

A Formal Approach Towards Measuring Trust in Distributed Systems

Guido Schryen
RWTH Aachen
Templergraben 55
52062 Aachen
schryen@winfor.rwth-aachen.de

Melanie Volkamer, Sebastian Ries,
Sheikh Mahbub Habib
CASED / TU Darmstadt
Mornewegstrasse 32
64293 Darmstadt
firstname.lastname@cased.de

ABSTRACT

Emerging digital environments and infrastructures, such as distributed security services and distributed computing services, have generated new options of communication, information sharing, and resource utilization in past years. However, when distributed services are used, the question arises of to what extent we can trust service providers to not violate security requirements, whether in isolation or jointly. Answering this question is crucial for designing trustworthy distributed systems and selecting trustworthy service providers. This paper presents a novel trust measurement method for distributed systems, and makes use of propositional logic and probability theory. The results of the qualitative part include the specification of a formal trust language and the representation of its terms by means of propositional logic formulas. Based on these formulas, the quantitative part returns trust metrics for the determination of trustworthiness with which given distributed systems are assumed to fulfill a particular security requirement.

1. INTRODUCTION

Emerging digital environments and infrastructures have rapidly generated new ways and services of communication, information sharing, and resource utilization for individuals, organizations, and societies in past years. For example, it has become common for individuals to use security services, such as *I2P Anonymous Network* and *TOR*. Organizations have started to explore the opportunities of web services, including storage services (e.g., *Amazon Simple Storage Service*) and computing services (e.g., Microsoft's *Azure Services Platform* and *Google App Engine*). While the aforementioned services are realized with cloud computing, services can also be requested from multiple administrative domains (*grid computing*). Even whole societies are involved in scenarios with shared information and transaction processing, as political elections with Internet voting systems show [31].

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

SAC'11 March 21-25, 2011, TaiChung, Taiwan.
Copyright 2011 ACM 978-1-4503-0113-8/11/03 ...\$10.00.

What all these services have in common is that some kind of distributed information processing and/or information sharing occurs, across private, organizational, or national boundaries. Often, consumers of these services have no control over their data, and they need to trust service providers not to violate their security policies. For example, scientific computation results can be modified or provided to third parties. In some cases, organizational, legal, and/or technical countermeasures have been taken in order to prevent or to mitigate the consequences of data abuse. For example, in Internet voting the separation of duties is quite common in order to realize the separation of voter's identity and his/her vote. In such cases, the abuse of data by a single party does not disclose confidential information. However, what happens when multiple parties maliciously cooperate and join their information? This leads to scenarios where a voter's ID can be assigned to his/her vote, where the identity of a user is disclosed through the cooperation of parties of an anonymity mix net, etc. Attacks like these are referred to as "insider attacks"¹, which are committed by service providers in isolation or jointly. Consequently, when distributed services are used, the question arises of whether service customers can trust the aggregated service and the underlying distributed system, or, to be more precise, of how they can determine the level of trust. Answering these questions is crucial for designing trustworthy distributed systems and selecting trustworthy service providers.

Similar to the field of security where it is widely argued that measurement and metrics are necessary to justify investments in security and to manage security [1,24,25], trust assessment should capture quantitatively the perceived ability of the system to resist attacks against a particular security requirement. The main purpose of this paper is to present a novel approach for trust measurement in heterogeneous distributed systems that relies on propositional logic and probability theory, and that includes the proposition of trust metrics.

The remainder of this paper is structured as follows: Section 2 presents related work. In Section 3, we describe our research framework. Section 4 proposes a formal trust language and demonstrates its applicability. Section 5 shows how language terms can be mapped on propositional logic terms, which are used in Section 6 to develop trust metrics. Finally, we discuss our approach and propose further research paths.

¹We do not consider outsider attacks, such as those committed by hackers.

2. RELATED WORK

There is a substantial body of literature on concepts, models, evaluation, and management of trust in digital environments (see [13, 17, 27] for a detailed overview). Analyzing this body, [22] identifies two lines of research: The first strand is based on a technical understanding coined by [6] and includes the “access-control list” approach and the “credential-based” approach, e.g., [5, 7]. The second strand is “experience-based” and assumes that an entity’s trust in another one is based on past behavior [22]. This strand focuses on general models for deriving trust from previous evidence that are especially capable of expressing the uncertainty that is associated to the derived trust values, e.g., [16, 26], on the robust aggregation of direct experience and recommendations [28, 34, 39], and on the application of those models in different scenarios like eCommerce [3], mobile ad hoc networks [10], or P2P networks [20]. Security-related applications of trust have especially been proposed in the field of public key infrastructures starting in the 90s [2, 23] up to now, e.g., [21] is a valuable example and presents a model for reliability and public-key authenticity that is based on logic and probability theory.

Our paper also draws on this theoretical basis, but it focuses on the evaluation of the trustworthiness of distributed systems based on the knowledge of the trustworthiness of its components. To this end, we build on current “experience-based”, probabilistic approaches for modeling trust, however, we provide a novel concept for deriving the trustworthiness of a complex system from the trustworthiness of its components. The proposed approach is especially capable of dealing with the dependencies between the components in the system and their redundancy with respect to the security requirements under consideration. In this way, our contribution is also related to approaches for evaluating security requirements, e.g., [14, 18, 30]. However, attack trees require knowledge on the implementation [18, 30], and the evaluation of system reliability addresses a particular security requirement only [14].

Regarding metrics, the literature often provides joint taxonomies of security and trust metrics. We identified two perspectives on such metrics: The first perspective is aligned to the objects the metrics refer to. [36] propose a taxonomy for information assurance metrics consisting of organizational security metrics and metrics for “Technical Target of Assessment”. [29] suggest a high-level information security metrics taxonomy that divides business-level security metrics into five categories: (1) trust metrics for business collaboration, (2) security metrics for business-level risk management, (3) security metrics for information security management in the organization, (4) security metrics for cost-benefit analysis, and (5) security, dependability and trust metrics for ICT products, systems and services. The metrics proposed in this paper fall into categories 1 and 5. The literature review of [33] distinguishes between metrics based on policy compliance and technical metrics. The NIST *Performance Measurement Guide for Information Security* [24] focuses on ex post security measures (the document does not sharply distinguish metrics and measures). [9] surveys economic approaches for security metrics. The authors identify two main areas of research, where one has its roots in investment and decision theory and is mainly pursued in the field of information technology-oriented business administration, and the other area has ancestors in micro-economics and deals with

market concepts to gather security-relevant information.

3. RESEARCH FRAMEWORK

Before we present our research framework, we operationalize key notions of our paper, *measurement*, *metric*, and *trust*, none of which are consistently defined in the literature. We refer to

- *measurement* as an abstraction process that “[...] reduces the complexity of one or several attributes of the original system to a single symbol [or to several symbols]” [8, p. 7], where a symbol can be at various numerical scale levels (e.g., at nominal level for the purpose of classification, or at ordinal or cardinal level for the purpose of comparison),
- *metric* as a well defined function that takes a particular system from the set of systems S and maps it to an element of an (ordered) set V , i.e. $f : S \Rightarrow V$ [8, p. 7],
- *trust* as “[...] a particular level of the subjective probability with which an agent assesses that another agent or group of agents will perform a particular action [...]” [11, p. 217f].

In our context, “agents or group of agents” refers to service providers, who provide the requested service by contemporaneously meeting a particular security requirement r (e.g., anonymity, confidentiality). The assurance of r corresponds to what is referred to as “perform a particular action”.

An overview of our research approach is shown in Figure 1. The basic concept of the model we propose in this paper draws on [15, 37], who use secret shares [4] and the concept that k out of n entities are required for reconstructing the secret key and correspondingly decrypting a cyphertext. We adapt this concept to the context of trust, and propose an approach for modeling distributed, heterogenous systems where “ k out of N entities need to show trustworthy behavior” (compare to [38]), in order to not compromise the trustworthiness of the overall system. Note, in case of Shamir’s secret sharing [32] the property k out of n means that you need to trust k out of N regarding availability and $n - k + 1$ out of N regarding secrecy (where N is the set of share holders and n the number of share holders). Thus, the choice of the actual value of k depends on the security requirement under consideration.

In contrast to the aforementioned papers, which regard entities as homogeneous, we account for heterogeneity of entities in terms of security requirements by explicitly itemizing them in a set N . We use this model for the proposition of a formal trust language, which contains *trustworthiness terms*, which formally describe trustworthiness properties of a system (regarding a particular security requirement). It should be noticed that different security requirements on a system can lead to different trustworthiness terms. E.g., in a system that implements a mixnet that routes messages sequentially through a set N of n anonymizing nodes, each node must be trustworthy with regard to availability (n out of N), while only one node needs to be trustworthy with regard to achieving anonymity (1 out of N). Consequently, a system needs to deal with a trade-off between different security requirements.

While trustworthiness terms are a useful representation of trustworthiness, they are less appropriate for comparing alternative systems and for determining the probability with which systems meet a particular security requirement. We

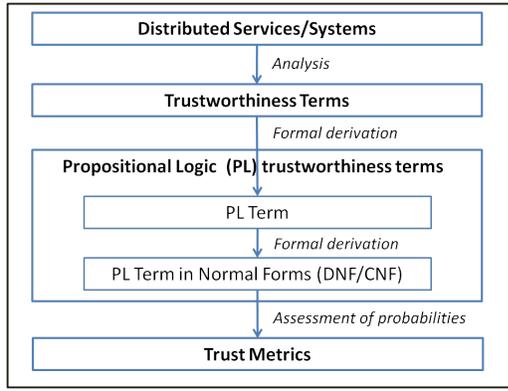


Figure 1: Research framework

prove that each trustworthiness term can be mapped on a propositional logic term such that both terms are semantically equivalent, and we show that this mapping is a useful way to overcome these limitations. We also show that propositional logic terms (given in conjunctive normal form) are a representation that allows to determine straightforward the overall probability with which a system is assumed to fulfill a particular security requirement. The determination of those probabilities result in the definition of trust metrics.

4. TRUSTWORTHINESS TERMS

As our trustworthiness terms address distributed systems, we first define distributed systems: A distributed system is either an “atomic system” or is composed of other (sub-)systems. We define a system as “atomic” if it contains only (atomic) components that are not being split any further. These components can be persons, computers, or even organizational units.

The definition of trustworthiness terms in terms of syntax and semantics follows the inductive definition of systems and is provided by definitions 1-4. In order to keep definitions short, we introduce the abbreviation “*wrts. r*” (with regard to security requirement r).

Let S be an atomic system with the set of atomic components $A = \{A_i\}_{i=1}^n$.

Definition 1. A system S can be described by the trustworthiness term (k out of N), $k \in \{1, \dots, |N|\}$, $N \subseteq A$, wrts. r

$:\Leftrightarrow \left\{ \begin{array}{l} \text{At least } k \text{ components out of } N \text{ need to show trust-} \\ \text{worthy behavior wrts. } r \text{ so that } S \text{ meets require-} \\ \text{ment } r. \end{array} \right.$

In order to get more flexible representations of requirements on atomic systems, we define the following trustworthiness terms:

Definition 2. A system S can be described by the trustworthiness term a) $((k_1 \otimes \dots \otimes k_m)$ out of (N_1, \dots, N_m)), b) $((k_1 \otimes \dots \otimes k_m)$ out of (N_1, \dots, N_m)), $k_i \in \{1, \dots, |N_i|\}$, $N_i \subseteq A \forall i$, wrts. r

$:\Leftrightarrow \left\{ \begin{array}{l} \text{For a) each } i \in \{1, \dots, m\}, \text{ b) any } i \in \{1, \dots, m\}, \\ \text{at least } k_i \text{ components out of } N_i \text{ need to show} \\ \text{trustworthy behavior wrts. } r \text{ so that } S \text{ meets re-} \\ \text{quirement } r. \end{array} \right.$

With regard to non-atomic systems, we define trustworthiness terms similarly: Let $\{S_i\}_{i=1}^n$ be (sub)systems of a

system S , and let system S_i be described by the following trustworthiness term l_i for all $i \in \{1, \dots, n\}$.

Definition 3. A system S can be described by the trustworthiness term (k out of $\{l_{i_1}, \dots, l_{i_m}\}$), $k \in \{1, \dots, m\}$, $\{i_1, \dots, i_m\} \subseteq \{1, \dots, n\}$, wrts. r

$:\Leftrightarrow \left\{ \begin{array}{l} \text{At least } k \text{ systems out of } \{S_{i_1}, \dots, S_{i_m}\} \text{ need to} \\ \text{show trustworthy behavior wrts. } r \text{ so that } S \text{ meets} \\ \text{requirement } r. \end{array} \right.$

Definition 4. A system S can be described by the trustworthiness term a) $((k_1 \otimes \dots \otimes k_m)$ out of (Q_1, \dots, Q_m)), b) $((k_1 \otimes \dots \otimes k_m)$ out of (Q_1, \dots, Q_m)), $k_i \in \{1, \dots, |Q_i|\}$, $Q_i \subseteq \{l_1, \dots, l_n\} \forall i$, wrts. r

$:\Leftrightarrow \left\{ \begin{array}{l} \text{For a) each } i \in \{1, \dots, m\}, \text{ b) any } i \in \{1, \dots, m\}, \\ \text{at least } k_i \text{ systems out of the set of systems} \\ \text{for which } Q_i \text{ contains trustworthiness terms need} \\ \text{show trustworthy behavior wrts. } r \text{ so that } S \text{ meets} \\ \text{requirement } r. \end{array} \right.$

We now illustrate the analysis and the determination of trustworthiness terms with an example.

Example 1. We use a web service scenario, in which a retailer uses three web services in order to identify customers’ behavior. Service A offers data mining capabilities and stores sales data, including customer IDs. Service B is offered by a financial service provider, who provides credit ratings of customers. Service C provides storage capacities and stores master data on customers, including their customer IDs and identities. In this example, we consider secrecy with regard to information on which customer has bought what under which financial conditions. Secrecy is kept if one of the providers A and B is trustworthy, or if one of B and C is trustworthy. With regard to provider A , we assume that this provider accounts for secrecy by storing data on two components (A_3 and A_4) and implementing a secret share mechanism [4]. Components A_1 and A_2 are responsible for distributed computation in terms of data mining; both components get data from A_3 and A_4 . With regard to financial service provider B , customer IDs generated by B (they differ from customer IDs stored at A) are stored on B_1 and B_2 together with financial data by implementing a secret share mechanism. Components B_3 and B_4 store names of customers and customer IDs (generated by B) redundantly. Analogous to A and B , storage provider C implements a secret share mechanism when storing customer data. Figure 2 shows the overall system S . In the following, l refers to the complete system and l_i with $i \in \{1, 2, 3\}$ to its subsystems as indicated in Figure 2. Applying definitions 1, 2a, 2b, and 4b, we yield the following trustworthiness terms with respect to the secrecy requirement:

- A : $\underbrace{((2 \otimes 1) \text{ out of } (\{A_1, A_2\}, \{A_3, A_4\}))}_{=:l_1}$ (def. 2a)
- B : $\underbrace{((1 \otimes 2) \text{ out of } (\{B_1, B_2\}, \{B_3, B_4\}))}_{=:l_2}$ (def. 2b)
- C : $\underbrace{(1 \text{ out of } \{C_1, C_2\})}_{=:l_3}$ (def. 1)
- S : $((1 \otimes 1) \text{ out of } (\{l_1, l_2\}, \{l_2, l_3\}))$ (def. 4b)

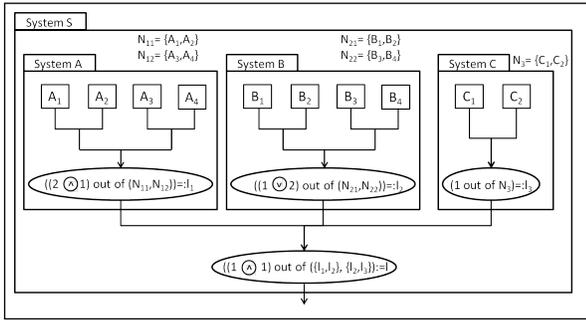


Figure 2: Example of the inductive determination of trustworthiness terms

5. PROPOSITIONAL LOGIC TRUSTWORTHINESS TERMS

As example 1 shows, trustworthiness terms can become complex, even for small systems. In order to yield representations that are comfortable to interpret for persons and appropriate for the computation of the probability with which a system fulfills a specific requirement r , we transform trustworthiness terms into propositional logic formulas. Particularly useful is the subsequent transformation of formulas into semantically equivalent formulas in normal form, such as the disjunctive normal form (DNF) or the conjunctive normal form (CNF). These normal forms show different strengths: while the CNF allows to determine “weak points”, such as single points of failure, the DNF is useful for identifying “strong points”, such as components or subsystems where trustworthiness results in the trustworthiness of the overall system, regardless of the trustworthiness of other components and subsystems. Thus, both normal forms should be applied complementarily. Due to limited space, we decided to use only CNF in this paper.

THEOREM 5.1. *Let system S consist of basic components $A = \{A_1, \dots, A_n\}$, and let $\{X_{A_1}, \dots, X_{A_n}\}$ be literals with $X_{A_i} = \text{true} \forall i$, if A_i is trustworthy wrts. r . Then, the trustworthiness term l of S can be mapped on a propositional logic formula $f(l)$ such that S is trustworthy wrts. r if and only if $f(l)$ is true. (For the proof see Appendix A.)*

We use the example shown in Figure 2 to illustrate how to determine the propositional logic formula of particular trustworthiness terms, namely l_1, l_2, l_3 , and l .

Example 2.

- $l_1 = ((2 \otimes 1) \text{ out of } (\{A_1, A_2\}, \{A_3, A_4\}))$
 $\Rightarrow f(l_1) \stackrel{(10)}{=} (f((2 \text{ out of } \{A_1, A_2\}))) \wedge$
 $(f((1 \text{ out of } \{A_3, A_4\})))$
 $\stackrel{(9)}{=} ((A_1 \wedge A_2)) \wedge ((A_3) \vee (A_4)) = A_1 \wedge$
 $A_2 \wedge (A_3 \vee A_4) =: f_A$
- $l_2 = ((1 \otimes 2) \text{ out of } (\{B_1, B_2\}, \{B_3, B_4\}))$
 $\Rightarrow f(l_2) \stackrel{(11)}{=} (f((1 \text{ out of } \{B_1, B_2\}))) \vee$
 $(f((2 \text{ out of } \{B_3, B_4\})))$
 $\stackrel{(9)}{=} ((B_1 \vee B_2)) \vee ((B_3) \wedge B_4)$
 $= B_1 \vee B_2 \vee (B_3 \wedge B_4) =: f_B$

- $l_3 = (1 \text{ out of } \{C_1, C_2\})$
 $\Rightarrow f(l_3) \stackrel{(9)}{=} (C_1) \vee (C_2) = C_1 \vee C_2 =: f_C$
- $l = ((1 \otimes 1) \text{ out of } (\{l_1, l_2\}, \{l_2, l_3\}))$
 $\Rightarrow f(l) \stackrel{(14)}{=} (f((1 \text{ out of } \{l_1, l_2\}))) \vee$
 $(f((2 \text{ out of } \{l_2, l_3\})))$
 $\stackrel{(12)}{=} (((f(l_1))) \vee ((f(l_2)))) \vee (((f(l_2))) \vee$
 $((f(l_3))))$
 $= (f(l_1)) \vee (f(l_2)) \vee (f(l_3))$
 $= (f_A) \vee (f_B) \vee (f_C)$
 $= (A_1 \wedge A_2 \wedge (A_3 \vee A_4)) \vee$
 $(B_1 \vee B_2 \vee (B_3 \wedge B_4)) \vee (C_1 \vee C_2)$

Finally, we convert the resulting propositional logic term given in (1) into CNF.

$$\bigwedge_{\substack{X \in \{A_1, A_2, A_3\} \\ Y \in \{A_1, A_2, A_4\} \\ Z \in \{B_3, B_4\}}} (X \vee Y \vee B_1 \vee B_2 \vee Z \vee C_1 \vee C_2) \quad (2)$$

(2) can be easily derived when we first transform (1) into DNF, given by

$$(A_1 \wedge A_2 \wedge A_3) \vee (A_1 \wedge A_2 \wedge A_4) \vee B_1 \vee B_2 \vee (B_3 \wedge B_4) \vee C_1 \vee C_2$$

The CNF formula given in (2) reveals that system S can be trustworthy with respect to secrecy requirement r if at least one of the components B_1, B_2, C_1, C_2 is trustworthy wrts. r , which is a sufficient, but not necessary condition. While propositional logic terms are useful to describe requirements on the trustworthiness of distributed systems, they do not provide a metric because they are unordered (and at nominal scale level only). In order to quantitatively assess propositional logic terms, we need to assign (trust) values to them. A straightforward approach is to draw on trustworthiness terms of the system components, which are represented by literals, and then to aggregate the attached trust values according to the way the literals are arranged in the propositional logic term.² The following section proposes a probabilistic approach towards using and aggregating trust values.

6. TRUST METRICS

In order to determine trust in the complete distributed system, i.e., the subjective probability with which the complete distributed system fulfills a particular security requirement r , we need to consider the trustworthiness of each component regarding r as well as the dependencies and requirements which are expressed in the trustworthiness terms. In the following, we assume that trust values of the atomic components regarding security requirement r are given. In the simplest case, the probability that each atomic component fulfills the security requirement r can be derived from historical data only. However, in order to take into account

²As trust is related to a particular requirement, we get for each component and (sub)system as many trust values as requirements exist.

uncertainty when modeling these probabilities, we propose to use the models presented in [16, 26]. Besides expressing uncertainty those models build on a Bayesian approach for deriving subjective probabilities from subjective knowledge and past experience³.

Regardless of the particular model for determining atomic trust values, the advantage of having propositional logic terms available is that we only need to define how the probabilities of the respective atomic trust values are aggregated depending on the logical operators \vee, \wedge, \neg .

According to probability theory, we regard each atomic formula/component L as a Bernoulli random variable X_L with $E(X_L) = p$, where p is the probability that component L and the respective service provider fulfill a particular security requirement. As the trust concept adopted in this paper also considers attacks due to malicious cooperation of entities, we draw on joint distributions of malicious cooperation. In addition, it seems reasonable to assume that the bias of entities to maliciously cooperate with other entities depends on who or what the other components are. Thus, we account for joint distributions of stochastically dependent variables, and we consequently apply conditional probabilities⁴. We define

$$g : \text{PLT} = \{\text{propositional logic terms}\} \rightarrow [0, 1] :$$

If $F_1, F_2 \in \text{PLT}$, then

$$g(\neg F_1) := 1 - g(F_1) \quad (3)$$

$$\begin{aligned} g(F_1 \wedge F_2) &:= g(F_2|F_1) \cdot g(F_1), \\ g(F_2|F_1) &:= P(X_{F_2}|X_{F_1}) \end{aligned} \quad (4)$$

$$\begin{aligned} (2), (3) \Rightarrow g(F_1 \vee F_2) &= g(\neg(\neg F_1 \wedge \neg F_2)) \\ &= 1 - g(\neg F_1 \wedge \neg F_2) \\ &= 1 - [g(\neg F_2|\neg F_1) \cdot g(\neg F_1)] \\ &= 1 - [(1 - g(F_2|\neg F_1))(1 - g(F_1))] \end{aligned} \quad (5)$$

If a propositional logic term F is in CNF, i.e. $F = F_1 \wedge \dots \wedge F_n$, $F_i = F_{i,1} \vee \dots \vee F_{i,n_i}$, $i = 1, \dots, n$, then the application of equations (3)-(5) yields

$$\begin{aligned} g(F) &= g(F_n|F_1 \wedge \dots \wedge F_{n-1}) \cdot g(F_1 \wedge \dots \wedge F_{n-1}), \quad n \geq 2 \\ g(F_i) &= 1 - [(1 - g(F_{i,n_i}|\neg(F_{i,1} \vee \dots \vee F_{i,n_i-1}))) \\ &\quad (1 - g(F_{i,1} \vee \dots \vee F_{i,n_i-1}))], \quad n_i \geq 2 \forall i \end{aligned} \quad (6)$$

When the trustworthiness term of a system S is transformed into the respective CNF, then the application of equation (6) allows for the recursive computation of the probability with which S is assumed to fulfill requirement r .

Applying the aforementioned evaluation steps results in a trustworthiness term, propositional logic formulas, and a trust value for each requirement r . In order to assess and potentially compare distributed systems with regard to a particular security requirement r , we need to address two challenges: First, in the presence of multiple requirements we get a multi-criteria decision problem, where trade-offs

³Both approaches can be applied for experiments with binary outcomes.

⁴See Eq. (4) where $P(X_{F_2}|X_{F_1})$ is a conditional probability.

between security requirements need to be addressed. Second, assessing and comparing distributed systems wrts. r only in terms of their probabilities assumes that the decision maker does not take into account how the probability was computed and how the architecture of the overall system looks like. For example, an overall probability 0.75 can result from a) one basic component where the probability of meeting requirement r is 0.75, or b) from a system with two components A, B where at least one component need to fulfill requirement r (trustworthiness term $= (1 \text{ out of } \{A, B\})$) in order to trust system S wrts. r and where the probability with regard to fulfilling r is 0.5 for both A and B . When comparing alternative systems with equal probability values, we suggest to also draw on trustworthiness terms and their respective CNF representations in order to rank these systems.⁵

As a propositional logic term in CNF is a conjunction of disjunctive clauses, we define redundancy-oriented preference relations ($\prec_{\text{red+}}, \prec_{\text{red-}}$) by drawing on the numbers of literals of disjunctive clauses:

Let F, G be two finite propositional logic terms in CNF, where in each term its' disjunctive clauses are arranged in ascending order of the numbers of their literals (e.g., $F = A \wedge (B \vee C) \wedge (A \vee B \vee D)$):

$$\begin{aligned} F &= F_1 \wedge \dots \wedge F_n, F_i = F_{i,1} \vee \dots \vee F_{i,p_i}, \\ &\quad i = 1, \dots, n; p_1 \leq \dots \leq p_n \\ G &= G_1 \wedge \dots \wedge G_m, G_j = G_{j,1} \vee \dots \vee G_{j,q_j}, \\ &\quad j = 1, \dots, m; q_1 \leq \dots \leq q_m \end{aligned}$$

Then

$$F \prec_{\text{red+}} G := \Leftrightarrow \exists k p_k < q_k \text{ and } p_l = q_l \forall 1 < l < k \quad (7)$$

$$F \prec_{\text{red-}} G := \Leftrightarrow \exists k p_k > q_k \text{ and } p_l = q_l \forall 1 < l < k \quad (8)$$

7. DISCUSSION AND CONCLUSION

This paper presents a novel formal approach towards the trust assessment of distributed systems. The approach includes the specification of trustworthiness terms and the proposition of quantitative trust metrics, which allow to assess and compare systems. Especially, the provided concept provides a means for deriving trust in a complex system from trust in its components and subsystems. The proposed approach is especially capable of dealing with the dependencies between the components and subsystems and their redundancy with respect to the security requirements under consideration. Thus, having formal descriptions and trust figures available, designers of distributed systems and customers of distributed services have quantitative information at hand that allow them to determine the impact of selecting specific services or specific system designs based on their trust perceptions. A key advantage of our approach lies in its generality and flexibility: it is scalable with regard to the granularity and type of atomic components (persons, PCs, machines, organizational units, services etc.), and capable of covering systems that go beyond boundaries of organizations.

⁵While our ranking procedure applies two criteria (probability and structure) sequentially, an alternative approach would be to assess systems in a multi-criteria sense by considering both criteria contemporaneously. Due to space limitation, we do not follow this approach here.

On the other hand, we need to discuss assumptions and limitations of our approach. First, we have to accept that any metric simplifies any complex socio-technical situation down to numbers or partial orders [29]. Second, by adopting propositional logic we assume that trust in a (sub)system with regard to a particular security requirement r depends only on r -related trust in parts of the system. When interdependencies between different requirements occur (e.g., the availability of subsystem A affects the anonymity of system S), our model needs to be extended. One option is to apply first-order logic, where $A(r)$ means that system A meets requirement r . Third, our evaluation is based on a centralized perspective, i.e. the assessment of systems is done by a central party.

Beyond research on the aforementioned topics, we suggest the following paths for further research: (a) Currently, we have to address different requirements in separate expressions. In the future, trustworthiness terms and trust metrics need to be extended to enable trust statements for all addressed requirements at once. Research can draw upon a comprehensive set of methods proposed in the field of multi-criteria decision making [35]. This extension should be flexible enough to allow different weightings of requirements regarding different stakeholders. For instance, one might emphasize on privacy while the other one might emphasize on the quality of the service. (b) The application of our metrics can lead to the development of distributed architectures patterns [12], which would support the design of trustworthy architectures. Pattern have already proven to be useful in other fields, such as in software development and in security design. (c) The operators for \vee, \wedge, \neg , which are currently defined for probability values only, could be replaced by the operators defined for ‘subjective logic’ in [19]. Thus, one could benefit from modeling uncertainty not only when determining and modeling the trust values of the atomic components, but also when deriving the trust value of the complete system. (d) The extension of our approach on other kinds of uncertainty theories, such as fuzzy set theory, allows to deal with scenarios where no probabilities are available [40, 41]. (e) In the economic context, further work on the application of our metrics is useful: Given a formal relationship between investments and the increase of trust in a particular component or system, decision makers can build on our metrics to determine those components with the highest marginal benefit of security investments, i.e. decision makers know the optimal “points” of investment. Given a budget for security investments, decision makers can also use our models to determine the optimal resource allocation, i.e. to decide on how much should be spent for each component in order to maximize the level of trust in the distributed system.

8. REFERENCES

- [1] A. Atzeni and A. Liroy. Why to adopt a security metric? A brief survey. In *1st Workshop on Quality of Protection*, volume 23 of *Advances in Information Security*, pages 1–12. Springer, September 2005.
- [2] T. Beth, M. Borcharding, and B. Klein. Valuation of trust in open networks. In *Proc. 3rd European Symposium on Research in Computer Security – ESORICS ’94*, pages 3–18, 1994.
- [3] H. Billhardt, R. Hermoso, S. Ossowski, and R. Centeno. Trust-based service provider selection in open environments. In *SAC ’07: Proceedings of the 2007 ACM Symposium on Applied Computing*, pages 1375 – 1380. ACM Press, 2007.
- [4] R. Blakley and G. Kabatiansky. *Encyclopedia of Cryptography and Security*, chapter Secret Sharing Schemes, pages 544–545. Springer, 2005.
- [5] M. Blaze, J. Feigenbaum, and A. D. Keromytis. Keynote: Trust management for public-key infrastructures. In *Security Protocols Workshop*, pages 59–63, 1998.
- [6] M. Blaze, J. Feigenbaum, and J. Lacy. Decentralized trust management. In *Proceedings from the 17th Symposium on Security and Privacy*, page 164i; $\frac{1}{2}$ 173. IEEE Computer Society Press, 1996.
- [7] M. Blaze, J. Feigenbaum, and J. Lacy. Decentralized trust management. In *Proceedings of the 1996 IEEE Symposium on Security and Privacy (SP ’96)*, pages 164–173. IEEE Computer Society, 1996.
- [8] R. Böhme and F. Freiling. On metrics and measurements. In *Dependability Metrics*, pages 7–13. Springer, 2008.
- [9] R. Böhme and T. Nowey. Economic security metrics. In I. Eusgeld, F. C. Freiling, and R. Reussner, editors, *Dependability Metrics*, pages 176–187, 2008.
- [10] S. Buchegger and J.-Y. Le Boudec. A Robust Reputation System for Peer-to-Peer and Mobile Ad-hoc Networks. In *P2PEcon 2004*, 2004.
- [11] D. Gambett. *Trust: Making and Breaking Cooperative Relations*, chapter Can we Trust Trust?, pages 213–237. Basil Blackwell, New York, 1988.
- [12] E. Gamma, R. Helm, R. E. Johnson, and J. M. Vlissides. Design patterns: Abstraction and reuse of object-oriented design. In *ECOOP ’93: Proceedings of the 7th European Conference on Object-Oriented Programming*, pages 406–431. Springer, 1993.
- [13] T. Grandison and M. Sloman. A survey of trust in internet applications. *IEEE Communications Surveys and Tutorials*, 3(4):2–16, 2000.
- [14] D. Heimann and N. Mittal. Availability and reliability modeling for computer systems. *Advances in Computers*, 31:175–233, 1990.
- [15] T. Hofmeister, M. Krause, and H. Simon. Optimal k out of n secret sharing schemes in visual cryptography. *Theoretical Computer Science*, 240:471–485, 2000.
- [16] A. Jøsang. A logic for uncertain probabilities. *International Journal of Uncertainty, Fuzziness and Knowledge-Based Systems*, 9(3):279–212, 2001.
- [17] A. Jøsang, R. Ismail, and C. Boyd. A survey of trust and reputation systems for online service provision. *Decision Support Systems*, 43(2):618–644, 2007.
- [18] A. Jürgenson and J. Willemson. Computing exact outcomes of multi-parameter attack trees. In *Proceedings of the OTM Confederated International Conferences*, pages 1036 – 1051, 2008.
- [19] A. Jøsang and D. McAnally. Multiplication and comultiplication of beliefs. *International Journal of Approximate Reasoning*, 38(1):19–51, 2004.
- [20] S. D. Kamvar, M. T. Schlosser, and H. Garcia-Molina. The eigentrust algorithm for reputation management in P2P networks. In *Proc. of the 12th Int. Conf. on World Wide Web*, pages 640–651. ACM Press, 2003.
- [21] R. Kohlas, J. Jonczyk, and R. Haenni. A trust

evaluation method based on logic and probability theory. In *2nd Joint iTrust and PST Conferences on Privacy, Trust Management and Security*, volume II, pages 17–32, 2008.

- [22] K. Krukow and M. Nielsen. Trust structures: Denotational and operational semantics. *International Journal of Information Security*, 6(2-3):153–181, 2007.
- [23] U. Maurer. Modelling a public-key infrastructure. In *Computer Security — ESORICS 96*, pages 325–350. Springer, 1996.
- [24] NIST. Performance measurement guide for information security, May 2008. NIST Special Publication 800-55 Revision 1.
- [25] S. C. Payne. A guide to security metrics. Technical report, The SANS Institute, 2006.
- [26] S. Ries. Extending bayesian trust models regarding context-dependence and user friendly representation. In *Proceedings of the 2009 ACM Symposium on Applied Computing*. ACM Press, 2009.
- [27] S. Ries. *Trust in Ubiquitous Computing*. PhD thesis, Technische Universität Darmstadt, 2009.
- [28] S. Ries and A. Heinemann. Analyzing the robustness of CertainTrust. In *2nd Joint iTrust and PST Conference on Privacy, Trust Management and Security*, pages 51 – 67. Springer, 2008.
- [29] R. Savola. Towards a taxonomy for information security metrics. In *QoP '07: Proceedings of the 2007 ACM workshop on Quality of protection*, pages 28–30. ACM, 2007.
- [30] B. Schneier. Attack trees: Modeling security threats. *Dr. Dobbs' journal*, 24:21 – 29, 1999.
- [31] G. Schryen and E. Rich. Security in Large-Scale Internet Elections: A Retrospective Analysis of Elections in Estonia, The Netherlands, and Switzerland. *IEEE Transactions on Information Forensics & Security*, 4(4):729–744, 2009.
- [32] A. Shamir. How to share a secret. *Commun. ACM*, 22(11):612–613, 1979.
- [33] R. Sharman, R. Rao, and S. Upadhyaya. Metrics for information security - a literature review. In *Proceedings of Americas Conference on Information Systems*, pages 1436–1440, 2004.
- [34] W. T. L. Teacy, J. Patel, N. R. Jennings, and M. Luck. TRAVOS: Trust and reputation in the context of inaccurate information sources. *Autonomous Agents and Multi-Agent Systems*, 12(2):183–198, 2006.
- [35] E. Triantaphyllou. *Multi-Criteria Decision Making Methods: A comparative Study*, volume 44 of *Applied Optimization*. Springer, 2000.
- [36] R. Vaughn, R. Henning, and A. Siraj. Information assurance measures and metrics: State of practice and proposed taxonomy. In *Proceedings of 36th Hawaii International Conference on System Sciences*, 2003.
- [37] E. Verheul and H. van Tilborg. Constructions and properties of k out of n visual secret sharing schemes. *Designs, Codes and Cryptography*, 11(2):179–196, 1997.
- [38] M. Volkamer and R. Grimm. Determine the Resilience of Evaluated Internet Voting Systems. In *Proceedings of 1st RE-VOTE conference*, pages 47–54, 2009.
- [39] A. Whitby, A. Jøsang, and J. Indulska. Filtering out

unfair ratings in bayesian reputation systems. *The ICFAIN Journal of Management Research*, 4(2):48 – 64, 2005.

- [40] L. A. Zadeh. Fuzzy logic and approximate reasoning. *Synthese*, 30:407–428, 1975.
- [41] H.-J. Zimmermann. An application-oriented view of modelling uncertainty. *European Journal of Operational Research*, 122:190–198, 2000.

APPENDIX

A. PROOF FOR THEOREM 5.1

PROOF. We prove the theorem along the inductive definition of trustworthiness terms, and we provide for each definition of trustworthiness terms the corresponding propositional logic formula. The principal idea of the proof is that we reformulate the expression “ k out of a set L ” by explicitly considering all combinations of elements of L , where L can be either a set of basic components or of trustworthiness terms of subsystems. The provision of such a mapping f (of trustworthiness terms on propositional logic terms) proves the theorem.

- If $l = (k \text{ out of } N)$, $k \in \{1, \dots, |N|\}$, $N \subseteq A$ (def. 1), then

$$f(l) := \bigvee_{\substack{\{A_{i_1}, \dots, A_{i_k}\} \subseteq A \\ |\{A_{i_1}, \dots, A_{i_k}\}| = k}} \left(\bigwedge_{j=i_1}^{i_k} A_j \right) \quad (9)$$

- If $l = ((k_1 \otimes \dots \otimes k_m) \text{ out of } (N_1, \dots, N_m))$, $N_i \subseteq A \forall i$ (def. 2a), then

$$f(l) := \bigwedge_{i=1}^m (f((k_i \text{ out of } N_i))) \quad (10)$$

- If $l = ((k_1 \otimes \dots \otimes k_m) \text{ out of } (N_1, \dots, N_m))$, $N_i \subseteq A \forall i$ (def. 2b), then

$$f(l) := \bigvee_{i=1}^m (f((k_i \text{ out of } N_i))) \quad (11)$$

- If $l = (k \text{ out of } \{l_{i_1}, \dots, l_{i_m}\})$, l_{i_j} trustworthiness terms, $\{i_1, \dots, i_m\} \subseteq \{1, \dots, n\}$ (def. 3), then

$$f(l) := \bigvee_{\substack{\{j_1, \dots, j_k\} \subseteq \{i_1, \dots, i_m\} \\ |\{j_1, \dots, j_k\}| = k}} \left(\bigwedge_{j=j_1}^{j_k} (f(l_j)) \right) \quad (12)$$

- If $l = ((k_1 \otimes \dots \otimes k_m) \text{ out of } (Q_1, \dots, Q_m))$, Q_i set of trustworthiness terms (def. 4a), then

$$f(l) := \bigwedge_{i=1}^m (f((k_i \text{ out of } Q_i))) \quad (13)$$

- If $l = ((k_1 \otimes \dots \otimes k_m) \text{ out of } (Q_1, \dots, Q_m))$, Q_i set of trustworthiness terms (def. 4b), then

$$f(l) := \bigvee_{i=1}^m (f((k_i \text{ out of } Q_i))) \quad (14)$$

□