability. [Acr02, MMH02] are other prominent social science works about trust from a business management perspective. With respect to the computer science scope of this survey, [Mar94] is one of the first prominent, formal, computational models of trust. [Mar94] combines a subjective set of variables to derive a continuous trust value in the range [-1, 1], ranging from complete distrust to full trust. In addition three types of trust are identified: basic (over all contexts), general (between two people and all their contexts occurring together), and situational (between two people in a specific context). Beyond this, time is identified as being relevant to each of the variables used to comprise trust.

Following [Mar94] many researchers have endeavored to model, refine, and explain properties of trust in a computational setting. Before going into the details of those studies, some common terms and definitions as well as descriptions of trust properties are introduced.

1.2 Terms and Definitions

The term “trust” is used in literature with a variety of meanings [McK96]. According to [Jos05], a distinction between context independent trust (defined as “reliability trust”) and context dependent trust (defined as “decision trust”) can be recognized in the literature, although usually not explicitly expressed in those terms.

- **Reliability Trust** refers to past encounters and can be interpreted as the reliability of something or somebody independently of the context. [Jos05] and [Aga07] provide examples of how trust can be formulated on the basis of Gambetta [Gam88]: “Trust is the subjective probability by which an individual, A, expects that another individual, B, performs a given action on which its welfare depends” and Mui et al. [MMH02]: “[Trust is] a subjective expectation an agent has about another’s future behavior based on the history of their encounters”.

- **Decision Trust** is seen within a context and is unique in referring to the “competence” to act instead of actions themselves. Appropriate definitions are given in [Aga07]: “[Trust is] the firm belief in the competence of an entity to act dependably, securely, and reliably within a specified context” or [Jos05]: “Trust is the extent to which a given party is willing to depend on something or somebody in a given situation with a feeling of relative security, even though negative consequences are possible”.

- **Alternative Definitions** Additionally, there are also trust definitions combining some aspects of reliability and decision trust. The definition from Olmedilla et al. [ORMN05] is context dependent but refers to actions and not competence: “Trust of a party A to a party B for a service X is the measurable belief of A in that B behaves dependably for a specified period within a specified context (in relation to service X)”.

1.3 Properties of Trust

Merging properties assigned to trust in [RKM06] and [Gol05], the most relevant aspects of trust are covered when transferring the concepts to computer science. The following properties are regularly assigned to trust.

- **Subjectivity** Trust is based on personal experience with the interaction partner in the context of concern, on his reputation, or on recommendations. Furthermore, trust is connected to the presence of a subjective notion of uncertainty and depends on the personally expected risk associated with an interaction [RKM06].

- **Asymmetry** Since trust is subjective, it is asymmetric too. This property is of a significant importance when modeling trust in very anonymous settings like the World Wide Web. It is conceivable to trust a person to do his job well, the same person might not even know the person assigning the trust. In the real world many trust relationships are in fact mutual, but even there people trust each other to different extents.

- **Transitivity** Transitivity is a sensitive, controversial property most existing models assign to trust [GKRT04], [Guh03], [Gol05]. The basic idea of transitive trust is the following: Assuming Alice trusts Bob and Bob trusts Gil, what can be said about Alice’s trust in Gil? The question on how much Alice should trust Gil in comparison to someone who she trusts directly is an important aspect of modeling. [Mar94] points out that trust is not transitive over arbitrarily long chains, since if the chain becomes too long, people do not really know each other anymore.

  The propagation of distrust poses an interesting problem since distrust is not necessarily transitive [Guh03]. If Alice does not trust Bob and Bob does not trust Gil, it is not clear weather Alice should trust Gil or not. Among others [GKRT04] presents and evaluate different possibilities of modeling distrust.

- **Composability** The basic idea of composability of trust can be derived studying human behavior. Composing the same recommendations from different trusted sources, the trust assigned to that recommendation will be higher. A more difficult situation occurs when trust and distrust values shall be accumulated into one single trust value. Different ways to handle this problem are summarized in [Guh03].

- **Personalization** According to [FPHKH00] also made prominent is the idea that people trust people, not technology, which itself earns (or loses) our trust as an extension of trust in people. However, depending on subjective understanding of trust, the same person will be trusted by some and mistrusted by others. For this, personalization is a very important aspect of trust in computer science.

- **Dynamic** In [FC04] a key idea is dealing with the dynamic nature of trust making the realization that an agent that knows he is trusted may act differently from one who does not know his level of trust.

  Looking at another aspect of trust dynamics, [JSTT04] reports on human experiments showing how positive and negative experiences can change negative and positive trust, respectively. Key results from this work suggest that trust changes with different experiences, and that distrust may be harder to overcome than one would expect.

1.4 Classification of Trust

1.4.1 Principle Categories of Trust

According to [McK96] there are three principle categories of trust:

- **Personal** / **interpersonal trust** describes trust between people or groups and is closely related to the mutual experiences of the acting people.

- **Impersonal** / **structural trust** is not bound to a person but raises from social or organizational situation.

- **Dispositional trust** can be described as a person's general attitude toward the world.

These three categories can also be recognized in other works like in Marsh's model [Mar94]. Although Marsh uses in his work the concepts general, situational, and dispositional trust, the three principal categories stay the same. In this case the general trust corresponds to personal / interpersonal trust and situational trust to structural trust. In [AR04] it is shown that much work is done on transferring interpersonal trust to computer science, whereas there is little work supporting the other categories. To that effect this deliverable gives a survey over existing literature with the focus on modeling interpersonal trust. Nevertheless, in the context of NeOn all three categories of trust could be useful and should be used for determining the quality of an ontology.

1.4.2 Dimensions of Trust in Computer Science

In [Aga07] several dimensions are identified along which trust in computer science can be described:

- **Target**: The target specifies the entity whose trustworthiness is to be evaluated. It can be distinguished whether a user, a network or a service is the target of the trust evaluation. On the Web, one can either trust the agents, which provide the content, or the content itself.

- **Representation**: This deals with the representation of trust. Trust can be digitally encoded in many different ways. The two common ways of determining trust, credentials and reputation, are described

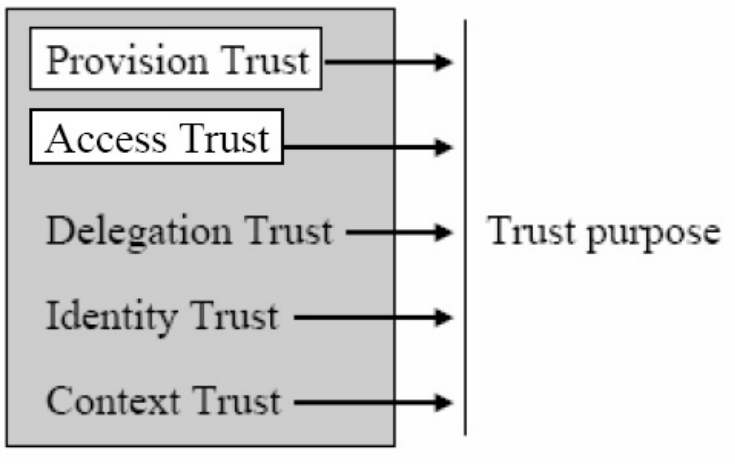

Figure 1.2: Trust Classes according to [Jos07]

in the following in detail (see section 1.5 and chapters 2 and 3). Regarding representation, credential-based trust systems use digital signatures and tokens for representing trust. In reputation-based trust systems, trust can be represented in the history of past interactions with other agents or users. For that a social network or a web of trust is used in order to determine trust in an unknown entity.

- **Method**: Trust can be determined through different methods. Corresponding to the previous paragraph, on the one hand there are the credential-based methods, which exchange credentials to establish trust before a transaction, and on the other hand there are the reputation-based methods, which use their history of past interactions or other entities' histories to determine trust through reputation.
- **Management**: The management dimensions of a trust system describe if the system is centrally controlled by a trust third party (as in traditional systems) or if the system uses a decentralization of control for the trust decision.
- **Computation**: Trust can be quantified and computed in many ways. Some approaches use discrete trust values, while others choose a continuous numerical range. Algorithms for how trust is transferred, combined or resolved can range from a simple average to computing eigenvalues. Many approaches compute trust assuming that time is static, while in other approaches trust may change over time. In cases where trust information is large or always changing, several approaches argue for a local computation of trust, instead of a globally consistent value. That means, that each entity computes its own trust values locally. Thus, each entity can have different trust values for the same target entity and a globally consistent value cannot be available.
- **Purpose**: The need for trust spans all aspects of computer science and each aspect places different requirements on trust. For instance, trust can be used to protect data, to find accurate information or to get the best quality service.

1.4.3 Trust Classes

With respect to more detailed trust semantics, [Jos07] introduces on the basis of [GS00] different trust classes (see Figure 1.2). For information, [GS00] use the terms service provision trust, resource access trust, delegation trust, certification trust, and infrastructure trust.

The highlighting of provision trust in Figure 1.2 is done to illustrate the focus in this study. Namely, in view of the following description of the single trust classes, the focus of this survey is on reputation-based and credential-based trust which correspond to a great extent to provision and access trust.

- **Provision trust** specifies the relying party's trust in a service or resource provider. According to chapter 2, this class is realized in reputation-based trust systems and is relevant whenever the relying party is a user seeking protection from malicious or unreliable service providers (confer "soft security" in section 1.5). [Jos07] shows that also similar concepts like "business trust" [Boe03] in the Liberty Alliance Project can be interpreted as provision trust. Although "business trust" describes mutual trust between companies emerging from contract agreements that regulate interactions between them, when for example a contract specifies quality requirements for the delivery of services, then this business trust would correspond to provision trust.
- **Access trust** describes trust in principals for the purpose of accessing resources owned by the relying party. This relates to the authentication and access control of traditional security mechanisms (confer the "hard security" concept in [RJ96]). [GS00] provides a good overview over access trust systems which are discussed in chapter 3 in more detail.
- **Delegation trust** describes trust in a delegate (agent) making decisions on behalf of the relying party. [GS00] points out that acting on one's behalf can be considered to a special form of provision trust.
- **Identity trust** describes the belief that an agent's identity corresponds to the claimed. The implementation of identity trust can typically be found in authentication schemes such as X.509 and PGP [Zim95].
- **Context trust** describes to which extent the relying party believes that the necessary systems are in place to support the transaction and provide a safety net in case of something going wrong. Applications for this type of trust can be found in critical infrastructures, legal systems, insurances, and stability of society in general.

According to [Jos07], conceptually, identity trust and provision trust can be seen as two layers on top of each other, where provision trust normally can not exist without identity trust. If so, it is only possible to have a baseline provision trust in an agent.

1.5 Classification of Trust Systems

There are fundamental differences between traditional and online environments when using trust for decision making purposes. Among all [Jos05] identifies the following two differences as most important. First, there exist several properties of trust which are only available in the physical world (i.e. the traditional environment of trust) and which are missing in online environments. For those traditional trust indicators electronic substitutes are needed. Secondly, sharing information related to trust in the physical world is relatively difficult and usually constrained to local communities. In contrast, IT systems supported by the Internet can be leveraged and designed for collecting and exchanging information on a global scale. In order to build good trust systems [Jos05] therefore suggests:

- To find adequate online substitutes for the traditional leads of trust that we are used to in the physical world, and identify new information elements which are suitable for deriving measures of trust in online applications. For instance, in the traditional environment trust can be established if we know someone for years. In the online environment trust can be established if several past transactions (e.g. in a online shop system) were successful.
- To take the advantage of IT and the Internet for collecting information globally, and to derive measures of trust in order to support decision making. The technical principles for building trust systems are reflected in the network architectures which two main types are centralized and distributed architectures.

The system architectures being introduced in section 1.5.1 are basis for different ways of determining trust. Two common ways of determining trust, the reputation-based and the credential based trust systems are

shortly introduced in sections 1.5.2 and 1.5.3 respectively, before being discussed more detailed in chapters 2 and 3.

1.5.1 System Architectures

In the literature, there exist two fundamentally different system architectures for trust:

- In **centralized trust systems**, information about a given participant is collected by a central authority from other members in the community who have had direct experience with the participant. The central authority typically derives a trust score for every participant and makes all scores publicly available. The other members of the community can then use each other's scores, e.g. when deciding if or not to interact with a particular member.
- **Distributed reputation systems** act without any centralized functions for calculating trust scores of other community members. Instead, there can be distributed stores where each member simply records the opinion about each experience with other community members and provides this information on request from relying party. Before making decision on a transaction with a given target party, a relying party collects opinions from as many distributed members as possible who have had direct experience with that target party.

1.5.2 Reputation-based Trust

Decentralized system architecture is the basis for a reputation-based trust system. The scores about trust are collected in terms of estimated reputation of a given party based on ratings of all the other community members. In this context, a reputation is an assessment based on the history of interactions with or observations of an entity, either directly with the evaluator (personal experience) or as reported by others (recommendations) [Aga07]. Since recommendations are trust decisions made by other users in a social network, additional work is to be done for computing personalized trust of the relying party. Appropriate models for computing personalized trust as well as models for combining past interactions or performance for an entity to assess its future behavior are presented in chapter 2.

1.5.3 Credential-based Trust

Contrary to the estimation of trust used in reputation systems, the "hard evidence" used in policies describes the conditions necessary to obtain trust, but can also prescribe actions and outcomes if certain conditions are met [BO05]. Policies usually involve the exchange of verification of credentials, which are information issued by an entity, for example describing qualities or features of another entity. Sometimes, credentials are endorsed using a digital signature.

An example shall point out the meaning of credentials. Identifying a "credential" of a university degree means that its holder has been recognized by the issuing university as having a specific education level. This way the holder is associated with the university and to those educated in the same field. Another party can use this credential when trust in the holder of the degree is unknown, but there is still existing trust in what is associated through the entity's credentials. In general, policy-based trust assumes that trust is established simply by obtaining a sufficient amount of credentials concerning a specific party, and is usually called credential-based trust.

1.6 Classification of Trust Models

To port trust from the physical world to the computer science world adequately it is necessary to provide adequate modeling of the traditional leads of trust and to identify new elements suitable for deriving measures

of trust in online applications (confer section 1.5). Those measures of trust are called “metrics” and will be introduced in the following, before some general aspects on modeling of traditional trust are depicted in section 1.6.1.

1.6.1 Classification of Trust Metrics

[ZL04] equates trust metrics in social networks with quantitative estimates of how much trust an agent A should accord to its counterpart B, taking into account trust ratings from other members in the network (community). There is a plethora of trust metrics, but few are confined to the Semantic Web. First proposals for trust metrics have been developed to support the Public Key Infrastructure (PKI) [Zim95]. New research fields apart from PKI, like P2P networks, ubiquitous, mobile computing and rating systems for online communities have raised the research interest in appropriate trust metrics.

[ZL04] characterizes all the available metrics along three classification axes with distinctive features. Since these axes are not orthogonal, various features impose restrictions on the feature range of other dimensions. The three principal dimensions are defined with reference to network perspective, computational locus (the place where the computation is executed), and link evaluation.

Regarding the network perspective, trust metrics may basically be subdivided into ones with local and ones with global scope. While global trust metrics take into account all peers and trust links connecting them in a network, trust metrics with local scope, on the other hand, take into account personal bias only. Since global trust metrics compute overall reputation rather than personalized trust [MMH02], some researchers claim that only local trust metrics are “true” ones.

The second axis refers to the computation locus, meaning to the place where trust relationships between individuals are evaluated and quantified. Centralized approaches perform all computations in one single machine and hence are granted full access to trust information. In a globally acting network, additionally to centralized metrics also distributed metrics for computation of trust can be deployed. The computation load is equally distributed on every trust load in the network, resulting in a decreased computation load with respect to centralized computation approaches. Receiving trust information from its predecessor nodes in the network, an agent merges the data with its own trust information and provides synthesized values to its successor nodes. For this, all the nodes in a network need to store trust information about any other node in the system.

Link evaluation axis distinguishes scalar and group trust metrics. According to [Lev03] scalar metrics analyze trust information independently, while group metrics analyze groups of assertions “in tandem”. Moreover, global group and local group trust metrics are distinguished.

Most of the trust metrics belong to the category of scalar ones. They track trust paths from sources to targets and do not perform parallel evaluation of groups of trust assertions. An example for global group metrics is given in PageRank [PBMW98], which, for the assessment of reputation of one page, comprises the ranks of referring pages, thus causing parallel evaluation of relevant nodes thanks to mutual dependencies. Examples for local group trust metrics are given by applications like Advogato4 [Lev03], which will be discussed more detailed in chapter 4.

[ZL05] compares different trust metrics, where the concept of local group trust metrics is advocated, as a compromise between local and global trust computation. The authors claim that trust is a “subjective expectation” and propose a method, Appleseed, for performing local group trust computation.

1.6.2 Implementation of Trust Models in Different Areas of Computer Science

[RKM06] gives a survey how the concept of trust is realized in different areas of computer science. Contrary to the equation of terminology trust modeling and trust management in the literature, [RKM06] distinguishes between the three categories trust modeling, trust management, and decision making.

In this classification, trust management focuses on the collection of evidence and risk evaluation only, and separates it from decision making, since the latter is such an important aspect.

In this terminology, trust modeling corresponds to the common research about trust in the literature. [RKM06] provides a more-fine grained classification for trust modeling, especially on the aspects domain, dimension, and semantics of trust values.

- The domain of trust values can be binary, discrete, or continuous, since trust values are usually expressed as number or labels. To express the two states of “trusted” and “untrusted”, a binary representation can be used, which is similar to certificate- or credential-based access control. To represent more than two discrete values, a set of labels or natural (discrete) numbers is used. Continuous trust values are calculated by means of mathematical theories depending on the semantics of the trust values.
- The dimension of trust values can be one- or multi-dimensional. In one-dimensional approaches the value describes the degree of trust a party assigns to another one, whereas multi-dimensional approaches allow to introduce a notion of trust uncertainty of the trust value.
- The semantics of trust values belong to the following set: rating, ranking, probability, belief or fuzzy value. Rating represents values which are directly linked with a trust related semantics, e.g. on a scale of natural numbers in the interval [1,4], where 1 can be linked to “very untrusted” and 4 to “very trusted”:
  - Ranking trust values are computed in ranking based models, e.g. [KSGM03], and are not directly associated with a meaningful semantics, but only in a relative way. Assuming that a higher values means higher trustworthiness, it is only possible to assign an absolute meaning to a value, if this value can be compared to a large enough set of trust values of other users.
  - Probability trust values calculated by probability models express the probability that an agent will behave expected.
  - Belief and fuzzy semantics are integral parts of the applications Subjective Logic [Jos01] and ReGreT [Sab03], and will be introduced in applications oriented chapter 4.
  - Additionally, [Jos07] describes semantics of trust values in terms of a specificity-generality dimension and a subjectivity-objectivity dimension. Subjective and specific trust values are for example used in survey questionnaires, where people are asked to express their opinion over a range of specific values. Subjective and general values are given in eBay’s reputation, where a member can express his subjective opinion in a general description. Objective measures are often applied in diverse product tests, either regarding a specific criterion, e.g. noise level, or general criteria, e.g. a general score on consumer satisfaction. Using objective measures, the correctness of ratings can be verified by others, or automatically generated.

Chapter 2

Reputation-based Trust

Traditional security mechanisms typically protect resources from unwanted access by means of authorizing users. However, in many situations the problem is in fact reverse. In these cases, we have to protect ourselves from those who offer resources, causing a whole range of security challenges, which are not covered by traditional security mechanisms. For example, deceitful users can provide false or misleading information, and traditional mechanisms are not able to protect against this type of threat.

Contrariwise, reputation-based trust systems provide a social control mechanism (in [RJ96] called “soft security”) and with it can protect against this type of threat. This chapter explores work in reputation-based trust. First, some relevant terms and definitions for reputation are introduced. Following, principles of reputation-based systems are presented by means of selected models and metrics. Concluding, implementation of reputation-based trust is presented by means of relevant applications, this way providing a basis for discussions in chapter 4.

In the context of NeOn, reputation-based trust could be used in order to select only those ontologies out of a repository of ontologies which are of a high quality. One could easily look at the trust value of the organization who had developed the selected ontology. If the trust value is high, which means that a lot of other persons were satisfied with the work of this organization, the quality of the selected ontology should also be high.

2.1 Terms and Definitions

As already mentioned in chapter 1, reputation-based trust uses personal experience or the experience of others, possibly combined, to make a trust decision about an agent. To be more specific, according to [ARH97a] and [ARH97b] reputation can be defined as: “an expectation about an individual’s behavior based on information about or observations of its past behavior”.

In the World Wide Web, an individual usually has less information to determine the trustworthiness of others; their reputation is typically used to determine to which extent they can be trusted. Someone who is more reputed is considered to be more trustworthy. His reputation is built on feedback from those who have had direct interactions with him. Given a set of feedback, [ZY04] introduces the term of “trust functions” being used for inferring one’s trustworthiness. However, those trust functions correspond to the trust metrics in section 1.6.1, which are used to quantify estimates of trust values. However, [ZY04] provides two essential contributions for future research in reputation-based trust. On one hand, [ZY04] summarizes relevant literature to a conceptual framework for reputation-based trust; on the other hand, it gives a classification scheme for reputation- based trust functions. In the following, the work of [ZY04] is detailed.

2.1.1 A Conceptual Framework for Reputation-based Trust

[ZY04] proposes a framework for reputation-based trust, assuming that in a decentralized environment, several entities interact with other entities in transactions. A transaction must be unidirectional, meaning there

is a service-provider (server) and a service-consumer (client). In this framework the terms trustworthiness, feedback, opinion, and source and destination of trust evaluation are defined.

- **Trustworthiness** is an indicator of the quality of an entity's services. In most trust models, the domain of trustworthiness is assumed to be in the interval [0, 1]. These values can be discrete or continuous.
- **Feedback** is a client's statement about the quality of a service that was provided by a server in a single transaction. A feedback may be multi-dimensional, reflecting the client's evaluation on a variety of aspects of a service, like price, product quality, etc. For simplicity, [ZY04] assumes that feedback is one-dimensional with the domain range in [0, 1].
- **Opinion** is a client's general impression about a server. It is derived from client's feedbacks on all the transaction that are conducted with the server. Opinions differ from trustworthiness although opinions are also assumed to be one-dimensional with the domain range [0, 1].
- **Source and destination of trust evaluation** "If an entity A is interested in knowing the trustworthiness of another entity B, then A is the source and B is the destination of a trust evaluation".

The framework for reputation-based trust can be modeled as a directed multigraph G(V,E). V is the set of vertices and E is the set of labeled edges. There are two types of edges: transaction edges and opinion edges. A label contains either A's feedback on a transaction with B, or A's opinion on B's service.

2.1.2 Trust Functions and Metrics in Reputation-based Systems

In the following subsections two classification schemes for reputation-based trust metrics or functions are provided. Using one of the both classification schemes, the essential features of the most well-known models for reputation-based trust can be comprehended.

Classification According to Zhang et al. [ZY04] Besides the proposed framework (confer section 2.1.1), [ZY04] proposes a classification scheme for reputation-based trust functions. This classification scheme has four dimensions:

- subjective vs. objective trust
- complete vs. localized information
- transaction- vs. opinion-based trust
- rank- vs. threshold-based trust.

Subjective vs. objective: The first dimension corresponds to the in [Jos07] proposed classification of trust semantics (confer section 1.6.2). If the quality can be objectively measured, then an entity's trustworthiness is called objective trust. Otherwise an entity's trustworthiness is called subjective trust. If an objective trust function is used, an entity's trustworthiness is independent of the source of the trust evaluation. By using a subjective trust function, an entity's trust may vary greatly depending on the source of the trust evaluation.

Complete vs. localized: This dimension corresponds to local vs. global classification in [ZL04]. If every entity has access to all the transaction or opinion information, the trust function is called global trust function. In this case, the trust function owns the complete information. If the trust function is only applied to a subgraph of the complete trust graph (maybe to its neighbors), the trust function is called a local or localized trust function. The local trust function uses only localized information. For a local trust function, each entity has access to different information. Because of that, a local trust function is also subjective [ZL04].

Transaction- vs. opinion based: Some trust models rely on the information of individual transactions in order to infer an entity's trustworthiness, whereas other trust models only rely on opinions. Unlike opinion-based

trust functions, transaction-based trust functions require more information to infer an entity's trustworthiness. Opinion-based trust functions give each entity more freedom to form their own opinions, which causes that opinion-based trust functions may be easily influenced by malicious users.

Rank- vs. threshold based: This dimension deals with the trust decision. For some trust functions, it is appropriate to define a threshold of trustworthiness in order to make trust decisions. These functions are called threshold-based. Other trust functions use the relative ranking of an entity in order to make trust decisions. Thereby, an entity's calculated trustworthiness is compared to other entities giving a relative ranking of an entity (confer also the semantics of trust values in section 1.6.2).

Classification According to Maresch et al. [Mar05] Similar to [ZY04], in [Mar05] a trust metric calculates the trust between a source and a destination entity. Even though not that widespread in the literature like the previous classification according to Zhang, [Mar05] provides a basis for the comprehension of [BCGM05] and [BO04], which describe implementations of trust systems for the Semantic Web (for details see section 2.3.2).

For the classification of reputation-based systems five main dimensions are given:

- **General properties** of the system imply common properties as the place of the evaluation (centralized or distributed), or the interpretation of the inferred trust values.
- **Database** is used to store data of the system exclusively. The rating schema specifies the structure of the data, and the way in which the ratings are done, is determined by the rating process (see below). There are two properties of the database that are of special interest: Visibility can be open or closed and locality can be central or distributed. If the database is visible, i.e. open for all entities, the trust in the trust evaluation is increased. Regarding the locality feature, the database can be stored centrally or in distributed entities.
- **Rating scheme** organizes the structure of the data, on which the reputation scheme relies. The rating scheme defines the relationships between the entities and the form of information that the reputation-based system needs. Furthermore, the rating scheme can be distinguished according to three subclasses: the form of ratings, the properties of the graph and the identity of the rater.

  The form of ratings can be very different. A rating can be direct or indirect. Direct ratings consist of a concrete statement about the strength of a relationship to another entity. The statement can be either a number or a verbal phrase that is defined in the vocabulary of a trust ontology. Additionally, the trust statement can be context-aware or not. A context-aware statement holds only in the mentioned context. A serious problem of a direct rating is that it is not visible how the rating was created. The direct rating can be a personal opinion about another entity or a synthetic value (summarization of transaction feedbacks).

  Indirect ratings can either be based on feedbacks or on transactions. A feedback is a rating of the entity's behavior in a single transaction. Transaction-based information is recorded data of transactions. They are kept separately from feedbacks because they do not consist of a explicit rating, but express the existing of a transaction and its objective circumstances.

  The current dimension of the form of ratings is closely related to [ZY04]. While in [ZY04] transaction-based and opinion-based metrics are distinguished, in [Mar05] direct and indirect ratings are differentiated. Direct ratings are similar to opinions, but they are not restricted to opinions. They can also consist of a synthetic value. In [ZY04] transaction- based and feedback-based metrics are not distinguished because they are both subsumed under the term transaction-based. In [Mar05] indirect ratings can either be transactions or feedbacks.
- **Properties of the graph** The rating scheme also depends on the properties of the underlying graph in which entities are nodes and links between nodes contain trust values. There are two major types of

graphs in the literature: Closed graphs do not contain leaves, while social networks model the social relationships between individuals.

- **Identity of the rater** is the last important property of the rating scheme. The identity of the rater can be known, anonymous or the rater can use pseudonym In open systems it is very important that the identity of a rater can be checked. If the rater has a high trustworthiness himself, its ratings are likely to be very trustworthy. If pseudonyms are used within a system, they can also be exploited in order to check the associated rating. However, this is only possible if it is not so easy to change one's own pseudonym. If the rater rates anonymously, it is not possible to relate the ratings to raters. Thus, it is not possible to ensure the quality of ratings

- **The rating process** explains the actions and circumstances under which the ratings of entities are created. The rating process is responsible for controlling the rating schema and the identity of the raters. Thereby, the authorization of the raters and the quality of rating should be checked. Reputation-based systems can fully control the recently created content for the database if the identity of the rater, the rating scheme and the plausibility of the information are checked. Therefore, a trustworthy central entity is needed. Some systems need less control by only applying an access control. They only check the identity of the rater. The rater ensures the compliance with the rating scheme and the quality of the ratings. Finally, there are some systems without any access control. Everyone is free to express whatever he wants.

- **The rating algorithm** is the core of a reputation-based system. The rating algorithm can have a certain perspective, a selection procedure, and a certain evaluation of entities' relationships. "The Perspective" of a rating algorithm can either be global or local.

  A global trust metric considers all entities of a network. Its result is universally valid and objective. Thus, such global trust metrics are well suited for centralized environments, which possess the complete information about all entities.

  A local trust metric also regards the personal interests of all entities. The identity of the user is exploited as another parameter for the local trust metric. Hence, local trust metrics are subjective and depend on the source of the trust evaluation.

- **Selection procedure** Trust metrics should support the decision making of an application. Therefore, the algorithms use selection procedures that can be either rank-based or threshold-based. Rank-based procedures evaluate an entity relatively to their concurrent entities and are used in metrics that evaluate the relationships between entities in the whole group. Threshold-based procedures need a threshold value as an input parameter in order to distinguish if an entity is trustworthy or not. The results of a trust metric can correspond to the relationship between two entities or to the relationships within a group of entities.

- **Evaluation of entities' relationships** Group trust metrics evaluate relationships between entities within a group in parallel. If group trust metrics are used, the relationship to another group member is influenced by the relationships that the other group members have. Local trust metrics compute the trustworthiness for a subset of all entities and use the relationships within this subset for the computation.

  Scalar trust metrics, in contrast to group trust metrics, only evaluate the relationship between the source and the destination of the trust evaluation.

2.2 Fundamental Reputation-based Trust Models

In the following, items and classifications of metrics and semantics introduced in section 2.1.2 and chapter 1 are clarified by means of some selected trust models. These models provide a basis for implementations

of trust systems. This way, interdependences between different models are highlighted, although formal definitions of the models stem from different sources.

Few reputation-based trust systems are introduced in the following section 2.3, before being discussed and compared with credential-based trust systems in chapter 4.

2.2.1 Trust Model Proposed by Marsh [Mar94]

Marsh' work [Mar94] is the seminal work on trust in computer science. Direct trust between only two agents is modeled without the collection of recommendations provided by other agents. Marsh introduces the items knowledge, utility, importance, risk, and perceived competence, in order to model trust between two agents. His model answers three questions: "with whom should an agent cooperate, when, and to which extent?" The trust values are expressed as real numbers in the range [-1,..,1] with threshold based decision making. Marsh uses three kinds of trust:

- **Dispositional trust** is trust of an agent x without involvement in the situation and possible cooperation partner.
- **General trust** expresses the trust of x in y, but is still independent from the situation.
- **Situational trust** involves the situation a in order to describe the trust of agent x in agent y.

Risk, which describes the probability that an entity behaves in an improper and bad way, is calculated based on costs and benefits of the considered engagement. Perceived risk and the competence of the possible interaction partner determine the cooperation threshold. Having the situational trust above the value calculated for the cooperation threshold, cooperation takes place, otherwise not.

2.2.2 Trust Model Proposed by Golbeck [Gol05]

Golbeck [Gol05] proposes a reputation based trust model for social networks called TidalTrust. Golbeck approximates continuous trust values by ten discrete trust values in the interval [1..10], claiming that humans are better in rating on a discrete scale than on a continuous one. The model is evaluated in the social network FilmTrust, where the users have to rate movies [Gol06a]. Moreover, one can rate friends in terms of "if the person were to have rented a movie to watch, how likely it is that you would want to see that film" [Gol05].

2.2.3 Trust Model Proposed by Abdul-Rahman and Hailes [ARH00]

The trust model proposed by [ARH00] is a kind of merge between the two models discussed in sections 2.2.1 and 2.2.2, although the formal definition of trust is based on [Gam88]. The model comprehends the direct trust of an agent in another one based on direct experience and the recommender trust of an agent in the ability of another agent to provide good recommendations.

Alike [Gol05], [ARH00] also represents trust by discrete values. However, to express direct trust, the labeled trust levels "very trustworthy", "trustworthy", "untrustworthy", and "very untrustworthy" are used, as well as "very good", "good", "bad", and "very bad" for recommender trust.

The direct trust values are only used to calculate the semantic distance to other agents. The semantic distance corresponds to the difference between two labels in case that an agent A labels an agent C to be very "trustworthy" based on personal experience, and at the same time A knows that an agent B labels the same agent C to be "very trustworthy". The semantic distance value can then be used to adjust further recommendations of B. The recommender trust determines weights used for combination of recommendations by means of a weighted summation. The model drops recommendations of unknown agents for the calculation of the recommended trust value. However, providing their recommendations, those agents get known and are considered for future calculations.

For the case of unknown agents or insufficient experiences with other agents, uncertainty is introduced, but without a clear statement how to take benefit from this in computation process. Moreover, the model does not deal with risk and it is not explicitly described how to introduce recommendations of recommendations.

2.2.4 Mathematically More Sophisticated Trust Models

Mathematically more sophisticated trust models do not calculate the trust values by simple weighted summations, but usually involve aspects of probability theory. The "subjective logic" model from [Jos01] combines elements of Bayesian probability theory with belief theory.

Belief theory represents a possibility to deal with uncertainty. [Jos01] introduces belief theory to express opinions as a triple (b, d, u), where b models a human notion of belief, d the disbelief, and u the uncertainty about a specific statement. The three parameters are interrelated by the term $b + d + u = 1$.

Linguistically fuzzy concepts are another approach of modeling trust and reputation, where membership functions describe to what extent an agent can be assigned as trust- worthy or not trustworthy. Reasoning with fuzzy values of this type can be provided by fuzzy logic. Some examples for this approach are described in [Man98] as well as the REGRET reputation system described in [SS01], [SS02a], and [SS02b].

2.3 Implementation of Reputation-based Trust Systems

By using a model similar to those presented in section 2.2, arbitrary scientific, commercial, etc. reputation models can be developed and used in real-life contexts. Most of those systems compute trust or reputation by transitive iteration through looped or arbitrarily long chains and are often summarized by the term "flow models".

Some flow models assume a constant reputation weight for the entire community, and this value is divided among all the members of the community. In those systems, a participant can only increase his reputation at the cost of others. For instance, in section 1.6.1 mentioned Advogato's reputation scheme [Lev03] and Google's PageRank [PBMW98] belong to this category.

Other flow models, like EigenTrust model in [KSGM03], do not require a constant sum of the reputation scores. EigenTrust computes reputation scores through repeated and iterative multiplication and aggregation of scores along transitive chains until the reputation scores for all community members converge to stable values.

In chapter 4 a more detailed discussion about applications of reputation- and credential-based trust systems is provided. In preparation for this, subsequently, a rough classification of reputation-based systems into P2P networks, on the one hand, and Web oriented reputation systems, on the other hand, is given.

2.3.1 Reputation Systems in P2P Networks and Grids

Peer-to-Peer networks represent a target application of reputation-based trust in order to address the problems of data quality. In contrast to web navigation, in a P2P network, every user plays actively the role of both, client and server, at the same time. Since there are no barriers and requirements to publish a file in the network, anyone can publish anything with any lack of quality. Additionally, the availability and reliability of any given node in the network is not guaranteed, thus preventing reliable data transfer. To solve the problem, in some systems like Napster, searching for nodes where the requested resource resides is centralized on a resource directory server; in other pure P2P networks like Gnutella and Freenet also the searching is distributed. Intermediate architectures like iMesh and Grokster also exist.

Downloading an arbitrary named file from a malicious servant can easily bypass firewalls and be used for spreading e.g. viruses or Trojan horses. To overcome these security threats many authors have proposed reputation systems for P2P networks [AD01, CDdV+02, DdVP+02, KSGM03, Fah02, Lia03, GJA03].

On the basis of the PageRank algorithm in [PBMW98] for ranking Web sites by authority, the EigenTrust algorithm in [KSGM03] or the P2PRep system in [CDdV$^+$02] are derived. The global reputation value for each agent represents the quality of a peer's uploads, e.g. calculated by counting the number of successful uploads.

In contrast to [YS02], [AD01] claims a more scalable system, since other reputation- based systems require the maintenance of a growing performance history. Still using reputation information, this approach uses statistical analysis to characterize trust, this way reducing computation effort.

Next improvements of reputation-based systems in P2P networks are presented in [DdVP$^+$02] and [ORMN05]. With the XRep protocol, [DdVP$^+$02] provides a more robust method for reputation management, which allows an automatic vote using user's feedback for the best host for a given resource. [ORMN05] describes requirements supporting trust in P2P networks and argues that semantic representations can address the requirements outlined. Moreover, the limitations of existing work on trust in P2P networks are discussed. In addition, extensive studies on the most relevant problems of reputation-based systems, as well as solutions are given in [Jos07] for the following problems:

- Low incentive for providing rating
- Bias toward positive rating
- Unfair ratings
  - Endogenous discounting of unfair ratings
  - Exogenous discounting of unfair ratings
- Change of identities
- Quality variation over time
- Discrimination
- Ballot box stuffing

2.3.2 Reputation Systems in a Web of Trust

In Web environments, a trust decision is mostly modeled as a transitive process, where trusting one information source requires trusting another associated one. The majority of transitive trust computation has concentrated on using reputation. Reputation is defined as a measure of trust, whereas each entity records reputation information on other entities, thus creating a so called "web of trust". In the following two subsections trust concerns on the Web are discussed regarding the significant distinction between Semantic and hyperlink-based Web.

Trust Using Hyperlinks Transferring trust into web of trust is also key contribution of [Ste99] and [SZ03]. They describe a set of hypotheses and experiments for testing how trust is transferred between hyperlinks on the Web. Specifically, it is computed how much trust is transferred from a trusted Web resource to an unevaluated one in the context of a consumer trusting a business for purchasing a product. Calculations are performed considering different types of links, types of resources and types of trust in the known source.

A similar approach is realized with the tool TrustRank in [GGMP04]. With a given small data set of decisions made by users about weather or not few Web sites are spam, TrustRank uses the link structure to other pages to derive whether or not they are also spam.

Hitherto presented models assume in accordance with Google's PageRank [PBMW98] that all Web links are positive endorsements and indications of trust. Since the assumption does not always hold true, [MH05] propose a minor addition to HTML which enables the author to specify that a link has a positive, negative, or undefined value.

However, all the models have in common that they do not consider context and thus do not differentiate between "topic specific" and referral trust. Solely, [DKG+04] and [DZF03] provide a method of computing within a web of trust that also considers the domain of knowledge (context), and does so separately from referral trust. To use context-based trust more advanced systems specially designed for the Semantic Web are proposed.

Trust on the Semantic Web With the Semantic Web in mind, [BO04] makes several claims for the reputation-based trust:

- All statements in the Semantic Web are to be considered as claims rather than facts until trust can be established.
- It is too difficult to provide trust information that is current.
- A Semantic Web architecture should use all trust relevant information available instead of using trust ratings only. Combining different trust mechanisms, users can formulate subjective and task-specific trust policies. Especially the usage of context- and content-based trust mechanisms within Semantic Web is considered a promising path for future research:
  - Context-based trust refers to the circumstances and associations of the target of the trust decision. An example of a context is an entity providing a description for an item, where the entity may be a vendor selling the same item, or a customer recommending the item.
  - Content-based trust can be used by applying common sense rules to make a trust decision, e.g. generally not to trust prices below 50 percent of the average price.

To prove the advantage of the proposed concept for the Semantic Web, [BCGM05] provides a browser implementation which filters content based on a user specified policy. Written in the TriQL.P language, the policies allow specification of requirements with reference to the context, content, and source of information.

Also [DZF03] proposes a system, where agents use both context and reputation to decide what information to trust in the Semantic Web. In this work, referral trust is employed to collect reputation, but it relies on the Semantic Web features to determine context. As a result of this recommender system, it is possible to ask another agent for instance "which agent can I trust to get the weather forecast?". Those recommender systems are common on the Semantic Web, and help to filter information based on recommendations and trust rating. Another example considering the Semantic Web is given in [Zie04], where a "taxonomy" is used to evaluate the similarity between profiles of users' interests.

Key recent examples of modeling trust on the Semantic Web are given in [GH04] and [Gol04], which use ontologies to express trust and reputation information. In related work, [Gol06b] uses the Semantic Web and "provenance" to infer trust relations. The concept provenance refers to general details regarding the sources and origins of information to evaluate trust, e.g. author, citations, publisher, etc.

In [Gol06b], provenance establishes a relationship between people and information, and the Semantic Web contains social network data used to calculate trust between people. Another method for selection of information sources is presented in [DKF+05], where again provenance and computation over a web of trust are applied. Assuming a determined provenance, the method uses this information to determine more trusted sources, while taking into consideration the concept of ignorance, i.e. not having any information about trust.

Further key works like [RAD03] and [GKRT04] consider computing of trust transitivity for Web applications regarding distrust and system's robustness to noise, whereas [MA05] concentrates on methods to deal with controversial users, i.e. those who are both trusted and distrusted at the same time.

2.3.3 Application-specific Reputation Systems

Some reputation-based trust systems are specially designed for specific applications requiring unique ways to harness reputation.

An example of such a system is represented in [PM04] used for routing in ad-hoc networks, where some nodes may be more trustworthy for routing packets than others. To decide which nodes in a network to use for routing traffic, each node in the network indirectly monitors the performance of the other nodes nearby, and infers their trustworthiness for correct routing.

Other specific applications are [DRJ04] for allocating tasks to the best performing agent among several counterparts, or [Jos02] which combines reputation feedback data by means of a beta probability distribution in order to infer adherence to contracts in e-commerce. For similar trust-based decision making in e-commerce, [Jos99] shows how to use the concept of subjective logic (cf. section 2.2.4), and how the method can be integrated in policy based trust management [Jos06].

Chapter 3

Credential-based Trust

Contrary to reputation-based trust systems, the primary goal of credential-based trust systems is the same as in traditional security mechanisms, i.e. to protect resources from unauthorized access. Therefore the concept of credential-based trust management is limited to verifying credentials and controlling access to resources according to application defined policies. A resource-owner provides a requesting agent access to a restricted resource only if it can verify the credentials of the requesting agent either directly or through a web of trust.

Because of the enormous heterogeneity of the available information, information providers, and users on the World Wide Web, security becomes increasingly important. Security- related aspects are plentifully investigated in the literature. Those aspects are mostly classified in three categories: confidentiality, integrity, and availability [Bis03, Sam01].

Current access control is mostly based on identity-based "authentication", which means that the users must be known to the provider, for instance, by a previous registration. Since World Wide Web is an open, distributed, decentralized, dynamic and interoperable environment providing services that must be usable by anyone spontaneously and dynamically and users do not always wish to disclose their identities, security infrastructures that require registrations or any other central controlling components are not suitable. Therefore, instead of authentication-based access control rather "authorization"- based access control of Web services is proposed.

This chapter deals with authorization-based access controls as the supporting architecture for establishing credential-based trust systems for the Semantic Web. Introducing, relevant concepts and mechanisms of access control are presented and classified in section 3.1. Subsequently, a general approach of modeling credentials is presented in section 3.2, and finally, an overview on the implementation of credential-based trust systems for the Semantic Web is given.

In deliverable D4.4.1 a survey of access rights has been made. These access rights are similar and, thus, related to credential based trust. Further information about access rights and their use cased for NeOn can be found in D4.4.1.

3.1 General Concepts and Classifications

3.1.1 Classification of Access Control Mechanisms and Strategies

Merging the literature [Bis03, Eck04, Gol06c], three main kinds of access control can be distinguished:

- Discretionary Access Control (DAC) is defined by the TCSEC[1] as "a means of restricting access to objects based on the identity of subjects and/or groups to which they belong. The controls are discretionary in the sense that a subject with a certain access permission is capable of passing that

[1] Trusted Computer System Evaluation Criteria

permission (perhaps indirectly) on to any other subject (unless restrained by Mandatory Access Control)".

- Mandatory Access Control (MAC) also defined by the TCSEC as "a means of restricting access to objects based on the sensitivity (as represented by a label) of the information contained in the objects and the formal authorization (i.e. clearance) of subjects to access information of such sensitivity".

- Role-Based Access Control (RBAC) is the latest approach to restricting system access to authorized users. Among numerous others, [AS04] and [ASW04] also identify the role-based access control to be best suitable for Semantic Web services. The central idea of the Role-Based Access Control is that permissions are associated with roles, and users are assigned to appropriate roles whereby the management of permissions is greatly simplified. Moreover, users can be easily reassigned from one role to another, and roles can be granted new permissions as new applications and systems are incorporated, or permissions can be revoked from roles as needed.

Additionally to the above differentiation, with reference to different strategies of access control it is distinguished between:

- Identity-Based Access Control is currently the most used access control, also called authentication, where the users must be known to the provider, for instance via previous registration. In open and distributed environments like the Web, identity based access control is less appropriate, since there is no central instance to check the correctness of the identities. Moreover, the identities of the actors are often unknown. In such environments capability- or credential-based access control is more suitable.

- Credential-Based Access Control is based on authorization and especially appropriate for the open and distributed Semantic Web. In such systems, users prove their legitimacy for access by showing a set of credentials stating their capabilities, by means of a Web service provider verifies, whether the shown set of proven credentials satisfies all the required constraints. Authorization-based access control also includes authentications, but here the authentication is based on Public Keys and not on identities. However, establishing trust by using Public Keys as the basis of credential-management (e.g. managing credentials and credential chains as well as developing strategies for automated trust negotiation when interacting with other agents), requires a trustworthy Public Key Infrastructure (PKI), which provides services and cryptographic methods. For this, the essential properties of appropriate Public Key Infrastructures are explained in more details in section 3.1.2.

Additionally to this classification, [GS00] provides a more detailed overview on access trust systems.

3.1.2 Public Key Infrastructures (PKI)

A Public Key infrastructure constitutes the basis of credential-management since it provides services and cryptographic methods for trustworthy exchange of credentials. The trustworthiness of credential-based systems is discussed in detail in [LMW02, BFK99, HMM+00]. In summary it might be said that all well-known classifications of trust (confer chapter 2) can also be found in access control systems. Infrastructure trust refers to the trustworthy Public Key Infrastructure including hard- and software. Access trust describes the owner's attitude toward requesters, and provision trust describes exactly the opposite attitude. Certification trust implies the basis for a trustworthy relationship, namely the certifications and credentials, and delegation trust describes the possibility to delegate one's rights to authorized agents.

The purpose of the latter two classes of trust will become clearer, after the mechanisms of credential-based Public Key infrastructures and models are presented in the following sections. However, prior to this, for the sake of completeness, two common identity-based Public Key infrastructures are introduced shortly.

Identity-based PKI infrastructures To provide a trustworthy network environment, several identity-based Public Key infrastructures have been developed. The most famous ones are following:

- X.500 and X.509 were developed by the CCIT in 1988 to provide electronic directory services for globally distributed networks. By means of X.500, a worldwide distributed database with distinguished names for directories has been built. X.509 was developed to support authentication of directory entries. For more details refer to [Hou99].
- The Privacy Enhanced Mail (PEM) standard was developed in the 80ies with reference to protection of emails by applying cryptographic methods. The main goals were the guarantee of trustworthiness, integrity, and authenticity of emails by also applying X.509. However, this standard was never widely deployed or used.
- Pretty Good Privacy (PGP) was originally created by Zimmermann [Zim95]. PGP is a software tool that provides cryptographic privacy and authentication. Hence, there is no real hierarchical certification infrastructure; rather each user corresponds to her own certification authority. Everyone can digitally sign the public keys of other users and this way build up a network of certifications also known as a Web of trust.

Credential-based PKI infrastructures In the previous section, the concept of certifications was used, although this concept is in the literature widely equated with the concept of credentials. Even if both concepts represent digital and digitally signed documents, according to [Bis02] there is still a significant difference between them. Namely, certificates are related to an identity whereas credentials are key-based exclusively. Certificates confirm a set of free attributes of an identity (holder of free property) and connect a public key to its owner's private key. This connection is authenticated by means of the certificate. Contrary to this, credentials are used for adjudication of authorizations. The issuer of a credential assigns this way an access admission to the credential assignee.

However, mostly both concepts are used to describe a credential-based Public Key infrastructure. In the following, for Semantic Web applications the most relevant credential-based Public Key infrastructure SPKI/SDSI is outlined in order to give a more concrete insight in the working of such an infrastructure. Other credential-based Public Key infrastructures, like PolicyMaker [BFL96], KeyNote [BFK99], and REFEREE [CFL$^+$97] are referenced only.

Simple Public Key Infrastructure/Simple Distributed Security Infrastructure (SPKI/SDSI) Unlike other Public Key infrastructures, the credential-based SPKI/SDSI infrastructure allows each principal to issue credentials and thus requires no central certification authorities. For this reason, each Web service provider can issue and trust credentials independently of other service providers and may even define her own trust structures. Acting as a verifier, each Web service provider can also locally and autonomously decide whether access to her service should be granted or not. Access decisions are based on the provider's interpretation of a user's capabilities or characteristics given by previously shown SPKI/SDSI certificates. Moreover, users can request Web services spontaneously without registering themselves with the individual Web service providers. Therefore SPKI/SDSI credentials are more suitable than the classical authentication based systems for specifying access control policies in the Semantic Web.

Usually, SPKI/SDSI for credential-based access control is applied as proposed in [Ell99a, Ell99b, CEE$^+$01, Riv96]. The infrastructure distinguishes between two kinds of credentials, namely "name certificates" to bind principals to names and "authorization certificates" to bind authorizations to names. In addition to the both, SPKI/SDSI also provides "access control lists" (ACL) to specify access control policies for a certain interface.

- Name certificates allow definitions of local name spaces associated with every Public Key, since in SPKI/SDSI all principals are represented by their Public Keys. It has to be mentioned that there are no global name spaces in SPKI/SDSI. A name certificate is a document of the form "Keyholder,Name,Subject,Validity".

  - "Keyholder" represents the issuing principal who certifies an agent with a signature.
  - "Name" is an identifier defined by the issuing Keyholder to form a local name or a local name space if several principals are associated by the same issuing Keyholder.
  - "Subject" corresponds to a principal's local name.
  - "Validity" denotes to all principals if the certificate is currently valid.

- "Authorization certificates" are used to bind an authorization to a name and represents a document of the form "Keyholder,Subject,Authorization,Delegation,Validity".
  - "Keyholder" represents again the issuing principal who certifies a body with a signature.
  - "Subject" denotes a set of grantees of the authorization.
  - "Authorization" specifies all granted permissions.
  - "Delegation" represents a Boolean flag being set when a grantee is allowed to forward the permissions specified in Authorization to other principals.
  - "Validity" denotes again if the certificate is currently valid.

- Access control lists (ACL) are applied for access decisions in an SPKI/SDSI infrastructure, since there are no algebras defined to specify more complex access conditions. In an ACL-based system, a principal's access to an object depends simply on whether the identity of the principal is on a list associated with the object. Consequently, access control lists can only demand existence of a certificate but not enforce the absence of another certificate. However, definitions of more complex access conditions are a part of the latest extension of the infrastructure to so-called "SPKI/SDSI plus". Again, to access an interface, a requester must prove her eligibility or show that her set of credentials fulfills the access control policy of the interface. To do so, she constructs an authorizing set (chain) of certificates from the ACL and her own set of certificates. For this construction, a certificate chain discovery algorithm is suggested in [CEE+01].

3.2 Modeling Credentials

Using the fundamental constructs presented in previous sections, more complex models as integral parts of a trustworthy access control system can be designed. Remembering that practically every participant can certify capabilities defined in her own name space or terminology, determining the semantics of certified capabilities and the trustworthiness of certification authorities are two major challenges in such a setting. In the following, a general approach to a solution is outlined by showing

- how end users can check automatically, whether they can be granted access to a Web service or not,
- how Web service providers can specify and check their access control policies, and
- how certification authorities and their policies can be modeled semantically.

Users The concept user captures all entities that want to access Web services. In case of restricted access they prove their eligibility by showing appropriate credentials. Sometimes, users also wish to infer automatically if they can fulfill the access control requirements. In other situations, users may compose some individual Web services that could belong to different administrative domains. For this, they need to know the access control requirements. To support those use cases, the access control policies and the credentials have to be specified formally. Hence, a large number of specifications are proposed for different use cases in the literature (confer chapter 5).

Web Service Providers The main goal of the Web service providers is to restrict access to their services to eligible users only. For this, they specify and enforce individual access control policies in terms of capabilities that have to be proved by credentials. To specify those policies, questions like the following need to be addressed:

- Clarify the meaning of the individual terminology used by certification authorities in their certificates.
- Which credentials of which certification authority shall be trusted?
- Are the credentials still valid or maybe expired?
- How to specify a satisfiable access control policy whose conditions can all be fulfilled by at least one user?
- Is the specified access control policy consistent with laws and similar conditions?

Certification Authorities Certification authorities certify user properties by issuing certificates. Each certification authority defines its own terminology that it uses in its certificates, e.g. the names of certifiable properties and the relation between them. Currently, the certification authorities specify their policies explicitly in documents that are readable for humans only. In [Kel06] an extensive list of certification authorities is provided. Those documents are meant to be read by the service providers before defining the access control policies for their Web services.

Meanwhile, there are also first approaches to specify certification policies in a machine- readable form. For some specific use cases, a more detailed description of these approaches is provided, for instance, in [ASW04] and [Aga07].

3.3 Implementation of Credential-based Trust Systems for the Semantic Web

Regarding the implementation of and building trust in credential-based systems, an essential part is given by the policy specification for negotiating interactions, since the rules of negotiation determine how and if trust is achieved. These rules in turn are depending on the selected trust languages. Hence, in the following two subsections, first, an overview over most relevant works regarding trust negotiation and trust languages is provided.

3.3.1 Trust Negotiation

Trust negotiation and establishing trust causes the problem that revealing a credential may incur a loss of privacy or control of information. Hence, the focus of the works [Win00, WYS+02, YWS01, YW03] is on the trade-off between privacy and earning trust. According to these works, trust in a specific context is earned by revealing a certain number and type of credentials, but privacy of credential information is lost as the credentials are revealed.

TrustBuilder [WYS+02] represents an implemented architecture based on these principles, which provides mechanisms for addressing this trade-off. In TrustBuilder, trust is earned when sufficient credentials are shown, however, not too many to sacrifice privacy. Also applied in TrustBuilder, is the concept of a “credential chain”, where trust is transferred transitively through credentials, e.g. if A trusts the credentials of B, and B trusts the credentials of C, then A may have some trust in the credentials of C too. To perform credential chaining, the trust management method RT0 [LWM03] is designed explicitly and allows for an efficient distributed search to find such chains.

PeerTrust [NOW04] is a more recent policy and trust negotiation system that facilitates the automatic negotiation of a credential exchange. Following PeerTrust is PROTUNE [BO05], a provisional trust negotiation

framework. PROTUNE allows policies with “provisional predicates”, where actions can be specified in order to satisfy (currently unsatisfied) conditions.

To enable context-aware applications on the Semantic Web, [GS04] propose using ontologies for trust negotiation. Those context-aware applications shall only reveal credentials in the correct context. [LNO+04] also proposes ontologies to flexibly represent trust negotiation policies (rules used to negotiate trust). Ontologies have more flexibility than set standards. They simplify policy specification and they enable more information to be specified to control privacy during trust negotiation.

Others works in this area contribute ideas on client-server credential exchange [Win00], and protecting privacy through generalizing or categorizing credentials [SJ04].

3.3.2 Trust Languages

[Ton03] illustrates and compares several policy languages, designed for use in the Semantic Web.

One of those policy languages is known as KAoS [UBJ+03]. The major goal of KAoS is to enable the use of the same policy in distributed heterogeneous environments and to enable dynamic policy changes. Additionally, [UBJ+03] describes the KAoS “services” used to enforce its policies.

Another policy language known as Rei [Kag03] addresses security and privacy issues in the Semantic Web, while allowing each entity to specify their own policy. Using semantic representations, the Rei language separates policy from implementation and models “speech acts” (to programmatically “discuss” a policy at runtime) as a means of negotiation and dynamic policy manipulation.

Some policy languages, like the OASIS eXtensible Access Control Markup Language XACML [XAC05], still assume that trust is established through some external system, and this way keep trust and security separate. The extension to the OASIS Security Assertion Markup Language SAML [SAM05] provides a means for authentication and authorization, but is not able to represent or suggest trust. Consequently, SAML has the precondition that some external system is trusted.

To provide methods for the exchange of credentials, several standards for representation of policies and credentials have been proposed. WS-Trust [WS-05], an extension of WS-Security, specifies how trust is gained through proofs of identity, authorization, and performance. In this work, trust is approached from a hard security perspective, issuing a “security token” when trust is earned. WS-Trust does not address the trust negotiation process, only its representation.

The Cassandra system [BS04] applies a policy specification language that enforces how trust may be earned through the exchange of credentials. The system combines a role-based access control and context-based system for authorization. [Olm07] provides a comprehensive overview and a more detailed comparison of policy languages.

Chapter 4

Comparison of Trust Systems

Currently, there exist two different major approaches for managing trust: reputation and credential-based trust management. In previous chapters, the main features of reputation and credential-based trust systems have been introduced. In order to provide a better insight into those systems, also some applications have been described. In this chapter, at first, the significant differences between reputation- and credential-based applications are worked out.

Reputation-based and credential-based approaches have been developed within the context of different environments and targeting different requirements. Credential-based trust relies on objective “hard security” mechanisms like signed certificates and trusted certification authorities in order to control the access to services. Moreover, the access decision is often based on mechanisms with well-defined semantics providing strong verification and analysis support. The result of such an approach for managing trust usually consists of a binary decision according to whether the requester is trusted or not.

In contrast, reputation-based trust relies on a “soft computational” approach to the circumstance of trust. Here, trust is typically computed from local experiences combined with the feedbacks given by other entities in the network, e.g. users having used services of the same provider. This approach of trust management has been favored for environments such as Peer-to-Peer networks or the Semantic Web, generally spoken, where the existence of certifying authorities cannot always be assumed, but where a large pool of individual user ratings exists.

To provide a wider overview of trust applications, in the following two subsections, some selected applications for reputation-based and credential-based trust are given.

4.1 Reputation-Based Trust Applications

Reminding chapter 2, reputation-based trust systems are always useful when an entity needs to protect itself from those who offer resources. Corresponding use cases are present whenever the entity retrieves or downloads information from unknown Web sources. These downloaded data can, on the one hand, be provided from malicious users to spread dangerous software like viruses, Trojan horses etc. On the other hand, these data can be of very low quality or intentionally falsified to spread misleading information.

To overcome given or similar use cases, in chapter 2 several applications have been introduced which mainly focus either on trust management in P2P networks or trust management for Web applications like search-engines or e-commerce. In addition to the algorithms and applications already introduced in chapter 2, some more comparable approaches are itemized and referenced in the following:

- First of all, for a large number of Web applications [Jos07] provides extensive evaluations and descriptions of models and functionalities for the following applications being here just enumerated:
  - eBay’s Feedback Forum
  - Expert Sites like AllExperts, AskMe, Advogato

  - Product Review Sites
  - Epinions as a product and shop review site
  - BizRate for e.g. comparing prices and quality of products in Internet shops
  - Amazon
  - Discussion Fora like Slashdot, Kuro5in
  - Google's Web Page Ranking System
  - Supplier Reputation Systems like applied in Open Ratings.

- Regarding data exchange in P2P networks, where the trust systems provide mechanisms by which a peer requesting a resource may evaluate its trust in the reliability of the resource and the peer providing the resource itself.

Besides the mentioned examples of such systems, like EigenTrust, XRep, also the systems SPORAS [ZMM99], HISTOS [ZMM99], NICE2, DCRC/CORC [GJA03], Beta [Jos02], PeerTrust [XL03], EigenRep [Kam03] etc. establish trust relationships as a function of the combination of the peer's global reputation and the evaluating peer's perception of that peer.

4.2 Credential-Based Trust Applications

Credential-based trust management is usually proposed in the context of open and distributed services architectures. This approach provides a solution to the problem of authorization and access control in those open systems by issuing different access rights to different user groups. According to chapter 3, those access control policies can practically be defined by every Web service provider in a different way, so that it is unimportant to classify or to find similarities between all the different access policies. Rather, it can be stated that credential-based trust is applied in systems

- with strong protection requirements, or
- for systems whose behavior is determined by complex rules which
- must be easily changeable, as well as
- for systems where the nature of the information used in the authorization process is exact.

For example, such systems can be implemented in order to assign different access rights to different users to the same database. In a hospital, for example, a doctor would get different access rights to patient's documentation than a nurse or an IT administrator. In another conceivable scenario, a participant of a conference may be allowed to download data related to this conference from an IEEE server without being an IEEE member and thus previously logging in with his IEEE password, but only by showing credentials that he got from the conference organizer as a participant of the conference.

Chapter 5

Use Cases in NeOn

Already in [DDG+06], use cases were described where reputation- and credential-based trust played a role for a customized resolution of redundant and inconsistent information in integrated knowledge bases and for personalized views on a network of ontologies. In the following, we will briefly summarize the two use cases from [DDG+06] and discuss in how far either reputation- or credential-based trust can be used. Furthermore, we will describe two new use case about trust in e-invoicing and the semantic nomenclature.

5.1 Resolving Inconsistencies in and Personalized Views on a Network of Ontologies

Often inconsistencies occur when it comes to merging different knowledge bases even if each of the knowledge bases is itself consistent. A source for such inconsistencies are typically redundancies in the different knowledge bases. In [DDG+06], it was proposed to resolve them by taking contextual information of the particular user into account who does the merging. The contextual information may for example contain certain preferences like how the user rates the trustworthiness of information coming from the different knowledge bases. But also ratings of other users that the current user trusts can be taken into account. Once, less reliable information is identified based on its trust value, one can remove it step by step until the inconsistencies are resolved.

Another scenario described in [DDG+06] deals with personalized views on a network of ontologies. Already an ontology with removed inconsistencies based on the subjective trust ratings of a user can be seen as a personalized view. But further information may be removed from the network of ontologies based on the credentials owned by the user.

Thus, in the scenarios of [DDG+06] reputation-based trust as well as credential-based trust play a crucial role for providing customized views on networks of ontologies. On the one hand, we have a restriction of the view based on ratings of the ontologies in a network (i.e. based on their reputation) and on the other hand we have the restriction based on access rights or credentials of users.

5.2 Trust in e-invoicing

Trust is a very important issue in electronic invoicing. Potentially high amounts of money are exchanged each time an invoice is accepted by an organization. Thus, invoicing, in line with the general trends in electronic commerce, is a highly-regulated domain of application with a specific legal framework which needs to be observed at all times by the different organizations, in order to guarantee the integrity of the invoice data exchanged.

The EU council directive 77/388/EEC text recognizes the juridical validity of electronic invoice and fixes the modalities of its safeguard. According to this directive, invoicing conditions related to trust are given below:

- The receiver must give his approbation to the invoice when receiving it
- Authenticity, origin, integrity of the content and legality of the invoices must be safeguarded during the complete cycle of filing.
- Certification of the authenticity of origin and integrity of contents must also be safeguarded.
- E-invoices must be kept accessible on demand, with consistent format and in short delay.

On the other hand, business dynamic trends are evolving in the pharmaceutical sector. Laboratories and pharmacies are creating groups to improve the integration of the supply chain with integrated distribution. This concentration generates new opportunities for both and also for wholesalers. Wholesalers buy products from manufacturers (laboratories) and become owners of the medicine, then re-selling it to pharmacies.

On the other hand, clusters of laboratories built around an e-invoicing platform, like PharmaInnova, provide good conditions for exploitation of accumulated sectoral information in the form of the invoice flows of member laboratories. This information can be later exploited in a cooperative way in order to detect in real-time the current trends of the market, aiding the uptake of corrective measures by laboratories in order to e.g. safeguard their market share. However, like any other organization for profit, laboratories do no want their financial information to be revealed to other partners. Hence, ways need to be devised that allow exploiting such overall information on cluster-derived market trends while preserving specific data from being accessible to others but the owner organization. Real-time exploitation of market information for support of decision taking is the foundation of the so-called Business Intelligence 2.0, in contrast with traditional BI.

Accordingly to chapter 1, exchange of electronic invoices between two organizations is a transactional event that takes place under a model of reliability trust. Interacting organizations need to be reliable, in all occasions, independently from the context.

On the other hand, a scenario like the one described above where different laboratories agree on sharing their financial data for the common interest will not take place unless privacy of the data is guaranteed by a trust model which allows dealing with such data anonymously, profiting from the overall tendencies reflected in that data and not from the individual pieces of information.

Section 1.3 enumerates the most relevant properties that trust must have. Next, we contextualize them in the case of e-invoice exchange:

- **Subjectivity** is not an option in e-invoicing. The legal framework in which e-invoicing is embedded clearly and objectively determines that personal or past experience have little to do in this domain. There is no room either for uncertainty brought in by reputation-based trust models.
- **Assymetry** is a property which also holds in this domain. E-invoices need to be signed with a digital certificate by the emitting organization in order to deterministically provide the receiver with trust on the invoice content and on the organization that originated it.
- **Transitivity** is not a critical property for trust in e-invoicing. Additionally, since this kind of commercial transactions are made directly between peers, there is not much range for its application. We could, nevertheless, talk about hierarchy in e-invoicing due to the fact that certificates are issued by trusted, higher-level Certification Authorities.
- **Composability**: International e-commerce foresees the use of certificates issued by CAs from different organizations that can be equally used.
- **Personalization** is not possible in e-invoice exchange. E-invoicing is subject to strong regulations which tightly define trust therein.
- **Dynamic** is not an issue for invoicing exchange. Nevertheless, cooperative models of exploitation of private information like an eventual case of BI2.0 in PharmaInnova must contemplate the possibility of the realization that an agent that knows he is trusted may act differently from one who does not know his level of trust. Distrust in this scenario would be exceptionally hard to overcome.

Regarding how trust can be classified for e-invoicing (1.4), its definition suits best with a kind of interpersonal, or rather inter-organizational credentials-based kind of trust, which could be aligned as follows with respect to the dimensions of trust in Computer Science (1.4.2):

- **Target**: The target of trust are the contents of the invoice themselves and the identity of the invoice emitter.
- **Representation of trust**: Digitally encoded as a digital certificate, representing the credentials of the emitting organization.
- **Method**: The invoice emitter sends its credential, i.e. the digital certificate, to establish trust on the invoice.
- **Management**: Trust policies are centrally defined by the different CAs.
- **Computation**: Since trust depends on the validity of the certificate, trust, determined by such credentials, can only be granted or not, with no intermediate degrees of acceptance.
- **Purpose**: To validate and enable the commercial transaction represented by the invoice.

As a conclusion, the most general use of trust in e-invoicing, i.e. invoice exchange, is credentials-based, represented by the digital certificates signing the electronic invoices. Nevertheless, there is a number of added-value applications, like the support to decision-taking, in line with trends defined with BI2.0 regarding immediateness and accuracy in market information processing. Additionally, we have shown that the e-invoicing scenario presents a variation with respect to the expected properties of trust, defined in section 1.3. Finally, it should be stated that trust, as defined in NeOn should not only be applied to grant or denial of access to resources, but also to ensuring safe transactions between organizations and cooperating peers, as in the case of electronic invoice exchange.

5.3 Trust in the Semantic Nomenclature

The pharmaceutical industry is an important element of the medical assistance systems around the world; this sector is constituted by numerous public and private organizations dedicated to the research, development, manufacture and commercialization of medicines for the human and animal health.

In the semantic nomenclature case study, the information of the pharmaceutical products have different provenance: laboratories, government entities (Ministry of Health, Agemed, CCAA RegionsĚ), pharmacists associations (GSCoP[1]) and other sources of information like web pages, external databases, etc. Aspects like provenance and trust are of great importance for the Semantic Nomenclature, due to the fact that the case study is about aggregation of multiple and heterogeneous sources of information with different degrees of trust.

The pharmaceutical databases provided by the government are very important and useful for the pharma community, because they are the most trusted resources and provide data about the status of the different drugs in relation with the law (approved, withdrawn, etc). Of course, the reliability of the information coming from other external resources shall be ranked and validated, because these resources open new possibilities to infer knowledge and giving extra data to the Nomenclature's users and a new point of view of the pharmaceutical products. Also, the information managed in the Nomenclature scenarios is so critical.

Taking as a starting point the Nomenclature sketch, it could be identified two main scenarios where trust models could be applied. A reputation system and policies could be useful for trust and tracking the provenance of the drug information. There are several sources of information which provide drug content information, and

[1] GSCoP: General Spanish Council of Pharmacists

not all these sources should be dealt with the same reliability. All the information coming from the official government entities have the highest levels of trust, just like the information coming from certified laboratories, while drug information coming from other parties should have less trust in the reputation system.

Furthermore, it is identified a credential-based trust model in the Nomenclature scenario to gain access to the information provided by one of the actors of the pharmaceutical sector. The General Spanish Council of Pharmacists provides to its members a software tool called BOTPlus, developed by Atos. The database associated with BOTPlus contains homogeneous and updated information about medicines and sanitary products. It is a reference to the drug information for professionals. It offers information on diseases, symptoms, epidemics, treatments, detection of problems related to the medicines, etc. Some of this information can be publicly accessed, but the rest of the knowledge contained in BOTPlus is available just for the GSCoP and the associated pharmacists. Therefore, the BOTPlus Ontology that represents the knowledge of this source of information should have public and private parts, for which a credential-based trust mechanism is useful and necessary to allow the Nomenclature's users which have the correspondent and adequate credential to access the private part and obtain more information about the pharmaceutical products.

Chapter 6

Conclusions and Next Steps

In the development of modern open distributed and decentralized systems, trust management has become an important research line. A large number of different works studying trust in the context of reputation systems for P2P and Web applications, as well as Public Key certification and decentralized access control has been published. All the different approaches caused a plethora of meanings and concepts being used in newly designed systems (most of them designed from scratch).

Therefore, this work began with arranging and classifying all the terms and definitions of trust used in different scenarios. According to the currently existing two different major approaches for managing trust, namely reputation-based and credential-based trust management, these two approaches were outlined and discussed separately.

Finally, we listed several uses cases in NeOn that are related to trust. On the one hand, reputation- and credential-based trust may be used for providing personalized views on ontologies e.g. by (semi-)automatically resolving inconsistencies during merging knowledge bases or by removing those parts from ontologies for which the user doesn't have the required access credentials. Furthermore, we described the requirements in a more specific use case from WP8 that deals with sustaining trust during exchanging e-invoices. In the invoice scenario, trust is related to authenticating the identity of invoice emitters and signing the content of invoices. Those two requirements can be fulfilled by using existing public-key infrastructures (cf. section 3.1.2) and certification authorities. Furthermore, the access to financial data of laboratories has to be restricted and thus also a framework for treating access rights in ontologies is needed.

Within NeOn, already two proposals exist that cover aspects of reputation- and credential-based trust that are specific to ontologies and their lifecycle. Both proposals will help in fulfilling the requirements from the previously described use cases:

- **Reputation-based trust** In [SAd+07], a framework for Open Rating Systems is proposed. They allow users of ontologies for providing reviews of the quality and usefulness of ontologies as well as for rating the reviews of other users. The framework for Open Rating Systems proposed in [SAd+07] employs several ideas also described in this deliverable like the propagation of trust and distrust. It also proposes how several ratings and reviews can be aggregated for providing a trust based ranking of ontologies. One of the main strength of this approach is the fact that trust computation are in fact global and local, i.e. as soon as a user has connected to the web of trust, local (personalized) trust can be computed.

- **Credential-based trust** In [DKG+07], a framework for treating access rights in NeOn is introduced. It contains a more detailed overview on existing mechanisms for access control than is provided in this deliverable. Furthermore, it identifies requirements that have to be fulfilled by an access rights framework for ontologies. Specific requirements exist with regard to the level of granularity of access rights (e.g. on the level of whole ontologies or also on the level of single ontology elements), the inheritance of access rights within an ontology (e.g. based on the concept- or type-hierarchy) and with regard to delegating and revoking access rights.

Both approaches for reputation- and credential-based trust will be further pursued within NeOn. The treatment of access rights will be done in the context of WP4 while the Open Rating Systems will be evaluated and further refined in the context of WP2.

Bibliography

[Acr02] B. Acrement. Elements for building trust: do your management skills measure up?, 2002. http://www.imakenews.com/smei/earticle000051474.cfm.

[AD01] Karl Aberer and Zoran Despotovic. Managing trust in a peer-2-peer information system. In Henrique Paques, Ling Liu, and David Grossman, editors, *Proceedings of the Tenth International Conference on Information and Knowledge Management (CIKM01)*, pages 310–317. ACM Press, 2001.

[Aga07] Sudhir Agarwal. *Formal Description of Web Services for Expressive Matchmaking*. PhD thesis, Universität Karlsruhe (TH), Universität Karlsruhe (TH), Institut AIFB, D-76128 Karlsruhe, 2007.

[AR04] A. Abdul-Rahman. *A Framework for Decentralized Trust Reasoning*. PhD thesis, University College London, UK, 2004.

[ARH97a] Alfarez Abdul-Rahman and Stephen Hailes. A distributed trust model. In *NSPW '97: Proceedings of the 1997 workshop on New security paradigms*, pages 48–60, New York, NY, USA, 1997. ACM Press.

[ARH97b] Alfarez Abdul-Rahman and Stephen Hailes. Using recommendations for managing trust in distributed systems. In *IEEE Malaysia International Conference on Communication*, November 1997.

[ARH00] Alfarez Abdul-Rahman and Stephen Hailes. Supporting trust in virtual communities. In *HICSS '00: Proceedings of the 33rd Hawaii International Conference on System Sciences-Volume 6*, page 6007, Washington, DC, USA, 2000. IEEE Computer Society.

[AS04] Sudhir Agarwal and Barbara Sprick. Access control for semantic web services. In *Proceedings of 1st International Conference on Web Services*, JUL 2004.

[ASW04] Sudhir Agarwal, Barbara Sprick, and Sandra Wortmann. Credential based access control for semantic web services. In *AAAI Spring Symposium - Semantic Web Services*, MAR 2004.

[BCGM05] Christian Bizer, Richard Cyganiak, Tobias Gauss, and Oliver Maresch. The triql.p browser: Filtering information using context-, content- and rating-based trust policies. In *Semantic Web and Policy Workshop at the 4th International Semantic Web Conference, Galway, Ireland*, November 2005.

[BFK99] Matt Blaze, Joan Feigenbaum, and Angelos D. Keromytis. KeyNote: Trust management for public-key infrastructures (position paper). *Lecture Notes in Computer Science*, 1550:59–63, 1999.

[BFL96] Matt Blaze, Joan Feigenbaum, and Jack Lacy. Decentralized trust management. In *SP '96: Proceedings of the 1996 IEEE Symposium on Security and Privacy*, page 164, Washington, DC, USA, 1996. IEEE Computer Society.

[Bis02] Joachim Biskup. Credential-basierte zugriffskontrolle: Wurzeln und ein ausblick. In *Informatik bewegt: Informatik 2002 - 32. Jahrestagung der Gesellschaft für Informatik e.v. (GI)*, pages 423–428. GI, 2002.

[Bis03] Matt Bishop. *Computer Security: Art and Science*. Addison Wesley, 2003.

[BO04] Christian Bizer and Radoslaw Oldakowski. Using context- and content-based trust policies on the semantic web. In *WWW Alt. '04: Proceedings of the 13th international World Wide Web conference on Alternate track papers & posters*, pages 228–229, New York, NY, USA, 2004. ACM Press.

[BO05] Piero Bonatti and Daniel Olmedilla. Driving and monitoring provisional trust negotiation with metapolicies. In *POLICY '05: Proceedings of the Sixth IEEE International Workshop on Policies for Distributed Systems and Networks (POLICY'05)*, pages 14–23, Washington, DC, USA, 2005. IEEE Computer Society.

[Boe03] S. Boeyen. Liberty trust models guidelines. Liberty Alliance Project, 2003.

[BS04] Moritz Y. Becker and Peter Sewell. Cassandra: Distributed access control policies with tunable expressiveness. volume 00, page 159, Los Alamitos, CA, USA, 2004. IEEE Computer Society.

[CDdV+02] Fabrizio Cornelli, Ernesto Damiani, Sabrina De Capitani di Vimercati, Stefano Paraboschi, and Pierangela Samarati. Choosing reputable servents in a p2p network. In *WWW '02: Proceedings of the 11th international conference on World Wide Web*, pages 376–386, New York, NY, USA, 2002. ACM Press.

[CEE+01] Dwaine Clarke, Jean-Emile Elien, Carl Ellison, Matt Fredette, Alexander Morcos, and Ronald L. Rivest. Certificate chain discovery in spki?sdsi. *J. Comput. Secur.*, 9(4):285–322, 2001.

[CFL+97] Yang-Hua Chu, Joan Feigenbaum, Brian LaMacchia, Paul Resnick, and Martin Strauss. Referee: trust management for web applications. In *Selected papers from the sixth international conference on World Wide Web*, pages 953–964, Essex, UK, 1997. Elsevier Science Publishers Ltd.

[DDG+06] Klaas Dellschaft, Martin Dzbor, Jose Manuel Gomez, Carlos Buil Aranda, Dunja Mladenic, and Alexander Kubias. Review of methods and models for customizing/personalizing ontologies. Deliverable D4.2.1, NeOn Project, 2006.

[DdVP+02] Ernesto Damiani, De Capitani di Vimercati, Stefano Paraboschi, Pierangela Samarati, and Fabio Violante. A reputation-based approach for choosing reliable resources in peer-to-peer networks. In *CCS '02: Proceedings of the 9th ACM conference on Computer and communications security*, pages 207–216, New York, NY, USA, 2002. ACM Press.

[DKF+05] Li Ding, Pranam Kolari, Tim Finin, Anupam Joshi, Yun Peng, and Yelena Yesha. On Homeland Security and the Semantic Web: A Provenance and Trust Aware Inference Framework. In *Proceedings of the AAAI SPring Symposium on AI Technologies for Homeland Security*. AAAI Press, March 2005. (poster paper).

[DKG+04] Li Ding, Pranam Kolari, Shashidhara Ganjugunte, Tim Finin, and Anupam Joshi. Modeling and Evaluating Trust Network Inference. In *Seventh International Workshop on Trust in Agent Societies at AAMAS 2004*, July 2004.

[DKG+07] Martin Dzbor, Alexander Kubias, Laurian Gridinoc, Angel Lopez-Cima, and Carlos Buil Aranda. The role of access rights in ontology customization. Deliverable D4.4.1, NeOn Project, 2007.

[DRJ04] Rajdeep K. Dash, Sarvapali D. Ramchurn, and Nicholas R. Jennings. Trust-based mechanism design. In *AAMAS '04: Proceedings of the Third International Joint Conference on Autonomous Agents and Multiagent Systems*, pages 748–755, Washington, DC, USA, 2004. IEEE Computer Society.

[DZF03] Li Ding, Lina Zhou, and Timothy Finin. Trust based knowledge outsourcing for semantic web agents. In *WI '03: Proceedings of the 2003 IEEE/WIC International Conference on Web Intelligence*, page 379, Washington, DC, USA, 2003. IEEE Computer Society.

[Eck04] Claudia Eckert. *IT-Sicherheit*. Oldenbourg, München [u.a.], 2004.

[Ell99a] C.M. Ellison. Simple public key certificate, 1999.

[Ell99b] C.M. Ellison. Simple public key certificate. Internet RFC 2693, 1999.

[Fah02] D. Fahrenholz. Transactional security for a distributed reputation management system. In *Prooceedings of the 3rd Int. Conf. on e-Commerce and Web Technologies*, 2002.

[FC04] Rino Falcone and Cristiano Castelfranchi. Trust dynamics: How trust is influenced by direct experiences and by trust itself. In *AAMAS '04: Proceedings of the Third International Joint Conference on Autonomous Agents and Multiagent Systems*, pages 740–747, Washington, DC, USA, 2004. IEEE Computer Society.

[FPHKH00] Batya Friedman, Jr. Peter H. Khan, and Daniel C. Howe. Trust online. *Commun. ACM*, 43(12):34–40, 2000.

[Gam88] Diego Gambetta. *Can We Trust Trust?* Basil Blackwell, 1988.

[Gef02] David Gefen. Reflections on the dimensions of trust and trustworthiness among online consumers. *SIGMIS Database*, 33(3):38–53, 2002.

[GGMP04] Zoltán Gyöngyi, Hector Garcia-Molina, and Jan Pedersen. Combating web spam with trustrank. In *VLDB*, pages 576–587, 2004.

[GH04] Jennifer Golbeck and James A. Hendler. Accuracy of metrics for inferring trust and reputation in semantic web-based social networks. In Enrico Motta, Nigel Shadbolt, Arthur Stutt, and Nicholas Gibbins, editors, *EKAW*, volume 3257 of *Lecture Notes in Computer Science*, pages 116–131. Springer, 2004.

[GJA03] Minaxi Gupta, Paul Judge, and Mostafa Ammar. A reputation system for peer-to-peer networks. In *NOSSDAV '03: Proceedings of the 13th international workshop on Network and operating systems support for digital audio and video*, pages 144–152, New York, NY, USA, 2003. ACM Press.

[GKRT04] R. Guha, Ravi Kumar, Prabhakar Raghavan, and Andrew Tomkins. Propagation of trust and distrust. In *WWW '04: Proceedings of the 13th international conference on World Wide Web*, pages 403–412, New York, NY, USA, 2004. ACM Press.

[Gol04] J. Golbeck. Inferring reputation on the semantic web. In *Proceedings of the 13th Int. World Wide Web Conf.*, 2004.

[Gol05] J. Golbeck. *Computing and Applying Trust in Web-Based Social Networks*. PhD thesis, University of Maryland, USA, 2005.

[Gol06a] J. Golbeck. Filmtrust: Moview recommendations using trust in web-based social networks. In *Proceedings of the 3rd IEEE Consumer Communications and Networking Conference*, 2006.

[Gol06b] Jennifer Golbeck. Combining provenance with trust in social networks for semantic web content filtering. In Luc Moreau and Ian T. Foster, editors, *IPAW*, volume 4145 of *Lecture Notes in Computer Science*, pages 101–108. Springer, 2006.

[Gol06c] D. Gollmann. *Computer Security*. John Wiley & Sons, 2006.

[GS00] Tyrone Grandison and Morris Sloman. A survey of trust in internet applications. *IEEE Communications Surveys and Tutorials*, 3(4), 2000.

[GS04] Fabien L. Gandon and Norman M. Sadeh. Semantic web technologies to reconcile privacy and context awareness. *Web Semantics: Science, Services and Agents on the World Wide Web*, 1(3):241–260, April 2004.

[Guh03] R. Guha. Open rating systems, 2003.

[HMM+00] Amir Herzberg, Yosi Mass, Joris Michaeli, Yiftach Ravid, and Dalit Naor. Access control meets public key infrastructure, or: Assigning roles to strangers. In *SP '00: Proceedings of the 2000 IEEE Symposium on Security and Privacy*, page 2, Washington, DC, USA, 2000. IEEE Computer Society.

[Hou99] R. Housley. Internet x.509 public key infrastructure: Certificate and crl profile. RFC 2459, 1999.

[Jos99] A. Josang. Trust-based decision making for electronic transactions. In *Proceedings of the 4th Nordic Workshop on Secure IT Systems*, 1999.

[Jos01] A. Josang. A logic for uncertain probabilities. *Int. Journal of Uncertainty*, 9:279–212, 2001.

[Jos02] A. Josang. The beta reputation system. In *Proceedings of the 15th Bled Conference on Electronic Commerce*, 2002.

[Jos05] A. Josang. Can we manage trust? In *Proceedings of the 3rd Int. Conf. on Trust Managemetn*, 2005.

[Jos06] A. Josang. A method for access authorization through delegation networks. In *Proceedings of the 4th Australasian Information Society Workshop*, 2006.

[Jos07] A. Josang. A survey of trust and reputation systems for online service provision. *Decision Support Systems*, 43:618–644, 2007.

[JSTT04] Catholijn M. Jonker, Joost J. P. Schalken, Jan Theeuwes, and Jan Treur. Human experiments in trust dynamics. In Christian Damsgaard Jensen, Stefan Poslad, and Theodosis Dimitrakos, editors, *iTrust*, volume 2995 of *Lecture Notes in Computer Science*, pages 206–220. Springer, 2004.

[Kag03] L. Kagal. A policy based approach to security for the semantic web. In *Proceedings of the 2nd Int. Semantic Web Conf.*, 2003.

[Kam03] S.D. Kamvar. Eigenrep: Reputation management in p2p networks. In *Proceedings of the 12th Int. Conf. on World Wide Web*, 2003.

[Kel06] S. Kelm. The pki-page - extensive list of certification authorities, 2006.

[KSGM03] Sepandar D. Kamvar, Mario T. Schlosser, and Hector Garcia-Molina. The eigentrust algorithm for reputation management in p2p networks. In *WWW '03: Proceedings of the 12th international conference on World Wide Web*, pages 640–651, New York, NY, USA, 2003. ACM Press.

[Lev03] R. Levien. *Attack Resistant Trust Metrics*. PhD thesis, UC Berkeley, USA, 2003.

[Lia03] C.Y. Liau. Effiecient distributed reputation scheme for peer-to-peer systems. In *Proceedings of the 2nd Int. Human.Society@Internet Conf.*, 2003.

[LMW02] Ninghui Li, John C. Mitchell, and William H. Winsborough. Design of a role-based trust-management framework. In *SP '02: Proceedings of the 2002 IEEE Symposium on Security and Privacy*, page 114, Washington, DC, USA, 2002. IEEE Computer Society.

[LNO+04] Travis Leithead, Wolfgang Nejdl, Daniel Olmedilla, Kent E. Seamons, Marianne Winslett, Ting Yu, and Charles C. Zhang. How to exploit ontologies for trust negotiation. In Jennifer Golbeck, Piero A. Bonatti, Wolfgang Nejdl, Daniel Olmedilla, and Marianne Winslett, editors, *ISWC Workshop on Trust, Security, and Reputation on the Semantic Web*, volume 127 of *CEUR Workshop Proceedings*. CEUR-WS.org, 2004.

[LWM03] Ninghui Li, William H. Winsborough, and John C. Mitchell. Distributed credential chain discovery in trust management. *J. Comput. Secur.*, 11(1):35–86, 2003.

[MA05] Paolo Massa and Paolo Avesani. Controversial users demand local trust metrics: An experimental study on epinions.com community. In *AAAI*, pages 121–126, 2005.

[Man98] Daniel W. Manchala. Trust metrics, models and protocols for electronic commerce transactions. In *ICDCS '98: Proceedings of the The 18th International Conference on Distributed Computing Systems*, page 312, Washington, DC, USA, 1998. IEEE Computer Society.

[Mar94] Stephen Paul Marsh. *Formalising Trust as a Computational Concept*. PhD thesis, University of Stirling, 1994.

[Mar05] Oliver Maresch. Reputationsbasierte trustmetriken im kontext des semantic web. Master's thesis, Technische Universität Berlin, 2005.

[McK96] D.H. McKnight. The meanings of trust, 1996.

[McK01] McKnight. A topology and e-commerce customer relationshops model. In *Proceedings of the 34th Hawaii Int. Conf. on System Sciences*, 2001.

[MDS95] Roger C. Mayer, James H. Davis, and David F. Schoorman. An integrative model of organizational trust. *The Academy of Management Review*, 20(3):709–734, 1995.

[MH05] Paolo Massa and Conor Hayes. Page-rerank: Using trusted links to re-rank authority. In *WI '05: Proceedings of the 2005 IEEE/WIC/ACM International Conference on Web Intelligence*, pages 614–617, Washington, DC, USA, 2005. IEEE Computer Society.

[MMH02] L. Mui, M. Mohtashemi, and A. Halberstadt. A computational model of trust and reputation for e-businesses. In *HICSS '02: Proceedings of the 35th Annual Hawaii International Conference on System Sciences (HICSS'02)-Volume 7*, page 188, Washington, DC, USA, 2002. IEEE Computer Society.

[NOW04] Wolfgang Nejdl, Daniel Olmedilla, and Marianne Winslett. Peertrust: Automated trust negotiation for peers on the semantic web. In Willem Jonker and Milan Petkovic, editors, *Secure Data Management*, volume 3178 of *Lecture Notes in Computer Science*, pages 118–132. Springer, 2004.

[Olm07] D. Olmedilla. Security and privacy on the semantic web. Security, Privacy and Trust in Modern Data Management, 2007.

[ORMN05] Daniel Olmedilla, Omer F. Rana, Brian Matthews, and Wolfgang Nejdl. Security and trust issues in semantic grids. In Carole A. Goble, Carl Kesselman, and York Sure, editors, *Semantic Grid*, volume 05271 of *Dagstuhl Seminar Proceedings*. Internationales Begegnungs- und Forschungszentrum für Informatik (IBFI), Schloss Dagstuhl, Germany, 2005.

[PBMW98] Lawrence Page, Sergey Brin, Rajeev Motwani, and Terry Winograd. The pagerank citation ranking: Bringing order to the web. Technical report, Stanford Digital Library Technologies Project, 1998.

[PM04] Asad A. Pirzada and Chris Mcdonald. Establishing trust in pure ad-hoc networks. In *ACSC '04: Proceedings of the 27th Australasian conference on Computer science*, pages 47–54, Darlinghurst, Australia, Australia, 2004. Australian Computer Society, Inc.

[RAD03] Matthew Richardson, Rakesh Agrawal, and Pedro Domingos. *Trust Management for the Semantic Web*, volume 2870. January 2003.

[Riv96] R.L. Rivest. Sdsi- a simple distributed security infrastructure, 1996.

[RJ96] Lars Rasmusson and Sverker Jansson. Simulated social control for secure internet commerce. In *NSPW '96: Proceedings of the 1996 workshop on New security paradigms*, pages 18–25, New York, NY, USA, 1996. ACM Press.

[RKM06] Sebastian Ries, Jussi Kangasharju, and Max Mühlhäuser. A Classification of Trust Systems. In R. Meersman, Z. Tari, and P. Herrero, editors, *On the Move to Meaningful Internet Systems 2006: OTM Workshops*, volume 4277 of *Lecture Notes in Computer Science*, pages 894 – 903, Montpellier, France, EU, October/November 2006. Springer, Berlin, Germany, EU.

[Sab03] J. Sabater. *Trust and reputation for agent societies*. PhD thesis, Institut d'Investigation en Intelligencia Artificial, Spain, 2003.

[SAd+07] Marta Sabou, Sofia Angeletou, Mathieu dŠAquin, Jesús Barrasa, Klaas Dellschaft, Aldo Gangemi, Jos Lehmann, Holger Lewen, Diana Maynard, Dunja Mladenic, Malvina Nissim, Wim Peters andand Valentina Presutti, and Boris Villazón. Methods for selection and integration of reusable components from formal or informal user specifications. Deliverable D2.2.1, NeOn Project, 2007.

[Sam01] P. Samarati. Access control: policies, models and mechanisms. *Foundations of Security Analysis and design*, 2171:137–196, 2001.

[SAM05] Saml: Oasis security assertion markup language, 2005.

[SJ04] Jean-Marc Seigneur and Christian Damsgaard Jensen. Trust enhanced ubiquitous payment without too much privacy loss. In *SAC '04: Proceedings of the 2004 ACM symposium on Applied computing*, pages 1593–1599, New York, NY, USA, 2004. ACM Press.

[SS01] Jordi Sabater and Carles Sierra. Regret: A reputation model for gregarious societies. In *Fourth Workshop on Deception Fraud and Trust in Agent Societies*, pages 61–70, 2001.

[SS02a] Jordi Sabater and Carles Sierra. Reputation and social network analysis in multi-agent systems. In *First International Conference on Autonomous Agents and Multiagent systems (AAMAS-02)*, pages 475–482, 2002.

[SS02b] Jordi Sabater and Carles Sierra. Social regret, a reputation model based on social relations. *ACM, SIGecom Exchanges*, 3.1:44–56, 2002.

[Ste99] Katherine J. Stewart. Transference as a means of building trust in world wide web sites. In *ICIS '99: Proceeding of the 20th international conference on Information Systems*, pages 459–464, Atlanta, GA, USA, 1999. Association for Information Systems.

[SZ03] Katherine J. Stewart and Yali Zhang. Effects of hypertext links on trust transfer. In *ICEC '03: Proceedings of the 5th international conference on Electronic commerce*, pages 235–239, New York, NY, USA, 2003. ACM Press.

[Ton03] G. Tonti. Semantic web languages for policy representation and reasoning: A comparison of kaos , rei, and ponder. In *Proceedings of the 2nd Int. Semantic Web Conf.*, 2003.

[UBJ+03] Andrzej Uszok, Jeffrey M. Bradshaw, Renia Jeffers, Niranjan Suri, Patrick J. Hayes, Maggie R. Breedy, Larry Bunch, Matt Johnson, Shriniwas Kulkarni, and James Lott. Kaos policy and domain services: Toward a description-logic approach to policy representation, deconfliction, and enforcement. In *POLICY*, pages 93–. IEEE Computer Society, 2003.

[Win00] W.H. Winsborough. Automated trust negotiation. In *Proceedings of the DARPA Information Survivability Conference and Exposition*, 2000.

[WS-05] Ws-trust: Web service trust language, 2005.

[WYS+02] Marianne Winslett, Ting Yu, Kent E. Seamons, Adam Hess, Jared Jacobson, Ryan Jarvis, Bryan Smith, and Lina Yu. Negotiating trust on the web. *IEEE Internet Computing*, 6(6):30–37, 2002.

[XAC05] Xacml: Oasis extensible access control markup language, 2005.

[XL03] Li Xiong and Ling Liu. A reputation-based trust model for peer-to-peer ecommerce communities [extended abstract]. In *EC '03: Proceedings of the 4th ACM conference on Electronic commerce*, pages 228–229, New York, NY, USA, 2003. ACM Press.

[YS02] Bin Yu and Munindar P. Singh. An evidential model of distributed reputation management. In *AAMAS '02: Proceedings of the first international joint conference on Autonomous agents and multiagent systems*, pages 294–301, New York, NY, USA, 2002. ACM Press.

[YW03] Ting Yu and Marianne Winslett. Policy migration for sensitive credentials in trust negotiation. In *WPES '03: Proceedings of the 2003 ACM workshop on Privacy in the electronic society*, pages 9–20, New York, NY, USA, 2003. ACM Press.

[YWS01] Ting Yu, Marianne Winslett, and Kent E. Seamons. Interoperable strategies in automated trust negotiation. In *CCS '01: Proceedings of the 8th ACM conference on Computer and Communications Security*, pages 146–155, New York, NY, USA, 2001. ACM Press.

[Zie04] Cai-Nicolas Ziegler. Semantic Web recommender systems. In Wolfgang Lindner, Mauro Mesiti, Can Türker, Yannis Tzitzikas, and Athena Vakali, editors, *EDBT 2004 Workshops (PhD, DataX, PIM, P2P&DB, and ClustWeb)*, volume 3268 of *LNCS*, pages 78–89, Heraklion, Greece, November 2004. Springer-Verlag. Revised selected papers.

[Zim95] Philip R. Zimmermann. *The official PGP user's guide*. MIT Press, Cambridge, MA, USA, 1995.

[ZL04] Cai-Nicolas Ziegler and Georg Lausen. Spreading activation models for trust propagation. In *Proceedings of the IEEE International Conference on e-Technology, e-Commerce, and e-Service*, Taipei, Taiwan, March 2004. IEEE Computer Society Press.

[ZL05] Cai-Nicolas Ziegler and Georg Lausen. Propagation models for trust and distrust in social networks. *Information Systems Frontiers*, 7(4-5):337–358, 2005.

[ZMM99] Giorgos Zacharia, Alexandros Moukas, and Pattie Maes. Collaborative reputation mechanisms in electronic marketplaces. In *HICSS '99: Proceedings of the Thirty-second Annual Hawaii International Conference on System Sciences-Volume 8*, page 8026, Washington, DC, USA, 1999. IEEE Computer Society.

[ZY04] Qing Zhang and Ting Yu. A classification scheme for trust functions in reputation-based trust management. In *ISWC Workshop on Trust, Security, and Reputation on the Semantic Web*, 2004.