A Local Qualitative Approach to Referral and Functional Trust

Krishnaprasad Thirunarayan, Dharan K. Althuru, Cory A. Henson, and Amit P. Sheth

Kno.e.sis Center, Department of Computer Science and Engineering
Wright State University, Dayton, OH 45435
{t.k.prasad,althuru.2}@wright.edu
http://www.cs.wright.edu/˜tkprasad

Abstract. Trust and confidence are becoming key issues in diverse applications such as ecommerce, social networks, semantic sensor web, semantic web information retrieval systems, etc. Both humans and machines use some form of trust to make informed and reliable decisions before acting. In this work, we briefly review existing work on trust networks, pointing out some of its drawbacks. We then propose a local framework to explore two different kinds of trust among agents called referral trust and functional trust, that are modelled using local partial orders, to enable qualitative trust personalization. The proposed approach formalizes reasoning with trust, distinguishing between direct and inferred trust. It is also capable of dealing with general trust networks with cycles.

1 Introduction

Trust relationships occur naturally in many diverse areas and contexts such as ecommerce, social interactions/networks, semantic web information retrieval systems, distributed systems, decision-support systems, semantic sensor web, etc. As the connections and interactions between humans, machines, and other resources evolve, and as the agents providing content and services become increasingly removed from the agents that consume them, the issue of trust inference and management will become increasingly significant. Trust information is necessary for conflict resolution, and for making informed and reliable decisions before acting. Unfortunately, there is neither a universal notion of trust that is applicable to all domains nor a clear explication of its semantics in many situations.

Even if we confine ourselves to interactions among humans and/or machines (collectively called agents) in a narrow domain, there are still several fundamentally different notions of trust.

1.1 Motivating Example

Consider the following adaptation of examples from Josang et al [16] and its abstraction in Figure 1.

Alice may trust Bob for recommending a good car mechanic because of Bob's experiences with car problems. Bob may trust Dick to be a good car mechanic because of

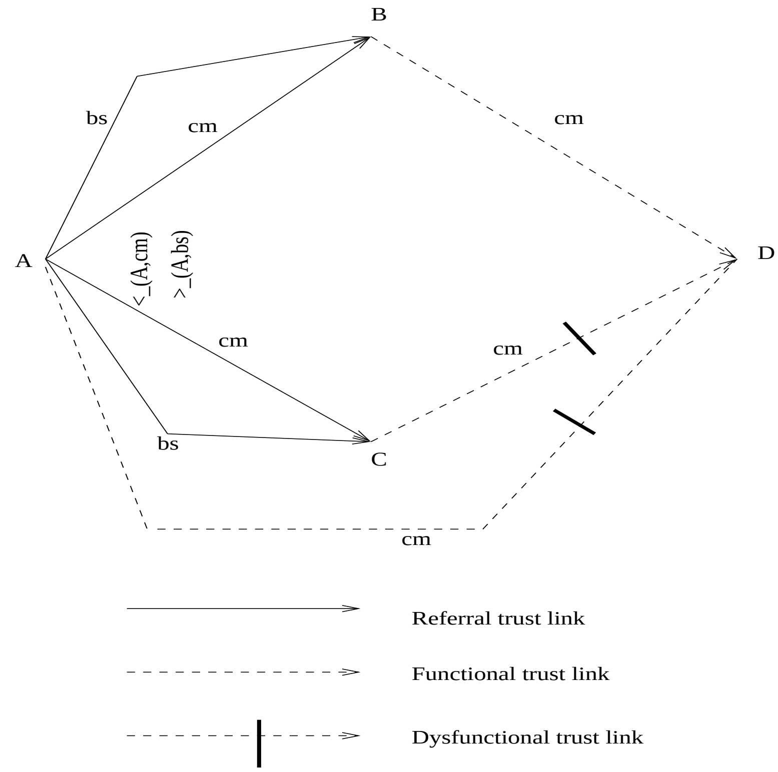

Fig. 1. Example DAG structured Trust Network

Bob's past experiences with Dick. On the basis of this, Alice may infer trust in Dick to be a good car mechanic upon recommendation from Bob.

Let us say that *trust scope* captures the domain/context/task/function for which the trust relationship is applicable. In general, an agent a1 may trust another agent a2 for agent a2's ability to provide good recommendations in a trust scope because agent a2 is knowlegeable in that trust scope. In this case, we say agent a1 has *referral trust* in agent a2 in the *trust scope*. Similarly, an agent a1 may trust another agent a2 for agent a2's ability to perform certain task/function in a trust scope. In this case, we say agent a1 has *functional trust* in agent a2 in the given *trust scope*.

Alice may also trust Charlie for recommending a good car mechanic, and Charlie may have a negative recommendation of Dick.

An agent a1 may distrust another agent a2 because of agent a2's inability to perform certain task/function in a trust scope. In this case, we say agent a1 has *nonfunctional trust* in agent a2 in the given *trust scope*.

Furthermore, Alice may trust Bob over Charlie for recommending a good car mechanic, possibly because of Bob's extensive experience with car problems. So, even if Bob and Charlie provide conflicting recommendations on Dick, Alice may prefer Bob's recommendations over Charlie's.

Thus, for a given *trust scope*, each agent a0 may have differential referral trust among its neighbors a1, a1 , a3, . . . , which can be formalized using a local partial ordering relationship among neighbors of agent a0. Recall that the partial order enables us to model incomparable trust, that is, it is not necessary to be able to, or to force, a linear total order on neighboring agents with respect to trust. This ordering can be used to resolve conflicts, minimizing ambiguity. Note that presence of ambiguity requires further investigation for resolution in order to permit decision or action.

Subsequently, Alice's direct experience with Dick's incompetence as a car mechanic may lead Alice to distrust Dick, irrespective of the recommendations of Bob and Charlie.

In general, an explicit functional (nonfunctional) trust link between agents a1 and a2 always *overrides* conflicting trust inference paths involving a sequence of referral trust links terminated by a nonfunctional (functional) trust link. In other words, a claim supported by trust link embodying strict knowledge always overrides a counter claim sanctioned by longer trust paths embodying defeasible knowledge. Thus, our approach can use local partial ordering and overriding to enable trust personalization.

Alice can have referral trust in Eric due to his knowledge about good car mechanics without having a functional trust in Eric about his being a good car mechanic. Similarly, Alice can have functional trust in Eric about being a good car mechanic without having a referral trust in Eric, possibly due to conflict of interest and competitive spirit that may lead Eric to be less than candid.

Thus, we permit agent a1 to have referral (functional) trust in agent a2 for a trust scope without having a functional (referral) trust in agent a2 for the same trust scope. That is, referral trust and functional trust are not coupled in general.

Even though Alice trusts Bob over Charlie for recommending a good car mechanic, Alice may trust Charlie over Bob for recommending a good baby sitter.

Effectively, the local partial ordering relationship on referral trust among neighboring agents a1, a1 , a3, . . . , of agent a0 can depend on *trust scope*, and can be different for different trust scopes.

Furthermore, our proposal enables setting *majority thresholds* for deriving functional/nonfunctional trust conclusions (for example, a majority functional trust threshold of 4 to model requiring 4 out of 5 stars, or 4 positive referrals for every negative referral on `amazon.com`), and remaining ambivalent if the thresholds are not crossed.

In Section 2, we provide brief background on the structure of trust values, and review existing work on trust networks, pointing out some of their shortcomings. In Section 3, we investigate our novel approach to representing and determining agent trust by combining and abstracting trust information that we claim is natural, efficient and effective. We propose a *local* framework to explore referral trust and functional trust among agents. In order to resolve conflicts and aggregate functional trust, we model relative differences in the level of trust among agents using local partial orders. We formally specify reasoning with referral trust and functional trust, distinguishing between direct and inferred trust. The approach is then generalized to deal with cycles in trust networks. This paper is not concerned with the domain/situation-specific issues that influence the acquisition of partial orderings, however. In Section 4, we conclude with suggestions for future work.

2 Background: The Structure of Trust

We discuss trust networks and existing trust models, to better situate our work in relation to other works in the literature. For completeness, we also recapitulate some of the motivation for our approach and for the proposed local framework from Thirunarayan and Verma [25].

Traditional approaches to formalizing trust between a pair of agents models trust as a real number in the closed interval [0,1]. Even though this facilitates trust computation, such as via aggregation and propagation, there are inherent difficulties in coming up with initial trust values and semantically justifying computed trust values. To paraphrase Guha et al [12]: *While continuous-valued trusts are mathematically clean [21], from the standpoint of usability, most real-world systems will in fact use discrete values at which one user can rate another.* Furthermore, it is not unreasonable to expect and allow users to specify relative trust information.

We explore "realistic" models of trust based on partially ordered discrete values. Our approach differs from popular works (such as [12, 18, 21, 16]) as follows:

- We distinguish both referral trust and (non)functional trust among agents implicitly as *discrete* values.
- Our approach is sensitive to *local, relative* ordering of trust values rather than their magnitudes.
- We distinguish between *direct* trust and *inferred* trust, letting direct information override conflicting inferred information.
- We regard *equal* or *incomparable* evidence in support and against functional trust in an agent as ambiguous trust, and represent the ambiguity explicitly.

We believe that our approach provides a natural and robust representation of *relative* trust information that an agent (aggregator) has. For instance, trust based on direct knowledge is superior to trust based on a stamp of approval from a certifying agency. However, it may not always be possible or desirable to impose arbitrary total order on trust values associated with agents in all contexts[1]. In such situations, our work enables representation of ambiguity as opposed to requiring one to break the tie. Note also that, in practice, trust relationships can change over time[2] as new information arrives, causing *nonmonotonic* changes to inferred trust information.

On `Epinions.com` [8], users can add other users to their "Web of Trust", i.e., reviewers whose reviews and ratings they have consistently found to be valuable and to their "Block list", i.e., authors whose reviews they find consistently offensive, inaccurate, or in general not valuable. Trust and distrust are materialized as 1 and -1. Richardson et al [21] start with `Epinions` user trust graph, *synthetically* generate

[1] However, note that a total order consistent with a partial order can always be generated.

[2] Tagging trust values with time stamps and the dynamic evolution of trust over a period of time is beyond the scope of the current work.

real-valued user trust values and statement belief information using user quality parameter and user reviews data, to study the relationship between user quality and trust propagation. Massa and Hayes [18] make a case for distinguishing (referential) hyperlinks into two categories: positive endorsement links and negative criticizing links. PageRank algorithm [4] is run on `Epinions` user trust graph with various combination of trust and distrust links, to analyze the effect of added expressiveness on user rankings. Guha et al [12] encode trust and distrust information as 1 and -1, and define four different atomic operations for propagating trust: direct propagation, co-citation, transpose trust and trust-coupling. These operations are captured via matrix operations. Their framework uses real-valued trust, and final trust/distrust values are determined using finite number of iterations (finite length paths) and thresholds for rounding. Zeigler and Lausen [31] presents a classification of trust metrics to evaluate "transitivity of trust through social networks". They discuss orthogonal issues such as privacy, scalability, local vs distributed nature, trust flow, etc. Wang and Singh [28, 29] propose a probability of probabilities approach to acyclic trust that represents uncertainty information explicitly and belief/trust as real numbers between 0 and 1. They study the formal properties of concatenation and aggregation operators.

Artz and Gil [1] surveys existing models of trust, different definitions of trust, trust metrics, and their specific determination using policies or reputation. Massa and Avesani [18] distinguish between Global ("one-size-fits-all") and Local ("personalized") Trust Metrics, and propagate trust over a limited length trust paths. They show the benefits of local trust metrics for controversial users (those who are trusted and distrusted by sizable population). The computed trust between a trustor and a trustee can potentially be effected by the numeric trust values on all trust links between the trustor and the trustee. In contrast, in our approach, claims supported via high trust links *override* conflicting claims via less reliable links, for forming trust conclusions. Massa and Avesani [20] analyze the variation in average trust values for different equivalence classes of users, determined on the basis of path length.

Golbeck and Hendler [11] describe a more sophisticated approach to locally inferring trust in web-based social networks that explicitly represents both trust and no trust on a fixed linear scale obtained from context-based ratings, and aggregates trusts from neighbors via weighted averaging. Our trust networks are more expressive, but on common networks, our approach resembles the rounding algorithm of Golbeck and Hendler [11]. In [17], Katz and Golbeck compute a partially ordered priority relationship among competing defaults using the trust computations described in Golbeck and Hendler [11]. They do not show how to leverage the partial order itself to get a new framework for computing trust values as we do in this paper.

In the approach of Bintzios et al [3], the trust in agent n by an agent u is determined by using the trust (weighted links) of the direct neighbors of agent n, weighted by the corresponding trust associated with the neighbors of agent n by agent u (weights associated with potentially indirect trust paths to neighbors of agent n from agent u) using path algebra operators such as maximum and multiplication.

Even though Golbeck [10], Bintzios et al [3] and our work can be applied to a common subset of trust networks, quantitative and qualitative comparisons to trust inference algorithms in Golbeck [10] and Bintzios et al [3] are hard to make because both present frameworks and algorithms that contain user tunable parameters and aggregation functions. As such, on a given example, we can always reverse engineer their parameters and make them agree or disagree with our definite semantics. In spite of such difficulties, we will present simple illustrative examples that highlight fundamental differences between these approaches. According to TidalTrust algorithm (with trust threshold = 6 and keeping in view the shortest path) [10], Figure 2 supports the conclusion that agent A trusts agent E. In contrast, our approach concludes that agent A distrusts (nonfunctional trust) agent E assuming that A (referral) trusts B more than A (referrals) trusts C. Similarly, according to [3], the top portion of Figure 3 can be interpreted as supporting the fact that agent A distrusts agent D, and the bottom portion of Figure 3 can be interpreted as both supporting or not supporting agent A trusts agent D, depending on whether MAX or MIN is chosen as the aggregation function. In contrast, according to our approach, in the former case, agent A referral/functional trusts agent D, while in the latter case, agent A is ambiguous about functional trust in agent D, assuming that the trust threshold is 0.5.

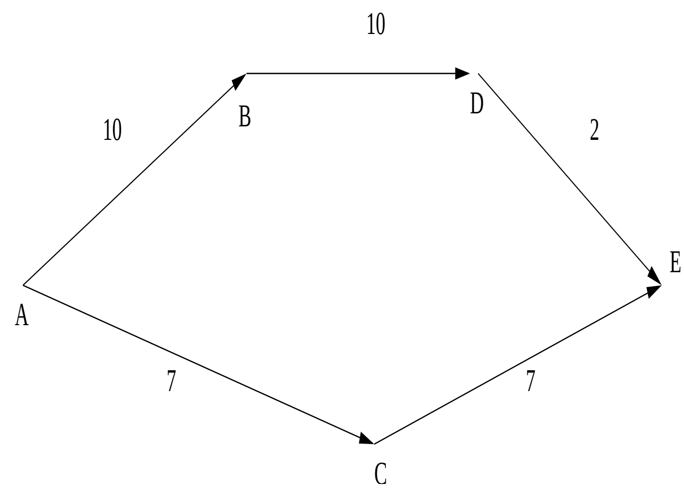

Fig. 2. Comparison with Golbeck's approach

Josang et al [16] presents a novel approach to trust that discriminates between two different forms of trust and reasoning with them. In fact, we have borrowed their terminology of referral trust, functional trust, and trust scope. Josang et al's approach transforms and filters a trust network into a canonical form (called Directed Series Parallel Graph), using heuristics. They use subjective logic to reason with opinions represented as 4-tuples (belief, disbelief, uncertainty, base rate). Even though this provides an expressive representational framework, it is unclear how quantitative opinions are obtained initially and what semantics to associate with "fine-grain" numeric values, compared to acquiring and reasoning with relative, binary trust information.

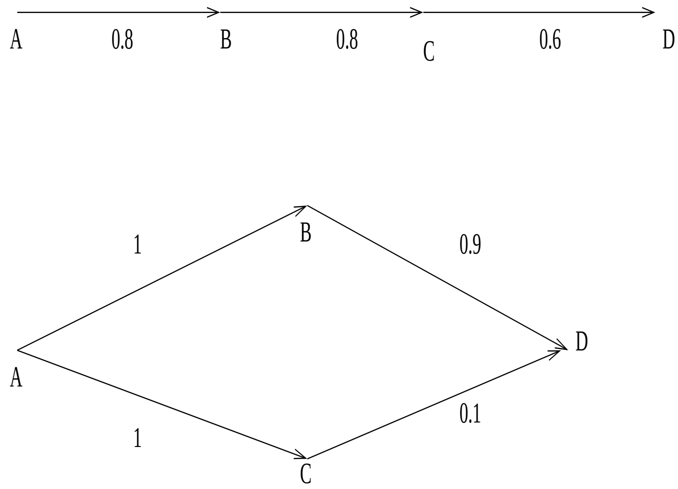

Fig. 3. Comparison with Bintzio's approach

Huang and Fox [15] introduce two different types of trust: *trust in belief* and *trust in performance* that roughly parallels referral trust and functional trust respectively. Within a specific context, the former is transitive, while the latter is not. The notion of *context* is analoguous to trust scope. They formalize the ontology of trust using situation calculus, distinguishing between *direct trust* and *indirect (social networks based) trust* and explicitly encoding trust formation rules. In Section 3, we capture the semantics of trust using a set-based, model-theoretic approach. Essentially, direct trust is supported by an explicit trust link, while indirect trust is supported by "unpreempted" paths in the trust network. (We do not separate out system trust (due to a certifying agency or standards body) from other forms of trust.)

In relation to our current work, Wang et al [30] discuss three orthogonal aspects pertaining to trust networks: (1) two ways of determining referral trust — one based on *similarity* and another based on *truthfulness* between agents; (2) characterization of small world networks; and (3) approachs to encourage/remunerate agents to share recommendations and ratings.

Both the current approach, and the earlier approach to trust-distrust-belief networks in Thirunarayan and Verma [25], use local framework. However, they differ significantly along the following lines: (i) The current approach ignores statements and beliefs, and refines trust by distinguishing between referral, functional and nonfunctional trust. (ii) In the current approach, an agent trusts another agent only in certain clearly delineated trust scopes. Furthermore, the local partial order associated with an agent and its neighborhood is specific to a trust scope.

Carroll et al [5, 6] discuss extensions to RDF to enable incorporation of trust and provenance information into Semantic Web. They also present different kinds of trust mechanisms (reputation-based, context-based and content-based) and several example trust policies that shed light on the nature of trust and that can be used to codify trust rules.

Schenk [22] proposes a framework for reasoning with partially ordered trust levels based on Belnap's bilattices [2], to provide additional information on the reliability of results to a user. Their focus is on developing bilattices-based semantics for logic programs and OWL. In contrast, our work is tailored for trust networks that uses the 4-values to capture the nature of the final conclusions and partial order information to resolve conflicts.

In summary, our work develops a computational model of referral trust and (non)functional trust among agents that abstracts weights on links through local partial ordering on links in trust scopes, and propagates it via local distributed computation. It is focussed on qualitative information that is natural and more readily available than on quantitative information. The binary approach also permits derivation of conclusions that permit taking action or decision. Our approach is robust with respect to redundant links obtained by replacing a node with a pair of synonymously named connected nodes. The discretization of trust values, trust scope dependent partial ordering, and trust aggregation via least-upper bound operation, enables us to readily see the semantic consequences of the trust network and the computational properties such as locality, convergence, etc. Eventually, such trust networks can be standardized using semantic web techniques and technologies, and used for information/resource retrieval.

3 Trust Networks and Their Semantics

We investigate representation and reasoning with different forms of trust among agents. We borrow the terms referral trust, functional trust and trust scope from Josang *et al* [16]. To us, *trust scope* captures the domain/context/task/function for which the trust relationship is applicable; *referral trust* captures trust in someone for "educated" recommendations; *(non)functional trust* captures (dis)trust in someone to carry out a task/function properly. In order to formalize these notions, we introduce trust networks as a graph containing agent nodes, connected by referral, functional, or nonfunctional [3] trust links.

The semantics of trust networks can be captured by formalizing trust relationship among agents and legitimate inferences from them, using local neighborhood for each agent node in the network. For instance, if agent a1 trusts agent a2 for referral in trust scope ts, and agent a2 trusts agent a3 for referral in the same trust scope, then one can infer that agent a1 trusts agent a3 for referrals in trust scope ts (that is, transitivity holds w.r.t. a trust scope). (This is analogous to someone seeking recommendations from their neighbors/friends about an electrician or a plumber that the neighbors/friends have heard of.) Similarly, if agent a1 trusts agent a2 for referral in trust scope ts, and agent a2 (dis)trusts agent a3 for carrying out the function in trust scope ts, then agent a1 has (non)functional trust in agent a3 by referral. (This is analogous to someone seeking recommendations from their neighbors/friends about an electrician or a plumber that the latter have experienced.) However, if agent a1 trusts agent a2 more than agent a1 trusts agent a3 in trust scope ts, and agent a2 has functional trust in agent a4, and agent a3 has nonfunctional trust in agent a4, then one can infer that agent a1 has functional

[3] For want of a better prefix to capture "incompetence"!

trust in agent a4 in trust scope ts. Furthermore, the local approach to trust formation can take into account *majority* opinion among the *most trusted* neighbors by counting support and dissent for functional trust in agent a4. (For instance, consider how we process votes when purchasing products from online vendors such as `Amazon.com`, choosing to buy a particular item only if it has a rating of 4 or 5 stars, or if at least 80% of the reviews are positive.) Note that, we interpret direct (non)functional links (paths of length 1) as *strict* while paths of lengths 2 or more as *defeasible*, overridden by stronger conflicting links.

There are many interesting topological and conceptual similarities between boolean trust networks and mixed inheritance networks [13, 14, 23, 24, 26, 27]. Both kinds of graphs have links that can be positive or negative, with potential for conflicts, ambiguity, and conflict resolution. For instance, referral trust links resemble instance/class to class links and functional trust links resemble instance/class to property links. Thus, we can explore adapting existing *local* theories of inheritance (that is, the semantics of a node is completely determined by the semantics of its neighboring nodes and the connecting links [24]) for reasoning with trust.

We formalize various intuitions about trust aggregation in terms of paths in the trust network as follows.

Definition 1 *A* trust network *is an ordered labelled graph G = (AN, RL, PFL, NFL, TS, TSF) containing agent nodes AN, trust links ($\subset AN \times AN$) of three types: (i) referral trust links RL, (ii) positive functional trust links PFL, and (iii) negative functional (non-functional) trust links NFL, trust scopes TS, and trust scopes labelling function TSF ($(RL \cup PFL \cup NFL) \rightarrow$ Powerset(TS)). Furthermore, for each agent node a_i in AN and for trust scope ts, a local referral trust ordering relation $\prec_{(a_i,ts)}$ on a_i's out-links (that is, $\{\ (a_i, a_j)\ \mid\ (a_i, a_j) \in$ RL$\}$). For completeness, PFL $\cap$ NFL = $\emptyset$, and $\prec$ is irreflexive/strict partial order.*

We may abbreviate agent node as *agent*, referral trust link as *referral link*, positive functional trust link as *functional link*, and negative functional trust link as *non-functional link*, when there is no ambiguity. Given a link (a_i, a_j), we say that a_i is a *predecessor* of a_j and a_j is a *successor* of a_i.

In order to develop formal semantics and efficient (one-pass, linear) computation procedure, we initially restrict the subgraph spanned by these links to be directed-acyclic graph (DAG). Subsequently, we relax this restriction, allowing cycles. Unfortunately, as explained later, this expressiveness brings with it additional computational complexity, in the worst case.

We model referral trust function $\mathcal{R}$ and functional trust function $\mathcal{F}$ *supported by* the trust network as:

$\mathcal{R} : AN \times AN \rightarrow Powerset(TS)$ and

$\mathcal{F} : AN \times AN \rightarrow Powerset(TS \times \{\perp, true, false, \top\})$.

The values $\perp, true, false$ and $\top$ correspond to no information, supporting information, opposing information, and ambiguous information respectively. These four values can be partially ordered on information-content scale, similarly to Belnap's 4-valued logic [2]: $\perp < true, \perp < false, true < \top$, and $false < \top$. Also, let [V1 $<$ V2 iff V2 $>$ V1], and [V1 $\geq$ V2 iff (V1 $>$ V2) or (V1 $=$ V2)].

We provide semantics of trust aggregation by defining when agent ai *can referral trust* agent aj and when agent ai *can (non)functional trust* agent aj in trust scope ts. The former semantics is "existence-based", while latter semantics is "majority-based". We define referral and functional trust functions $\mathcal{R}$ and $\mathcal{F}$ in each case as follows. (Actual detailed definitions of "*can referral trust*" and "*can (non)functional trust*" are given later.)

Reflexivity: Agents have referral trust and functional trust in themselves.

$\forall a \in$ AN: $\mathcal{R}(a, a) \;=\; TS$

$\forall a \in$ AN: $\mathcal{F}(a, a) \;=\; TS \times \{true\}$

Referral Trust-related: $\forall a_i, a_j \in$ AN:

$ts \;\in\; \mathcal{R}(a_i, a_j)$ if (a_i can referral trust a_j in trust scope ts)

(Non)Functional Trust-related: $\forall a_i, a_j \in$ AN:

Positive: $(ts, true) \;\in\; \mathcal{F}(a_i, a_j)$ if

(a_i can functional trust a_j in trust scope ts)

Negative: $(ts, false) \;\in\; \mathcal{F}(a_i, a_j)$ if

(a_i can nonfunctional trust a_j in trust scope ts)

Ambivalence: $(ts, \top) \;\in\; \mathcal{F}(a_i, a_j)$ if

(a_i can ambiguous trust a_j in trust scope ts)

Ignorance: $(ts, \bot) \;\in\; \mathcal{F}(a_i, a_j), \;\; \text{otherwise}$

The rationale is that if the trust is well-defined by majority then there is no harm in subscribing to it but when there is some doubt due to conflicting evidence, it is better to note the ambiguity for further investigation. "No information" is not the same as "ambiguous information". Note also that, as specified below, if there is direct positive (resp. negative) trust link from a_i to a_j then $\mathcal{F}(a_i, a_j) \;=\; true$ (resp. $\mathcal{F}(a_i, a_j) \;=\; false$), irrespective of any other information.

3.1 Majority-based Semantics for DAGs

We specify how trust can be propagated *top-down* through the DAG-structured trust networks. This approach takes into account both the polarity and the cardinality of the appropriate links. Note that $\| \ldots \|$ stands for set-cardinality operator.

Evidence in support of Referral Trust: a_i can referral trust a_j in trust scope ts if there is an explicit trust link from a_i to a_j, or there is a successor a_k of a_i that referral trusts a_j in trust scope ts.

$$\forall a_i, a_j \in AN : a_i \textbf{ can referral trust } a_j \textbf{ in trust scope } ts \;\text{ if}$$

$$[\,(a_i, a_j) \in \text{RL} \;\wedge\; ts \;\in\; TSF(a_i, a_j)\,] \;\;\vee$$

$$[\exists a_k \in \text{AN}: \;\; (a_i, a_k) \in \text{RL} \;\;\wedge\;\; ts \;\in\; TSF(a_i, a_k) \;\;\wedge\;\; ts \;\in\; \mathcal{R}(a_k, a_j)\,]$$

Ignorance about Functional Trust: The fact that a_i has no information to (non)functional trust a_j in trust scope ts can be specified as follows. (Note that the auxiliary definitions below are more general than necessary only because these will be reused later. They imply that there exists no suitable paths, which is sufficient.) Recall that $\| \ldots \|$ stands for set-cardinality operator.

$$\forall a_i, a_j \in AN : a_i \textbf{ is ignorant about } a_j \textbf{ in trust scope } ts \;\text{ if}$$

$$[\;\; ts \notin TSF(a_i, a_j) \;\;\vee\;\; (a_i, a_j) \notin \mathrm{PFL} \cup \mathrm{NFL} \;\;] \;\wedge$$

$$\|\text{ Undefeated functional trust of } a_i \textbf{ in } a_j \textbf{ for } ts \;\| \;\; = \; 0 \quad \wedge$$

$$\|\text{ Undefeated nonfunctional trust of } a_i \textbf{ in } a_j \textbf{ for } ts \;\| \; = \;\; 0$$

where

$$\|\text{ Undefeated functional trust of } a_i \textbf{ in } a_j \textbf{ for } ts \;\| \; =$$

$$\| \; \{(a_i, a_k) \in \mathrm{RL} \;\;|\;\; ts \in TSF(a_i, a_k) \;\;\wedge\;\; (ts, true) \;\in\; \mathcal{F}(a_k, a_j)$$

$$\wedge \;\neg\; \exists a_l \in \mathrm{AN} : \;\; (a_k \prec_{(a_i, ts)} a_l)$$

$$\wedge \; (a_i, a_l) \in \mathrm{RL} \;\wedge\; ts \in TSF(a_i, a_k) \;\wedge\; (ts, bf) \;\in\; \mathcal{F}(a_l, a_j)\} \wedge bf \geq false \;\; \|$$

and

$$\|\text{ Undefeated nonfunctional trust of } a_i \textbf{ in } a_j \textbf{ for } ts \;\| \; =$$

$$\| \; \{(a_i, a_k) \in \mathrm{RL} \;\;|\;\; ts \in TSF(a_i, a_k) \;\;\wedge\;\; (ts, false) \;\in\; \mathcal{F}(a_k, a_j)$$

$$\wedge \;\neg\; \exists a_l \in \mathrm{AN} : \;\; (a_k \prec_{(a_i, ts)} a_l)$$

$$\wedge \; (a_i, a_l) \in \mathrm{RL} \;\wedge\; ts \in TSF(a_i, a_k) \;\wedge\; (ts, bt) \;\in\; \mathcal{F}(a_l, a_j)\} \wedge bt \geq true \;\; \|$$

Note that if we were to set *bf* and *bt* equal to *false* and *true* respectively (as opposed to greater than or equal to *false* and *true* on the *information scale*), it amounts to allowing the unintuitive trust formation on the basis of less reliable sources that are certain in preference to more reliable sources that are ambivalent/skeptical [4].

[4] The whole problem with the world is that fools and fanatics are always so certain of themselves, but wiser people so full of doubts. — Betrand Russell

Evidence in support of Positive Functional Trust: a_i can functional trust a_j in trust scope ts if there is an explicit positive functional trust link from a_i to a_j, or there is majority of most referral trusted successors a_k of a_i that functional trust a_j rather than distrust a_j. In other words, for the purposes of a_j in trust scope ts, there are more endorsements than disapprovals via a_i's successors. We introduce a factor `Kp` to quantify the strength of majority for positive functional trust. Normally, its value is at least 1, and for simple majority, `Kp` is equal to 1.

$$\forall a_i, a_j \in AN : a_i \textbf{ can functional trust } a_j \textbf{ in trust scope } ts \text{ if}$$

$$(a_i, a_j) \in \mathrm{PFL} \quad \wedge \quad ts \in TSF(a_i, a_j) \quad \vee$$

$$\frac{\| \text{Undefeated functional trust of } a_i \textbf{ in } a_j \textbf{ for } ts \|}{\| \text{Undefeated nonfunctional trust of } a_i \textbf{ in } a_j \textbf{ for } ts \|} \quad > \quad Kp$$

Evidence in support of Negative Functional Trust: a_i can nonfunctional trust a_j in trust scope ts if there is an explicit negative functional trust link from a_i to a_j, or there is majority of most referral trusted successors a_k of a_i that functional distrust a_j rather than trust a_j. In other words, for the purposes of a_j in trust scope ts, there are more disapprovals than endorsements via a_i's successors. We introduce a factor `Kn` to quantify the strength of majority for negative functional trust. Normally, its value is at least 1, and for simple majority, `Kn` is equal to 1.

$$\forall a_i, a_j \in AN : a_i \textbf{ can nonfunctional trust } a_j \textbf{ in trust scope } ts \text{ if}$$

$$(a_i, a_j) \in \mathrm{NFL} \quad \wedge \quad ts \in TSF(a_i, a_j) \quad \vee$$

$$\frac{\| \text{Undefeated nonfunctional trust of } a_i \textbf{ in } a_j \textbf{ for } ts \|}{\| \text{Undefeated functional trust of } a_i \textbf{ in } a_j \textbf{ for } ts \|} \quad > \quad Kn$$

Evidence in support of Ambiguous Functional Trust:

$$\forall a_i, a_j \in AN : a_i \textbf{ can ambiguous trust } a_j \textbf{ in trust scope } ts \text{ if}$$

$$\textit{not } (a_i \text{ can functional trust } a_j \text{ in trust scope } ts) \quad \wedge$$

$$\textit{not } (a_i \text{ can nonfunctional trust } a_j \text{ in trust scope } ts)$$

To clarify, ambiguous case arises when both the positive trust threshold and the negative trust threshold are not crossed. For instance, if `Kp = 3` and `Kn = 1`, then 60% functional trust and 40% nonfunctional trust leads to ambiguous trust.

3.2 Properties of the Assigned Local Semantics

One can associate a unique meaning (in the form of referral trust and functional trust functions) with each trust DAG according to the semantics given in Section 3.1. The referral trust function $\mathcal{R}$ and functional trust function $\mathcal{F}$ can be computed in one-pass starting with agent nodes that have referral trust link out-degree of zero and processing agent nodes in reverse topological order. Furthermore, at each step of the computation, only the out-links order ($\prec_{(a,ts)}$) associated with referral links for an agent node a and the trust scope ts matters. The functional trust function $\mathcal{F}$ can be computed simultaneously with referral trust function $\mathcal{R}$. Note also that this approach employs majority to resolve potential conflicts only locally and propagates only final outcomes (not paths). For DAGs, the complexity of trust functions computation is linear in the size of the network (number of nodes and links).

In our trust model, the referral links are not weighted. So, in contrast with other approaches that associate real-valued weights with trust links, the impact of trust conclusions supported by longer trust paths do not decay. Implicitly, for each trust scope, we are relying on trustor's good judgement in gauging the *relative* trustworthiness of their neighbor's knowledge.

3.2.1 Dealing with Cycles The approach developed so far is not suitable for trust networks containing cycles that can cause apparent inconsistency or require iterative fixed point computation. However, one can assign a unique meaning to the network by defining a "personal" DAG for each agent that can be used to determine trust relationships for that agent as outlined in Thirunarayan and Verma [25]. The trust function over the agents with respect to the agent r can be computed using the topological order on agents (based on a breadth-first search [7] of the trust network). In fact, this computation can be carried out in parallel with respect to all agents. The computational complexity for trust networks with cycles is *quadratic*.

4 Conclusions and Future Work

In this paper, we developed a framework for describing semantics of trust networks containing referral trust links, (non)functional trust links, trust scopes, cycles, etc., by exploiting and adapting many evidence-based insights originally developed in the context of inheritance networks. It incorporates personalization by distinguishing between direct and inferred functional trust, and by using local trust relationship and majority to resolve conflicts. It localizes inconsistency and represents ambiguity explicitly. It formalizes trust values via partially ordered (implicit) discrete values over information and truth scales, and aggregation via least-upper-bound operation. The local partial ordering can vary with trust scope, and with respect to each user, to incorporate personalization. In the end, it provides a skeleton for enabling reasonable indirect trust inferences based on available direct, relative trust information. In practice, it may be necessary to attach summary justification with each trust conclusion. In future, we will explore extending current local trust framework by using a richer language of trust annotations, and generalizing local trust aggregation computation to take into account trust justification, trust path length (propagation horizon), and more expressive trust measures (subsuming real-valued trust), etc.

In the absence of standard data sets and benchmarks, we have not performed any substantial experiments beyond analyzing small-scale examples [9, 28, 29]. However, we have explained points of agreement and points of subtle disagreements with related works. Our trust model emphasizes clearer and more natural qualitative approach to trust in preference to quantitative approaches involving global, absolute weights. Eventually, trust networks analysis and formal specification can facilitate standardization using semantic web technologies, and used for information retrieval to make more informed decisions or take actions.

References

1. D. Artz and Y. Gil. A Survey of Trust in Computer Science and the Semantic Web. *Journal of Web Semantics*. Volume 5, Issue 2. pp. 58–71, 2007.
2. N. D. Belnap. How a computer should think. *Contemporary Aspects of Philosophy*, G. Ryle (ed.). Oriel Press, pp. 30–56, 1977.
3. V. G. Bintzios, T. G. Papaioannou, and G. D. Stamoulis. An effective approach for accurate estimation of trust of distant information sources in the semantic Web Second International Workshop on Security, Privacy and Trust in Pervasive and Ubiquitous Computing. 6 pages, 2006.
4. S. Brin and L. Page. The Anatomy of a Large-Scale Hypertextual Web Search Engine. WWW7 / Computer Networks 30(1-7), pp. 107–117, 1998.
5. J. J. Carroll, C. Bizer, P. Hayes, and P. Stickler. The Semantic Web Trust Layer. Developer's Day Talk at the 13th international conference on World Wide Web. 2004. `http://www4.wiwiss.fu-berlin.de/bizer/pub/CarrollBizer-Trust-WWW2004-DevDay.pdf`.
6. J. J. Carroll, C. Bizer, P. Hayes, and P. Stickler. Named graphs, provenance and trust. Proceedings of the 14th international conference on World Wide Web. pp. 613–622. 2005.
7. T. H. Cormen, C. E. Leiserson, R. L. Rivest, and C. Stein. Introduction to Algorithms (Second Edition) MIT Press and McGraw-Hill. 2005.
8. `http://www.trustlet.org/wiki/Epinions_dataset`, 2003.
9. K. Fullam, T. B. Klos, G. Muller, J. Sabater, A. Schlosser, Z. Topol, K. Suzanne Barber, J. S. Rosenschein, L. Vercouter, and M. Voss. A specification of the Agent Reputation and Trust (ART) testbed: experimentation and competition for trust in agent societies. In Proceedings of the 4th International Joint Conference on Autonomous Agents and MultiAgent Systems. pp. 512–518. 2005.
10. J. Golbeck. Computing and applying trust in web-based social networks. Doctoral Dissertation, University of Maryland, College Park, 2005. 2005.
11. J. Golbeck and J. Hendler. Inferring Trust Relationships in Web-based Social Network. *ACM Transactions on Internet Technology*. Vol. 6, Issue 4. pp. 497 - 529. 2006.
12. R. Guha, Ravi Kumar, Prabhakar Raghavan, Andrew Tomkins. Propagation of Trust and Distrust. *International World Wide Web Conference* (WWW2004), pp. 403–412, 2004.
13. J. Horty, R. Thomason, and D. Touretzky. Mixing strict and defeasible inheritance. Proceedings of the Seventh National Conference on Artificial Intelligence (AAAI-88), pp. 427–432, 1988.
14. J. Horty, and R. Thomason. A skeptical theory of inheritance in nonmonotonic semantic networks. Artificial Intelligence, Vol. 42, pp. 311–348, 1990.
15. J. Huang, and M. S. Fox. An ontology of trust: formal semantics and transitivity. Proceedings of the 8th international Conference on Electronic Commerce (ICEC'06). pp. 259–270. 2006.

16. A. Josang, R. Hayward, and S. Pope. Trust network analysis with subjective logic. ACSC '06: Proceedings of the 29th Australasian Computer Science Conference, pp. 85–94, 2006.
17. Y. Katz and J. Golbeck. Social Network-based Trust in Prioritized Default Logic. *Proceedings of The Twenty-First National Conference on Artificial Intelligence* (AAAI-06), 2006.
18. P. Massa and C. Hayes. Page-reRank: Using Trusted Links to Re-rank Authority. *Web Intelligence Conference*, pp. 614–617, 2005.
19. P. Massa and P. Avesani. Controversial Users demand Local Trust Metrics: an Experimental Study on Epinions.com Community. *Proceedings of The Twentieth National Conference on Artificial Intelligence* (AAAI-05), pp. 121–126, 2005.
20. P. Massa and P. Avesani. Trust-aware recommender systems. *Proceedings of the 2007 ACM Conference On Recommender Systems*, pp. 17–24, 2007.
21. M. Richardson, R. Agrawal and P. Domingos. Trust Management for the Semantic Web. *Proceedings of the Second International Semantic Web Conference*, pp. 351–368, 2003.
22. S. Schenk. On the Semantics of Trust and Caching in the Semantic Web. *Proceedings of the Seventh International Semantic Web Conference*. ISWC 2008, LNCS vol. 5318, Springer. pp. 533–549, 2008.
23. K. Thirunarayan and Michael Kifer. A Theory of Nonmonotonic Inheritance Based on Annotated Logic. *Artificial Intelligence Journal*, Vol. 60, pp. 23–50. March 1993.
24. K. Thirunarayan. Local Theories of Inheritance. *International Journal of Intelligent Systems*, Vol. 10(7), 617–645, July 1995.
25. K. Thirunarayan and R. Verma. A Framework for Trust and Distrust Networks. *Proceedings of Web 2.0 Trust Workshop (W2Trust)*, June 2008.
26. D. Touretzky, R. Thomason, and J. Horty. A skeptic's menagerie: conflictors, preemptors, reinstaters, and zombies in nonmonotonic inheritance. Proceedings of the Twelfth International Joint Conference on Artificial Intelligence (IJCAI-91), pp. 478 - 483, 1991.
27. X. Wang, J. You, and L. Yuan. A Default Interpretation of Defeasible Network. Proceedings of the Twelfth International Joint Conference on Artificial Intelligence (IJCAI-97), pp. 156 - 161, 1997.
28. Y. Wang and M.P. Singh. Trust Representation and Aggregation in a Distributed Agent System. Proceedings of the Twenty-First National Conference on Artificial Intelligence (AAAI-06). 2006.
29. Y. Wang and M.P. Singh. Formal Trust Model for Multiagent Systems. In Proceedings of the 20th International Joint Conference on Artificial Intelligence (IJCAI-07). pp. 1551 - 1556, 2007.
30. Y. Wang, Y. Hori, and K. Sakurai. Characterizing Economic and Social Properties of Trust and Reputation Systems in P2P Environment. Journal of Computer Science and Technology. pp. 129–140. 2008.
31. C. N. Ziegler and G. Lausen. Propagation Models for Trust and Distrust in Social Networks. Information Systems Frontiers, Vol. 7:4/5, pp. 337-358, 2005.