

Formalizing Trust-based Decision Making in Electronic Commerce Transactions

Tyrone Grandison¹, Han Reichgelt²

¹ IBM Almaden Research, 650 Harry Road, San Jose, CA 95120, USA
tyroneg@us.ibm.com

² Georgia Southern University, Statesboro, GA 30460, USA
han@georgiasouthern.edu

Abstract. Trust is the cornerstone of traditional business transactions and a lack of trust is one of the main inhibitors on the future growth of Electronic Commerce. However, trust is a nebulous social construct that is further complicated by its context dependence. Our assertion is that through study of the contextual artifacts and the process of trust formation in traditional business settings, we can devise a logic-based model that adequately captures the constructs needed to enable the automated creation of trusted Electronic Commerce transactions. In this paper, we introduce a model, called E2T2, which enables the creation of trust-based relationships between online merchants and their consumers.

Keywords: Logic, Trust, Trust Management.

1 Introduction

Trust is important to the transition to the next level of Electronic Commerce innovation. Without technology assurances that software is built on robust, verifiable trust technology, further consumer acceptance of online commercial innovations is likely to be slow.

The state of the art in commercial trust technology, e.g. KeyNote [1] and REFEREE [2], is very effective at mapping security credentials to access rights in static environments [3]. However, security constraints are only one facet of the trust problem. Other facets that are often ignored by security-focused trust solutions include reputation, risk propensity and past behavior [4]. Though systems have been built to address each of these and many other facets individually, there is no single underlying and unifying model that incorporates all the important aspects needed for Electronic Commerce.

The current set of logic-based trust models [5-7] suffers the same fate as their 'authorization-centric' trust management engineering counterparts. They focus on specific aspects of the trust problem, without taking all the dimensions of a business transaction into consideration.

The aim of our model, E2T2 (Epistemic Event Temporal Trust), is to show that it is possible to create a model that encapsulates the decision factors and that can be

implemented for any arbitrary E-Commerce application environment. While E2T2 is not complete, we believe that the model is flexible enough to allow adaptation to arbitrary application domains and facilitate the inclusion of emerging issues and concepts as they materialize.

We present the background notions needed for this discussion in section 2 and describe the basics of the E2T2 model in section 3. In section 4, we present how trust representation for Electronic Commerce is done in E2T2 and then round out the discussion on the auxiliary concepts needed for E-Commerce decision making (section 5). We conclude in section 6.

2 Background

Our work is grounded in several fields, namely: trust management [8,9], Electronic Commerce [4] and modal logic [10-15]. We introduce the base concepts in each of these areas.

2.1 Trust

Not all the definitions of trust in the computer science literature are applicable to Electronic Commerce [16]. For our context, we borrow from the definition used in [4] for trust in the E-Services environment, which states that “*trust is the quantified belief by a trustor with respect to the competence, honesty, security and dependability of a trustee within a specified context*”, where dependability encapsulates timeliness and reliability.

A trustor is defined as the subject of the relationship and a trustee the object. For example, if Luke wishes to initiate the purchase of a large shipment of fertilizer from RedCo Chemicals, then Luke is the trustor in this instance and the object is RedCo Chemicals. Luke’s decision to trust could be based on one or a combination of a myriad of different factors, e.g. RedCo’s reputation in the business community, a recommendation from one of Luke’s friends, RedCo’s SEC credentials or maybe even Luke’s personal past dealings with RedCo. These decision criteria are critical to the problem of establishing and evolving trust.

Currently, there are two prevailing views of trust management. The first is the one defined by Blaze et. al. [1] that advocates an authorization-based approach to trust management. The other is a business-oriented approach [16] that examines trust management through the core mechanism used in successful E-Commerce applications, e.g. reputation. We augment the latter approach in this paper.

2.2 Decision-Making in Electronic Commerce

When two entities, who may have no prior knowledge of each other, wish to engage in an Electronic Commerce transaction they must each ask themselves “*Should I trust this entity and engage in this transaction (or not)?*” The problem arises when each entity must choose the indicators that should be used to make this decision and request it from the other entity.

The common solution has been to use security credentials, such as authentication tokens and attributes, to establish either identity or determine the allowable trusted

tasks. However, security credentials are as good as the issuer of these credentials, do not give any indication of past or present behavior and only partially resolve the problem of how to interact with strangers. In a nutshell, a lot of contextual information is missing. This highlights the fact that the common solutions need to be enhanced.

In a standard electronic transaction e , where a wants to engage b , a must determine the level of trust or distrust and use this to take the most prudent course of action. Apart from assurances on the competence, honesty, security, timeliness and reliability of b , a has business considerations to include in his trust decision. From a purely commercial perspective, a favorable decision may depend on:

1. Intrinsic properties of the transaction.
2. Assurances about respecting a 's privacy.
3. a 's experience with b (with respect to similar transactions).
4. b 's past experience with others (who may be like a), with respect to similar transactions.
5. b 's reputation in the community of interest.
6. A recommendation on b 's behalf.

Generally, these factors can be classified as either transaction-specific or event-driven. Let's briefly present each.

2.2.1 Transaction-Specific Factors

Transaction-specific trust decision criteria are intrinsic properties of a transaction, and include the associated risk, the transaction value, the insurance coverage for the transaction and any associated economic incentive (or disincentive) associated with performing the transaction.

[4] purports that there are intuitive correlations between these factors and the final decision made. For example, transactions with a low risk value (in comparison to one's risk threshold) tend to lead to a positive trust decision. Under normal circumstances, a rational a considers a low-valued transaction trivial to his or her bottom-line and may not mind engaging with b , even if b is malicious and shatters all of a 's positive expectations. Generally, low-valued transactions are deemed as having very little risk. Normally, this is independent of an explicitly stated risk value for the transaction. Thus, a low-valued transaction has a high probability of being viewed as beneath one's risk threshold and often leads to positive trust decisions.

A transaction that is insured, at least to the transaction value, can be recovered (in the worst case). Having insurance lowers the risk threshold of a transaction. Normally, decisions involving insured transactions will not take the transaction's value into account. Instead the other factors involved in the decision will determine the outcome. If insurance is the only (or dominant) factor, then there is a higher chance of a favorable trust decision.

If b , or a third party, is willing to compensate a for engaging in e , to a level that a views as significant, then a will view the transaction as being below his risk threshold and thus be inclined to trust b and engage in the transaction. Similar reasoning can be applied to the disincentive case, where a third party pays a not to proceed with the transaction.

2.2.2 Event-Driven Factors

Event-driven factors relate to activity, agreements, arrangements or information that occur either prior to or outside the scope of the current trust decision with consequences that extend beyond the existing decision. For example, *a*'s body of knowledge concerning past encounters with *b* as it pertains to *e* constitutes information that occurred prior to the current decision and whose influence will extend beyond the decision. Thus, experience, reputation, recommendation and privacy concerns all fall into this category. We will discuss these later in this paper. Grandison [4] offers a discussion on the treatment of the correlations between these factors and the trust decision outcome. We will only state here that in the case of privacy constraints, *a*'s rights can be protected by 1) *a* not releasing any private information or releasing a limited view [17], or 2) *b* asserting that it won't violate *a*'s privacy and *a* actively monitoring the disclosures to and the activity performed by *b*, which seems to be infeasible in today's computing environment. If a privacy violation occurs, then this event will negatively impact the trust relationship. If no violation happens, then the relationship remains in tact.

2.3 Logic

In this section, we introduce the central concepts from epistemic logic and temporal logic.

2.3.1 Epistemic Logic

Epistemic logic [10,11] allows one to model complex scenarios involving knowledge and belief, by augmenting non-modal logic and leveraging the epistemic operators K_c and B_c such that $K_c A$ means Agent *c* knows *A* and $B_c A$ means Agent *c* believes *A* for some arbitrary proposition *A*. The fundamental concept behind any modal logic is that of a possible world, each of them essentially being a model (in the logical sense) of the underlying non-modal logic. In addition, each modal logic defines an accessibility relation between possible worlds, and different modal logics differ in the interpretation of this accessibility relation. Thus, in an epistemic logic, two worlds *w* and *w'* are accessible from each other for an agent *c* if both of them are compatible with *c*'s information state. The notation used is $R_c w w'$. World *w'* is said to be an *epistemic* or a *doxastic* alternative to world *w* for agent *c*, depending on whether knowledge or belief is the attitude under consideration.

2.3.2 Temporal Logic

Temporal logic [14,15] allows the representation and reasoning about propositions that are qualified with a time component. There are a number of different approaches in Computer Science to allow for temporal reasoning. They range from the situation calculus of McCarthy and Hayes, which essentially captures the temporal dimension by adding a temporal term to each function and predicate symbol, to modal temporal logic, to reified temporal logic. [19] provides an overview of the different approaches to incorporating a temporal dimension in a logical framework.

Because of its greater expressive power, we adopt a reified approach in E2T2. We distinguish between states and events, and we introduce the predicates **HOLDS** and **OCCURS** to indicate that a particular state is true over a period time, or that an event occurs at some point in time.

3 The Basics of the Model

From our definition of trust, we see that a trust decision is based on an entity's beliefs. Thus, our model must be founded on epistemic logic.

As there are two parties in a typical Electronic Commerce transaction, it is safe to assume that each party has their own set of beliefs and that these will evolve over time. Thus, the trust relationship is dynamic, in that the criteria or set of beliefs and one's behavior may change over time based on trigger events. This highlights the need to incorporate some form of temporal logic.

We will use a modal approach to modeling belief. However, since trust is a quantified belief, we cannot rely on the standard modal operator BEL to model belief. Instead, we introduce a modal operator schema BEL_{str} . By convention, $str \in \langle 0, 1 \rangle$, i.e., a number larger than 0 and less than or equal to 1, and indicates the level of confidence that the agent has in the belief.

We opt for a reified approach to temporal reasoning [18]. For our ontology, we adopt the formalism developed in [20], which introduces event and state tokens. We ignore subtleties regarding the choice of temporal primitives. For an in-depth discussion, see [21]. For the purposes of this paper, we model time as a set of points.

In the following sections, we introduce the basic vocabulary, individual terms and sentence construction in E2T2, then we present the semantic interpretation of the language constructs and highlight implications of our model's design assumptions.

3.1 The Language

We will start with the basics, i.e. presenting the syntactic primitives of E2T2 and showing proper sentence construction in this formalism.

3.1.1 The Basic Vocabulary

When representing trust, we need to be able to talk about different types of entity. We therefore define E2T2 as a sorted logic with the following sorts:

- I, Individuals ($i, j, k, i_1 \dots i_n$)
- A, Agents, a sub-sort of individuals, ($a, b, c, a_1 \dots a_n$)
- T, Times ($t_1 \dots t_n$)
- S, States ($s_1 \dots s_n$)
- E, Events ($e_1 \dots e_n$)
- N, the set of numbers between 0 and 1.

E2T2 contains a number of function symbols. Each n -place function symbol f has a signature $\Sigma_f: v_1, \dots, v_{n-1} \rightarrow v_n$ where v_1, \dots, v_n are sorts and v_1, \dots, v_{n-1} specify the sorts of the input arguments and v_n the sort of the output.

E2T2 contains a number of predicate symbols. Each n -place predicate p has a signature $\Pi_p: v_1, \dots, v_n$ where v_1, \dots, v_n are sorts. E2T2 contains the following special predicates:

- BEF, with signature $\Pi_{BEF}: T, T$. Intuitively, $BEF(t_1, t_2)$ is true iff t_1 is earlier than t_2 . Although a full temporal logic will have to contain more predicates between times, this one temporal predicate is sufficient for our current purposes.
- HOLDS, with signature $\Pi_{HOLDS}: S, T, T$. Intuitively, $HOLDS(s, t_1, t_2)$ is true iff state s is true between time t_1 and t_2 .

- OCCURS, with signature $\Pi_{\text{OCCURS}}: E, T$. Intuitively, OCCURS(e, t) is true iff event e occurs at time t .

Finally, E2T2 contains the 4-place modal operator BEL whose first three arguments are of type A, N and T respectively and whose final argument is a proposition. Intuitively, BEL(a, str, t, p) indicates that a believes p with confidence level str at time t .

3.1.2 Individual Terms

Individual terms in E2T2 are defined in the normal way. Thus:

- Every constant is an individual term;
- If f is a function symbol with $\Sigma_f: v_1, \dots, v_{n-1} \rightarrow v_n$ and Z_1, \dots, Z_{n-1} are terms of sorts v_1, \dots, v_{n-1} respectively, then $f(Z_1, \dots, Z_{n-1})$ is an individual term of sort v_n .

3.1.3 Sentences

In defining the syntax of E2T2, we make use of the following auxiliary notion:

- If ϕ is a sentence and z and x terms of sort v then $\phi[z/x]$ is obtained by replacing all occurrences of z in ϕ by x .

Sentences in E2T2 are formed in the normal way, that is:

- If p is an n -place predicate with signature $\Pi_p: v_1, \dots, v_n$ and Z_1, \dots, Z_n are terms of sorts v_1, \dots, v_n then $p(Z_1, \dots, Z_n)$ is a sentence;
- If ϕ and ψ are sentences, then so are $\neg\phi$, $(\phi \rightarrow \psi)$, $(\phi \& \psi)$ and $(\phi \vee \psi)$;
- If ϕ is a sentence that contains terms z and x of sort v , then $(\forall x:v)[\phi(z/x)]$ and $(\exists x:v)[\phi(z/x)]$ are sentences.
- If ϕ is a sentence and a, str and t are terms of type A, N and T respectively, then BEL(a, str, t, ϕ) is a sentence.

3.2 Semantics

The semantics for E2T2 is relatively complicated as it has to combine elements from the semantics of temporal logic with elements of the semantics of epistemic logic. Moreover, because we deal with quantified belief, we cannot rely on the standard possible worlds approach to provide semantics for the belief modality. We achieve the desired effect by complicating the accessibility relation between possible worlds.

From an intuitive point of view, an E2T2 model consists of a set of individuals, suitably partitioned to take account of the different sorts in the language, a set of time intervals and a set of possible worlds. The model contains an ordering relationship $<$ between time intervals. In addition, for each agent, we define an accessibility function that maps any two pairs of possible worlds at a particular point in time onto a number larger than 0 and less than or equal to 1. This number indicates the likelihood that a world w' is accessible from world w according to the agent at time t . It is loosely based on Lewis' complication of the standard possible world semantics to deal with counterfactuals [22] and will be used to deal with the confidence level that an agent has in different beliefs.

With this in mind, we can now formally define an E2T2 model as follows:

An E2T2 model \mathcal{M} consists of

- A tuple $\langle \mathbf{O}, \mathbf{W}, \mathbf{T} \rangle$ where
 - \mathbf{O} , the set of objects is partitioned into 3 non-overlapping sets, \mathbf{I} , \mathbf{S} and \mathbf{E} , the sets of individuals, states and events respectively, and \mathbf{I} contains a non-empty set \mathbf{A} , the set of agents;

- \mathbf{W} is the set of possible worlds
 - \mathbf{T} is the set of time intervals;
 - \mathbf{N} is the set of numbers in the interval $< 0, 1]$
- b. A temporal ordering relation $<$ over \mathbf{T} .
- c. A function \mathbf{Acc} that associates with each $\mathbf{a} \in \mathbf{A}$ a 3-place function from $\mathbf{W} \times \mathbf{W} \times \mathbf{T}$ into \mathbf{N} .
- d. An interpretation function $\mathfrak{I}^{\mathcal{M}}$.

The interpretation function $\mathfrak{I}^{\mathcal{M}}$ assigns an appropriate denotation to each expression in E2T2 in each world-time pair. We use $\mathfrak{I}^{\mathcal{M}}_{\mathbf{w},\mathbf{t}}(\text{exp})$ as a short-cut notation for $\mathfrak{I}^{\mathcal{M}}(\text{exp}, \mathbf{w}, \mathbf{t})$.

For primitive individual terms, it is defined in the standard ways:

- a. If z is a primitive term of sort v , then $\mathfrak{I}^{\mathcal{M}}_{\mathbf{w},\mathbf{t}}(z) \in \text{Corr}(v)$, where Corr is a function mapping a sort into the corresponding set of objects. Thus,
- $$\text{Corr}(\mathbf{A}) = \mathbf{A}, \text{Corr}(\mathbf{I}) = \mathbf{I}, \text{Corr}(\mathbf{S}) = \mathbf{S}, \text{Corr}(\mathbf{E}) = \mathbf{E}, \text{Corr}(\mathbf{T}) = \mathbf{T},$$
- $$\text{Corr}(\mathbf{N}) = \mathbf{N}.$$

In order to avoid the problem of trans-world identity, we stipulate that for all primitive terms $\mathfrak{I}^{\mathcal{M}}_{\mathbf{w},\mathbf{t}}(z) = \mathfrak{I}^{\mathcal{M}}_{\mathbf{w}',\mathbf{t}'}(z)$. In other words, the denotation of a primitive term does not vary from world-time pair to world-time pair.

- b. If f is a function symbol with $\mathbf{S}_i: v_1, \dots, v_{n-1} \rightarrow v_n$ then $\mathfrak{I}^{\mathcal{M}}_{\mathbf{w},\mathbf{t}}(f)$ is a function from $\text{Corr}(v_1) \times \dots \times \text{Corr}(v_{n-1})$ into $\text{Corr}(v_n)$, and $\mathfrak{I}^{\mathcal{M}}_{\mathbf{w},\mathbf{t}}(f(v_1, \dots, v_{n-1})) = \mathfrak{I}^{\mathcal{M}}_{\mathbf{w},\mathbf{t}}(f)(\mathfrak{I}^{\mathcal{M}}_{\mathbf{w},\mathbf{t}}(v_1), \dots, \mathfrak{I}^{\mathcal{M}}_{\mathbf{w},\mathbf{t}}(v_{n-1}))$

For predicates, $\mathfrak{I}^{\mathcal{M}}$ is defined in the expected way as well:

- a. If p is an n -place predicate with signature $\Pi_p v_1, \dots, v_n$, then $\mathfrak{I}^{\mathcal{M}}_{\mathbf{w},\mathbf{t}}(p) \subseteq (\text{Corr}(v_1) \times \dots \times \text{Corr}(v_n))$

The only complication is that the interpretation of the special predicate BEF is the ordering relation between points in time.

In order to define $\mathfrak{I}^{\mathcal{M}}$ for sentences, we introduce an auxiliary notion, namely the set of all possible worlds that are accessible from a given world \mathbf{w} according to actor \mathbf{a} at time \mathbf{t} with at least confidence level \mathbf{str} ; which is $\mathbf{Acc}(\mathbf{w}, \mathbf{a}, \mathbf{t}, \mathbf{str})$. Formally, this notion can be defined as follows:

$$\mathbf{w}' \in \mathbf{Acc}(\mathbf{w}, \mathbf{a}, \mathbf{t}, \mathbf{str}) \text{ iff } \mathbf{Acc}(\mathbf{a})(\mathbf{w}, \mathbf{w}', \mathbf{t}) \leq \mathbf{str}$$

Although this definition may at first glance seem counterintuitive, the reason for it will become clear in the following sections.

With this concept, we can now extend $\mathfrak{I}^{\mathcal{M}}$ to sentences. $\mathfrak{I}^{\mathcal{M}}$ maps each sentence to either 1 (true) or 0 (false) according to the following rules:

- a. $\mathfrak{I}^{\mathcal{M}}_{\mathbf{w},\mathbf{t}}(p(v_1, \dots, v_n)) = 1$ iff $\langle \mathfrak{I}^{\mathcal{M}}_{\mathbf{w},\mathbf{t}}(v_1), \dots, \mathfrak{I}^{\mathcal{M}}_{\mathbf{w},\mathbf{t}}(v_n) \rangle \in \mathfrak{I}^{\mathcal{M}}_{\mathbf{w},\mathbf{t}}(p)$ and 0 otherwise;
- b. $\mathfrak{I}^{\mathcal{M}}_{\mathbf{w},\mathbf{t}}(\neg\phi) = 1$ iff $(\mathfrak{I}^{\mathcal{M}}_{\mathbf{w},\mathbf{t}}(\phi) = 0)$, and 0 otherwise;
- c. $\mathfrak{I}^{\mathcal{M}}_{\mathbf{w},\mathbf{t}}(\phi \rightarrow \psi) = 1$ iff $((\mathfrak{I}^{\mathcal{M}}_{\mathbf{w},\mathbf{t}}(\phi) = 0) \text{ or } (\mathfrak{I}^{\mathcal{M}}_{\mathbf{w},\mathbf{t}}(\psi) = 1))$, and 0 otherwise;
- d. $\mathfrak{I}^{\mathcal{M}}_{\mathbf{w},\mathbf{t}}(\phi \& \psi) = 1$ iff $(\mathfrak{I}^{\mathcal{M}}_{\mathbf{w},\mathbf{t}}(\phi) = \mathfrak{I}^{\mathcal{M}}_{\mathbf{w},\mathbf{t}}(\psi) = 1)$, and 0 otherwise;
- e. $\mathfrak{I}^{\mathcal{M}}_{\mathbf{w},\mathbf{t}}(\phi \vee \psi) = 1$ iff $((\mathfrak{I}^{\mathcal{M}}_{\mathbf{w},\mathbf{t}}(\phi) = 1) \text{ or } (\mathfrak{I}^{\mathcal{M}}_{\mathbf{w},\mathbf{t}}(\psi) = 1))$, and 0 otherwise;
- f. $\mathfrak{I}^{\mathcal{M}}_{\mathbf{w},\mathbf{t}}((\forall x:v)[\phi]) = 1$ iff for all $\mathfrak{I}^{\mathcal{M}}$ which are exactly like $\mathfrak{I}^{\mathcal{M}}$ except for the value assigned to z , $\mathfrak{I}^{\mathcal{M}}_{\mathbf{w},\mathbf{t}}(\phi[x/z]) = 1$ where z is a term of sort v not occurring in ϕ , and 0 otherwise.
- g. $\mathfrak{I}^{\mathcal{M}}_{\mathbf{w},\mathbf{t}}((\exists x:v)[\phi]) = 1$ iff there is an $\mathfrak{I}^{\mathcal{M}}$ which is exactly like $\mathfrak{I}^{\mathcal{M}}$ except for the value assigned to z such that $\mathfrak{I}^{\mathcal{M}}_{\mathbf{w},\mathbf{t}}(\phi[x/z]) = 1$ where z is a term of sort v not occurring in ϕ .

- h. $\mathfrak{I}_{w,t}^{\mathcal{M}}(\text{BEL}(a, \text{str}, t, \varphi) = 1)$ iff (for all $w' \in \text{Acc}(w, a, t', \text{str})$, $\mathfrak{I}_{w',t}^{\mathcal{M}}(\varphi) = 1$) where $t' = \mathfrak{I}_{w,t}^{\mathcal{M}}(t)$ and 0 otherwise.

The decisions taken in designing the syntax and semantics have led to interesting consequences, which are presented in Appendix A.

4 Representing Trust

In this section we use the logic framework of E2T2 to formalize the notion of trust. However, before we do so, we first introduce some special function symbol and predicates.

4.1 Auxiliary Function and Predicate Symbols

The first function symbol that we need is **CANINDUCE** with signature $\Sigma_{\text{CANINDUCE}}: A, E \rightarrow S$. Intuitively, **CANINDUCE**(a,e) describes the state of a being able to bring about event e, either directly or through some third party. Often, the event will be some action that is performed by a.

A second required function symbol is **REQUEST** with signature $\Sigma_{\text{REQUEST}}: A, A, E \rightarrow E$. Intuitively, **REQUEST** maps two agents and an event into an event, namely the event of the first agent requesting the second agent to bring about the event.

A third function symbol is **COMMIT** with signature $\Sigma_{\text{COMMIT}}: A, E, T \rightarrow E$. Intuitively, **COMMIT** maps an agent, an event and a time into an event, namely the event of the agent committing itself to bring about event before that time.

The final function symbol that we need is slightly more complicated, namely **TPER**, with the signature $\Sigma_{\text{TPER}}: A, A, E, T \rightarrow T$. **TPER** takes two agents a and b, an event e, and a time t and maps this into a second time period t'. We assume that agent b has committed to agent a to perform e at time t. **TPER** returns the time by which agent a can reasonably expect agent b to have delivered on its commitment. It reflects the intuition that, when an agent commits to do something, one can reasonably expect him or her to perform the action within a certain time period. The time period that is reasonable clearly depends on the action committed to. For an event like passing the salt at a dinner table, the period will be fairly short; for a promise made to an incoming student in a four year program to help her complete her degree successfully, the time period is longer. We stipulate that if $x = \text{TPER}(a_1, a_2, e_1, t_1)$ then $\text{BEF}(t_1, x)$. This function is used to model timeliness, which is a crucial element of the trust formalism.

4.2 Trust

We are now in a position to give a preliminary definition of trust in terms of E2T2. We previously stated that for purposes, trust is the quantified belief by a trustor with respect to the competence, dependability, security and honesty of a trustee with respect to the trustee bringing about some event or state of affairs and in a specified context.

With this in mind, we introduce a predicate **TRUST** with the signature $\Pi_{\text{TRUST}}: A, N, A, T, E$. Intuitively, **TRUST**(a,str,b,t,e) means that a trusts b at trust level str at time t to bring about event e. Naively, we see that this trust relationship will only be

established when **a** believes that **b** is currently competent to bring about **e** (in a timely manner), when **a** believes that **b** will always carry out commitments with respect to **e** when requested (in a timely manner), when **a** believes that **b** is currently secure and when **a** believes that **b** is currently honest. It becomes clear that timeliness is an intrinsic requirement for all the other attributes, rather than a single explicit constraint.

Further examination of the clauses reveals that there are multiple dimensions of time present. For example, with the competence and security clauses, we assume that the trustee is competent and secure for a reasonable period, while for the dependability clause we assume that whenever the trustor makes a request of the trustee to bring about some event, the trustee forms, within a reasonable time of the request being made, the intention of bringing about **e** in a timely manner. All of this raises the question *what time period is considered reasonable?* The answer lies in a number of factors. One factor is the agent who considers the period to be reasonable, or not. Some agents are more patient than others. The second factor is the person who is to bring about the event. Although one might trust both a seven year old child and an adult to set the table, when asked to do so, one would expect the child to take longer than the adult. The final factor is the event itself. One would expect it to take longer to make a gourmet meal than it takes to make a cup of coffee.

Let's discuss the formulations of the exact clauses for each attribute.

Competence: A competent entity is one that is able to perform or lead to the performance of an event. In our framework, this translates to **a**'s belief that **b** is capable of bringing about **e** in a timely manner. We formally define this as:

$$\text{BEL}(\mathbf{a}, \text{str}_1, t, [(\forall x:T) (\forall x':T) \text{BEF}(x, x') \& \text{BEF}(x', \text{TPER}(\mathbf{a}, \mathbf{b}, \mathbf{e}, x)) \rightarrow \text{HOLDS}(\text{CANINDUCE}(\mathbf{b}, \mathbf{e}), x', \text{TPER}(\mathbf{a}, \mathbf{b}, \mathbf{e}, x))])$$

a does not have to believe that **b** is capable of bringing about **e** in perpetuity. **a** merely has to believe that **b** can bring about event **e** within the period in which it is reasonable for **a** to expect **b** to bring about **e**. Note that we also do not mean to imply that **a** needs to believe in **b**'s competence for the entire period in which it is reasonable for **a** to expect **b** to bring about **e**. **a** can change its opinion at any time. The clause merely states that at the time at which **a** trust **b** to perform **e**, **a** believes that **b** is competent to bring about **e** in a reasonable timeframe.

Dependability: Dependability refers to reliability and timeliness. A reliable entity is one that performs an event when it is requested to do so. A dependable entity is a reliable one that executes in a timely manner. We formalize this as:

$$\text{BEL}(\mathbf{a}, \text{str}_2, t, (\forall x:T) [\text{OCCURS}(\text{REQUEST}(\mathbf{a}, \mathbf{b}, \mathbf{e}), x) \rightarrow (\exists x':T) (\exists x'':T) [\text{BEF}(x, x') \& \text{BEF}(x', \text{TPER}(\mathbf{a}, \mathbf{b}, \text{REQUEST}(\mathbf{a}, \mathbf{b}, \mathbf{e}), x)) \& \text{BEF}(x', x'') \& \text{OCCURS}(\text{COMMIT}(\mathbf{b}, \mathbf{e}, \text{TPER}(\mathbf{a}, \mathbf{b}, \mathbf{e}, x)), x'')]])$$

The clause states that **a** believes that whenever it requests of **b** to bring about **e**, then **b** will, within a reasonable period of the request being made, form the intention to bring about event **e** in a timely manner. We use the function symbol **TPER** twice, to reflect two subtly different forms of timeliness. In order to be considered reliable, a trustee has to meet two requirements.

First, whenever the trustor makes a request, the trustee must, in a timely manner, respond to the request by making some commitment. How long a period is reasonable depends on the request. One would expect a friend to respond to a request to make a cup of coffee more quickly than he or she might respond to a request to

marry you. This element of timeliness is modeled in the first use of TPER. However, it is not enough for the trustee to merely make the commitment in a timely manner. The trustee must also commit to bringing about whatever the trustor requests within reasonable timescales. We model this aspect through the second use of TPER.

An example may further clarify the distinction. Assume that **a** trusts waitress **b** in Frank's Diner to serve a meal. **a** believes that whenever he orders a meal from **b**, i.e. requests **b** to serve him a meal, **b** will form the commitment to bring the meal after **a** makes the request. However, **a** would not expect **b** to bring it the meal immediately, although there clearly is some limit to the amount of time that **a** would be prepared for the meal to arrive.

Security: A secure entity is one that provides assurances on its safe operation and execution.

$$\text{BEL}(a, \text{str}_3, t, (\forall x:T) (\forall x':T) [\text{BEF}(x, x') \ \& \ \text{BEF}(x', \text{TPER}(a,b,e,x)) \rightarrow \text{HOLDS}(\text{SECURE}(b,e), x', \text{TPER}(a,b,e,x))])$$

As was the case with our analysis of competence, **a** does not have to assume that **b** will be secure with respect to **e** for ever, but merely for the period in which it is reasonable for **a** to expect **b** to perform **e** after **a** has requested **b** to perform **e**. Again, as in the case of the competence attribute, it is also true that **a** can change its view about **b**'s security at any time.

Honesty: A honest entity is one that is sincere with regards to its interactions, i.e. it does what it commits to doing.

$$\text{BEL}(a, \text{str}_4, t, (\forall x:T) (\forall x':T) [\text{OCCURS}(\text{COMMIT}(b,e,x'), x) \rightarrow (\exists x'':T) [\text{BEF}(x, x') \ \& \ \text{BEF}(x', x'') \ \& \ \text{OCCURS}(e, x'')]])$$

The above states that **a** believes that whenever **b** has committed to bringing about **e** before some time, **e** will indeed occur before that time. Thus, **b** is sincere with respect to its commitments.

The Complete Formulation: Putting all the pieces together, we get the following formulation of trust:

$$\begin{aligned} & \text{TRUST}(a, \text{CALC}(\text{str}_1, \text{str}_2, \text{str}_3, \text{str}_4), b, t, e) \leftrightarrow \\ & \text{BEL}(a, \text{str}_1, t, [(\forall x:T) (\forall x':T) \text{BEF}(x, x') \ \& \ \text{BEF}(x', \text{TPER}(a,b,e,x)) \rightarrow \\ & \quad \text{HOLDS}(\text{CANINDUCE}(b,e), x', \text{TPER}(a,b,e,x))] \\ & \ \& \ \text{BEL}(a, \text{str}_2, t, (\forall x:T) [\text{OCCURS}(\text{REQUEST}(a,b,e), x) \rightarrow \\ & \quad (\exists x':T) (\exists x'':T) [\text{BEF}(x, x') \ \& \ \text{BEF}(x', \text{TPER}(a,b,\text{REQUEST}(a,b,e,x))) \ \& \\ & \quad \text{BEF}(x', x'') \ \& \ \text{OCCURS}(\text{COMMIT}(b,e, \text{TPER}(a,b,e,x)), x'')]]) \\ & \ \& \ \text{BEL}(a, \text{str}_3, t, (\forall x:T) (\forall x':T) [\text{BEF}(x, x') \ \& \ \text{BEF}(x', \text{TPER}(a,b,e,x)) \rightarrow \\ & \quad \text{HOLDS}(\text{SECURE}(b,e), x', \text{TPER}(a,b,e,x))]) \\ & \ \& \ \text{BEL}(a, \text{str}_4, t, (\forall x:T) (\forall x':T) [\text{OCCURS}(\text{COMMIT}(b,e,x'), x) \\ & \quad \rightarrow (\exists x'':T) [\text{BEF}(x, x') \ \& \ \text{BEF}(x', x'') \ \& \ \text{OCCURS}(e, x'')]]) \end{aligned} \quad (1)$$

There are a number of issues that we wish to draw attention to.

First, the term $\text{CALC}(\text{str}_1, \text{str}_2, \text{str}_3, \text{str}_4)$ indicates that the trust level is dependent on the four quantified beliefs that make up the definition of trust. Clearly, this raises the question about the actual definition of **CALC**, which may be domain-specific and or entity-related. The exact mathematical formulae used to calculate the trust value may be as simple as arithmetic average or as complex as weighted statistical computation. This paper is not focused on trust calculation and we will leave this issue until another occasion. However, there are a number of stipulations that we need to put on the function **CALC**. First, we assume that the range of the function **CALC** is the interval [0,1]. Second, whatever function is chosen to implement **CALC**, it

must be monotonically increasing. We want to make sure that, as the confidence level in the trustor's beliefs of either the trustee's competence, dependability, security, or honesty with respect to e increases, then its trust level in the trustee to bring about e minimally does not decrease, and should probably increase.

Second, Note that the model theory and the increasing monotonicity of CALC imply that if a trusts b at time t to bring about e at trust level n , then a also trusts b at time t to bring about e at trust level n' where n' is less than n . After all, if a holds any belief at confidence level c , a also holds that belief at any confidence level less than c . If a trusts b to bring about e at level n , where n is the result of applying the function CALC on the confidence levels of the beliefs in the competence, dependability, security and honesty of b , then a also believes in the competence, dependability, security and honesty of b at lower confidence levels. Applying CALC to these lower confidence levels must, because of the increasing monotonicity of CALC, result in a value that is less than the original value. Thus,

$$\text{TRUST}(a, n, b, t, e) \rightarrow (\forall n') [(0 < n' \leq n) \rightarrow \text{TRUST}(a, n', b, t, e)] \quad (2)$$

Since n takes on a value greater than 0 and less than or equal to 1, it also follows that

$$\neg \text{TRUST}(a, n, b, t, e) \rightarrow (\forall n') [(n \leq n' \leq 1) \rightarrow \neg \text{TRUST}(a, n', b, t, e)] \quad (3)$$

Our formulation of trust lends itself to the modeling of constructs needed for Electronic Commerce, which for purposes of brevity are presented in Appendix B.

6 Conclusion

Epistemic Event Temporal Trust (E2T2) is a logical framework that leverages knowledge on the trusting behavior of businesses to create a unifying model that allows all the components of the trust decision to be factored into the process of trust establishment. The model's design allows it to be flexible enough to model the trusting behavior of any arbitrary business environment.

Our future works involves modeling the more troublesome notions that online entities must handle (e.g. partial information, uncertainty, etc.). We will also create E2T2 reasoning engines and evaluate their effectiveness. Finally, we will consider the best deployment mechanisms.

The overall focus on the work on E2T2 is to provide a foundation for discourse in the research community and to underscore the importance of the need to take a holistic approach to the trust problem.

References

- [1] Blaze, M., J. Feigenbaum and Keromytis, A.D.: *KeyNote: Trust Management for Public-Key Infrastructures*. in Security Protocols International Workshop. 1998. Cambridge, England. <http://www.cis.upenn.edu/~angelos/Papers/keynote-position.ps.gz>
- [2] Chu, Y.-H., J. Feigenbaum, B. LaMacchia, P. Resnick and M. Strauss: *REFeree: Trust Management for Web Applications*. 1997, AT&T Research Labs Research Report. <http://www.farcaster.com/papers/www6-referee/>

- [3] Blaze, Matt, Ioannidis, John, Keromytis, Angelos D.: *Experience with the Keynote Trust Management System: Applications and Future Directions*, Proceedings of Intl. Trust Management Conference (iTrust) 2003, Pg 284-300.
- [4] Grandison, T. : *Conceptions of Trust: Definition, Constructs and Models* in Trust in E-Services: Technologies, Practices and Challenge, Editor: Ronggong Song. Published by IDEA Group. 2007.
- [5] Xie, H.-b., Zhou, M.-t.: *PKI Trust Model Analysis Based on Probabilistic Model and Conditional Predicate Calculus Logic*, MINIMICRO SYSTEMS -SHENYANG, Vol 27, No 1, 2006.
- [6] Nefti, Samia, Meziane, Farid, Kasiran, Khairudin,: *A Fuzzy Trust Model for E-Commerce*, Proceedings of the Seventh IEEE International Conference on E-Commerce Technology (CEC'05), Pages 401-404, 2005.
- [7] Ramchurn, S. D., Sierra, C., Godo, L. and Jennings, N. R. : *Devising a trust model for multi-agent interactions using confidence and reputation*. International Journal of Applied Artificial Intelligence 18(9-10) pp. 833-852. 2004.
- [8] Ji, Ma, Orgun, M.A. : *Trust management and trust theory revision*, IEEE Transactions on Systems, Man and Cybernetics, Vol 36, Issue 3, May 2006, Pages 451- 460.
- [9] Ruohomaa, S., Kutvonen, L.: *Trust management survey*, Proceedings of the 3rd International Conference on Trust Management 2005, Pages 77–92.
- [10] Georg Henrik von Wright: *An Essay in Modal Logic*, 1951
- [11] Jaakko Hintikka: *The Logic of Epistemology and the Epistemology of Logic*.
- [12] R. Kowalski and M.Sergot: *A Logic-Based Calculus of Events* New Generation Computing, vol. 4 pp. 67–95, 1986.
- [13] R. Miller and M. Shanahan: *The event-calculus in classical logic — alternative axiomatizations*. Electronic Transactions on Artificial Intelligence, 3(1):77-105, 1999.
- [14] Venema, Yde: *Temporal Logic*, in Goble, Lou, ed., *The Blackwell Guide to Philosophical Logic*. Blackwell, 2001.
- [15] E.A. Emerson: *Temporal and modal logic*, Handbook of Theoretical Computer Science, Chapter 16, the MIT Press, 1990
- [16] Grandison, T. and Sloman, M.: *A Survey of Trust in Internet Applications*, IEEE Communications Surveys and Tutorials, Vol. 3 No. 4, Oct-Dec 2000.
- [17] K. Lefevre, R. Agrawal, V. Ercegovic, R. Ramakrishnan, Y. Xu, D. DeWitt: *Limiting Disclosure in Hippocratic Databases*. Proc. of the 30th Int'l Conf. on Very Large Databases (VLDB 2004), Toronto, Canada, August 2004.
- [18] Vila, L. and H. Reichgelt: *The token reification approach to temporal reasoning*. Artificial Intelligence, 83 (1996), 59-74.
- [19] Reichgelt, H. and L. Vila: *Temporal qualification in artificial intelligence*. In: M. Fisher, D. Gabbay and L. Vila (eds) *Handbook of Temporal Reasoning in Artificial Intelligence*. Amsterdam: Elsevier, 2005.
- [20] Shoham, Y.: *Temporal logics in AI: Semantical and ontological considerations*. Artificial Intelligence, 33 (1987), 89-104.
- [21] Vila, L.: *Formal Theories of Time and Temporal Incidence*. In: M. Fisher, D. Gabbay and L. Vila (eds) *Handbook of Temporal Reasoning in Artificial Intelligence*. Amsterdam: Elsevier, 2005.
- [22] Lewis, D.: *Counterfactuals*. Oxford: Blackwell, 1973.

Appendix A: E2T2 Design Consequences

Our definition of belief has a number of consequences that are worth mentioning. First, notice that, as any standard epistemic logic, our logic suffers from the problem of logical omnidoxasticity:

$$(\forall a:A)(\forall n:N)(\forall t:T) [((\phi \rightarrow \psi) \& \text{BEL}(a, n, t, \phi)) \rightarrow \text{BEL}(a, n, t, \psi)] \quad (4)$$

Second, agents can change their minds, i.e. if an agent believes at time t that some event will occur at time t' or that some state will hold at time t' , then it does not follow that it must believe at some other time t'' that this event occurs or state holds at t' . In other words,

$$\neg (\forall a:A)(\forall n:N)(\forall t, t', t'':T)(\forall e:E) [\text{BEL}(a, n, t, \text{OCCURS}(e, t')) \rightarrow \text{BEL}(a, n, t'', \text{OCCURS}(e, t'))] \quad (5)$$

$$\neg (\forall a:A)(\forall n:N)(\forall t, t', t'':T)(\forall s:S) [\text{BEL}(a, n, t, \text{HOLDS}(s, t')) \rightarrow \text{BEL}(a, n, t'', \text{HOLDS}(e, t'))] \quad (6)$$

After all, there is no requirement that all the worlds that are accessible from some given world at time t , are also accessible from that world at another time.

Third, if an agent believes a proposition with a certain confidence level, then it will also believe that proposition with a lower confidence level. In other words:

$$(\forall a:A)(\forall n, n':N)(\forall t:T) [(\text{BEL}(a, n, t, \phi) \& (n' < n)) \rightarrow \text{BEL}(a, n', t, \phi)] \quad (7)$$

This proposition follows directly from the way in which we defined $\mathbf{Acc}(w, a, t, \text{str})$ as the set of possible worlds that were accessible from w according to a at time t with a confidence level of at least str . Clearly, it follows that if a proposition ϕ is true in all worlds that are accessible from some world w according to a at time t with at least confidence level n , then it must also be true in all worlds that are accessible with a lower confidence level.

Fourth, while it does follow that if an agent believes the negation of a proposition with a certain confidence level, then it does not believe that proposition with any confidence level, it does not follow that if an agent does not believe a proposition with a certain confidence level, then it believes the negation of that proposition. In other words,

$$\neg (\forall a:A)(\forall t:T)(\forall n:N) [\neg \text{BEL}(a, n, t, \phi) \rightarrow (\exists n')[\text{BEL}(a, n', t, \neg \phi)]] \quad (8)$$

$$(\forall a:A)(\forall t:T)(\forall n:N) [\text{BEL}(a, n, t, \neg \phi) \rightarrow \neg (\exists n')[\text{BEL}(a, n', t, \phi)]] \quad (9)$$

Again, the above follows directly from the way in which the model theory has been defined. An agent not believing a proposition with a certain confidence level merely means that there is at least one possible world that is accessible with that confidence level in which the negation of that proposition is true; it certainly does not follow that the negation of the proposition must be true in all possible worlds that are accessible with an arbitrary confidence level. On the other hand, if an agent believes the negation of a proposition with an arbitrary confidence level, then the negation of that proposition must be true in all worlds that are accessible with a certain confidence level, and that, no matter what confidence level one considers, there is always at least one possible world in which the proposition is false.

Fifth, because denotations of primitives are unique across possible worlds and function symbols and predicates, other than BEF , varies from possible world to possible world, it is possible to tailor the model to any particular world of interest (e.g. an auction website, a manufacturing system, etc).

Appendix B: Electronic Commerce Constructs in E2T2

From [4], we know that mathematical properties such as transitivity, asymmetry, etc. cannot be axioms of any trust model simply because they are not universally applicable across domains. Given the way in which we have formulated the model

theory, none of these properties are axioms. However, each property could be modeled in E2T2 if it was applicable to the domain of interest.

In this section, we highlight the notions of distrust, trust dynamism and risk.

5.1 Distrust

Distrust can be conceptually viewed as simply a lack of trust. Formally, it is the quantified belief by a trustor that a trustee is incompetent, undependable, not secure and or dishonest with respect to bringing about some event or state of affairs for a specified context. For example, if **a** does not believe in **b**'s ability to competently execute event **e**, then **a** may distrust **b** to execute **e**. The same holds true for all the other attributes.

$$\begin{aligned}
& \neg \text{BEL}(a, \text{str}_1, t, [(\forall x:T) (\forall x':T) \text{BEF}(x, x') \& \text{BEF}(x', \text{TPER}(a,b,e,x)) \rightarrow \\
& \quad \text{HOLDS}(\text{CANINDUCE}(b,e), x', \text{TPER}(a,b,e,x))] \\
& \vee \neg \text{BEL}(a, \text{str}_2, t, (\forall x:T) [\text{OCCURS}(\text{REQUEST}(a,b,e), x) \rightarrow \\
& \quad (\exists x':T) (\exists x'':T) [\text{BEF}(x, x') \& \text{BEF}(x', \text{TPER}(a,b,\text{REQUEST}(a,b,e),x)) \& \\
& \quad \text{BEF}(x', x'') \& \text{OCCURS}(\text{COMMIT}(b,e, \text{TPER}(a,b,e,x)), x'')]]) \\
& \vee \neg \text{BEL}(a, \text{str}_3, t, (\forall x:T) (\forall x':T) [\text{BEF}(x, x') \& \\
& \quad \text{BEF}(x', \text{TPER}(a,b,e,x)) \rightarrow \text{HOLDS}(\text{SECURE}(b,e), x', \text{PER}(a,b,e,x))] \\
& \vee \neg \text{BEL}(a, \text{str}_4, t, (\forall x:T) (\forall x':T) [\text{OCCURS}(\text{COMMIT}(b,e,x'), x) \\
& \quad \rightarrow (\exists x'':T) [\text{BEF}(x, x') \& \text{BEF}(x', x'') \& \text{OCCURS}(e, x'')]]) \\
& \leftrightarrow \neg \text{TRUST}(a, \text{CALC}(\text{str}_1, \text{str}_2, \text{str}_3, \text{str}_4), b, t, e) \tag{10}
\end{aligned}$$

Note that the confidence levels now measure the lack of confidence.

Distrust is not the same as the trust not to do something. Thus, let us define two events **e** and **e'** as incompatible if whenever event **e** occurs, event **e'** does not occur and vice versa. More formally:

$$\text{INCOMPATIBLE}(e,e') \leftrightarrow (\forall t) [\text{OCCURS}(e,t) \leftrightarrow \neg \text{OCCURS}(e',t)] \tag{11}$$

Neither of the following hold:

$$(\text{INCOMPATIBLE}(e,e') \& \text{TRUST}(a,n,b,t,e)) \rightarrow \neg \text{TRUST}(a,n,b,t,e') \tag{12}$$

$$(\text{INCOMPATIBLE}(e,e') \& \neg \text{TRUST}(a,n,b,t,e)) \rightarrow \text{TRUST}(a,n,b,t,e') \tag{13}$$

In other words, the fact that **a** does not trust **b** to bring about **e**, does not mean that **a** trusts **b** to bring about the opposite of **e**, or the fact **a** trusts **b** to bring about **e** does not mean that **a** does not trust **b** to bring about the opposite of **e**. (12) does not hold because belief on my part in your competence, dependability, security and honesty with respect to **e** does not imply that I must believe you are either incompetent, undependable, insecure or dishonest with respect to the opposite of **e**. (13) does not hold because the fact that I do not believe in your competence, dependability, security or honesty with respect to **e** does not imply that I believe in your competence, dependability, security and honesty with respect to the opposite of **e**.

5.2 Trust Dynamism

A trust relationship evolves as new experiences and new information is added and incorporated.

There are two circumstances that can lead to a change in a trustor's trust level in a trustee to bring about some event. The first circumstance is simply a change in one of the constituent beliefs of the trust relationship, i.e., the belief that the trustor is competent, dependable, secure and honest. Thus, the trust that I had in WorldCom (in

2001) morphed into distrust when I discovered that their accounting practices were dishonest and misleading.

A second circumstance in which a trustor's level may change is the occurrence of some event, which although it changes the trust level, cannot immediately be traced to one of the constituent beliefs. For example, I may lose the trust in my insurance agent to get me the best rate on my car insurance, but I may not be sure whether this is due to his incompetence, or dishonesty. In order to model this situation, we introduce the predicate **IMPACTS** whose signature is $\Pi_{\text{IMPACTS}}: E, A, A, E, N$. Intuitively, $\text{IMPACTS}(e, a, b, e', r)$ is true if the occurrence of the event e impacts the level of trust that a has in b to bring about e' by r , where r is a ratio. The ratio is negatively correlated to the amount of confidence I already have. That is, the more I already trust you, the less my confidence increases when there is a new positive experience. This ensures that repeated occurrences of the same positive experience do not increase my confidence level to the same extent every time they occur. We can now introduce the following axiom to model trust dynamism:

$$(\forall t: T) [(\text{TRUST}(a, n, b, t, e) \ \& \ \text{OCCURS}(e', t) \ \& \ \text{IMPACTS}(e', a, b, e, w)) \rightarrow \text{TRUST}(a, n*(1+w), b, t+1, e)] \quad (14)$$

Thus, if a trusts b to perform e with trust level n at time t , and at time t some event e' occurs that impacts a 's trust in b to bring about e by ratio w , then a 's trust level in b to bring about e at the next point in time has changed by ratio w .

The notion of impact can also be used to distinguish between positive and negative experiences. An experience is an event that one has participated in. Trust interactions that are recorded and can be referenced in the future constitute one's experiences. The value of experience is that it can be used to enhance the quality of a decision. Experience is used to increase the probability of a favorable outcome and reduce the risk profile of a relationship. In essence, using experience to guide trust decisions makes life (and the trust decision process) simpler. A positive experience implies an increased confidence level, while a negative experience implies a lowered confidence level (even distrust). In other words, a positive experience is an event that positively impacts the trustor's trust level in the trustee, while a negative experience negatively impacts it. The trust level that a trustor has in a trustee with respect to some event is the result of the cumulative experiences that the trustor has had regarding the trustee with respect to the event.

There is a specific class of negative experiences that are worth mentioning. When a trusts b to bring about e , a believes in b 's competence, dependability, security and honesty with respect to e , and a therefore expects b to behave competently, dependably, securely and honestly with respect to e . b behaving competently, dependably, securely and honestly with respect to e is therefore likely to have no impact on a 's trust level. However, the same cannot be said for any violation of any of these expectations. Thus, if a experiences b behaving incompetently, undependably, unsecurely or dishonestly with respect to e , then a 's trust level in b to bring about e will be negatively impacted.

Finally, the experience that impacts the trustor's trust level in the trustee with respect to some action can be a purely cognitive experience, such as learning about an entity's reputation or receiving a recommendation from a third party. The concepts of experience, reputation, recommendations, business confidence, diffidence,

expectation, reliance and deception can all be modeled using the construct defined above.

5.3 Risk

There is a clear relationship between trust and uncertainty. After all, in general, when a trustor trusts a trustee to bring about an event, the trustor cannot be absolutely certain that a trustee will bring about the event. Because of the inherent uncertainty in a trust relationship, a trust relationship carries a certain risk, where risk can be regarded as the possibility or probability of incurring harm or loss. The extent to which a trustor is willing to take the risk associated with the trust relationship clearly depends on the trustor, the trustee and the nature of a particular transaction. Some trustors have a lower risk threshold than others, i.e. they are less risk-averse than others, and some trustees are inherently more trustworthy than others, no matter the transaction. Finally, trustors are more likely to trust trustees with respect to events that are less likely to cause harm to the trustor.

In order to model the relationship between risk and trust, we introduce the function symbol THRESHOLD with signature $\Sigma_{\text{THRESHOLD}}: A, A, E, T \rightarrow N$. Intuitively, $\text{THRESHOLD}(a, b, e, t)$ describes the upper bound on a 's willingness to take the risk to engage b to perform event e at time t . If a 's trust in b to perform e exceeds $\text{THRESHOLD}(a, b, e, t)$ then a is willing to engage b to perform e at time t as a is certain that any request for b to perform e will result in e occurring within a reasonable period of the request being made. In other words

$$\begin{aligned} (\forall t: T) [& (\text{TRUSTS}(a, n, b, t, e) \ \& \ (n > \text{THRESHOLD}(a, b, e, t) \ \& \\ & \text{OCCURS}(\text{REQUESTS}(a, b, e), t) \rightarrow \text{BEL}(a, (\exists t': T) \\ & (t' < \text{TPER}(a, b, e, t) \ \& \ \text{OCCURS}(e, t')))] \end{aligned} \quad (15)$$

Risk-averse entities will trust if their thresholds are not exceeded, whereas risk-loving entities do the opposite.

The important observation here is that confidence levels and risk thresholds are in the same numerical range and can be compared because there is a semantic equivalence. It should be noted that risk thresholds can be compared to explicitly derived risk values, from transaction-specific and event-driven trust decision factors.

Using Context in Security Design of a Search and Rescue System

Shivakant Mishra

Department of Computer Science
University of Colorado, Boulder, CO 80309-0430, USA
`mishras@cs.colorado.edu`

Abstract. With the emergence of small devices equipped with wireless communication, several sophisticated systems for search and rescue have been proposed and developed. However, a key obstacle in a wide deployment of these systems has been users' security and privacy. On one hand, such systems need to collect as much information about a user as possible in order to locate that user in a timely manner. On the other hand, this very capability drives users away from using such a system. This paper describes work-in-progress in building a security and privacy framework for CenWits, which is a new search and rescue system for people in emergency situation in wilderness areas. The paper focuses on the role of context in building this framework.

1 Introduction

Search and rescue of people in emergency situation in a timely manner is an extremely important service. In the past, it was difficult to build such a service because of a lack of timely information needed to determine the current location of a person who may be in an emergency situation. However, with the emergence of small computing devices such as PDAs, sensors and cell phones that have wireless communication capabilities, it has become feasible to build such a system. Indeed, several such systems have been proposed and prototypes of some of them have been implemented over the last five years [12, 1, 5, 6, 4].

We have designed and implemented a search and rescue system called CenWits[21] for a wilderness environment. A key differentiating feature of CenWits from the other recent search and rescue systems is that it is designed for a wilderness environment. In such an environment, there is no Internet connectivity, no cellular network, and building an adhoc network is infeasible due to an extremely sparse environment. Furthermore, GPS reception is only available at limited areas. CenWits (**C**onnection-less **S**ensor-Based Tracking System Using **W**itnesses) is comprised of three components: (1) mobile, in-situ sensors that are worn by people (e.g. hikers); (2) access points (AP) that collect information from these sensors; and (3) GPS receivers and location points (LP) that provide location information to the sensors. A subject uses GPS receivers (when it can connect to a satellite) and LPs to determine its current location. The key idea of CenWits is that it uses a concept of *witnesses* to convey a subject's movement

and location information to the outside world. This averts a need for maintaining a connected network to transmit location information to the outside world. In particular, there is no need for expensive GSM or satellite transmitters, or maintaining an adhoc network of in-situ sensors in CenWits.

Since a search and rescue system like CenWits must track the movement of people, there are some very obvious and important security and privacy issues. In fact, such systems must cope with two conflicting issues. On one hand, the system requires a collection of as much information about the location and movement of a person as possible. This is to ensure that a smaller and more accurate search area may be determined in case that person goes missing, or is in emergency situation. Indeed, it is in the interest of a person to give out as much information as possible about his/her location and movement to improve his/her chances of being located and rescued in case of emergency. On the other hand, a majority of people are not comfortable in giving out too much information about their location and movement for the fear that such information may be misused for malicious purposes, e.g. stalking. Indeed, this latter reason has proved to be a major hindrance in a wider deployment of CenWits.

It is clear that appropriate security and privacy support must be provided in a search and rescue system for wide acceptance. At present, we are designing a security framework for CenWits. In this paper, we describe a preliminary design of this framework, and discuss some important issues in the design and implementation of security and privacy support for a search and rescue system in general. An important observation is that there is no single security model that can be applied in such a system. The required security and privacy support varies based on individuals as well as context. In this paper, we focus on the role of context in the design of a security framework.

The rest of this paper is organized as follows. Section 2 provides a brief overview of CenWits. Section 4 describes the role of context in building a security and privacy framework for CenWits. Section 5 provides a high-level description of a multi-level security and privacy framework for CenWits. Finally, Section 6 concludes the paper.

2 CenWits: A Brief Overview

CenWits is a search and rescue system that makes use of smaller and cheaper sensor devices. It has several important advantages over the other search and rescue systems. These advantages include a loosely-coupled system that relies only on intermittent network connectivity, power and storage efficiency, and low cost. It solves one of the greatest problems plaguing modern search and rescue systems: it has an inherent on-site storage capability. This means someone within the network will have access to the last-known-location information of a victim, and perhaps his bearing and speed information as well. It utilizes the concept of witnesses to propagate information, infer current possible location and speed of a subject, and identify hot search and rescue areas in case of emergencies.

The concept of witness works as follows. Whenever two or more hikers are within a close range (say 100 meters) of one another, their sensors can exchange messages over a radio frequency. When two hikers, say *A* and *B* are in close range of each other, the following message exchange takes place. *A* generates a *witness record* that stores the following information: *B* was seen at this location at this time. Similarly, *B* generates a witness record storing *A* was seen at this location at this time. In addition, *A* sends all witness records in his/her memory to *B*, and similarly, *B* send all witness records in his/her memory to *A*. Whenever a hiker comes in close range of an access point (special computing devices that have Internet connection to a control center), he/she dumps all witness records in his/her memory to the access point.

A prototype of CenWits has been implemented using MICA2 sensor 900MHz running Mantis OS 0.9.1b. We have experimented with it in a number of indoor and outdoor environments.

3 Related Work

A survey of location systems for ubiquitous computing is provided in [11]. These include [17], [19], [8], and [1]. These systems are mainly designed for an indoor environment, and not useful for our purpose. A system that is viable in suburban area where a user can see clear sky and has GSM cellular reception at the same time is [5]. Since, cellular reception is not available in wilderness areas, this system is not useful for our purpose.

Personal Locator Beacon (PLB) that uses RF transmitter has been used for avalanche rescuing for years. Luxury version of PLB that combines a GPS receiver and a COSPAS-SARSAT satellite transmitter this also available [4]. Both of these devices are impractical in wilderness environment because of significantly large space and/or expensive satellite transmitter. Another related technology in widespread use today is the ONSTAR system [3], typically used in several luxury cars. Like PLBs, this system has several limitations for use in a wilderness environment, including heavy-weight, expensive, and requirement for a connected network. The Lifetch system uses a GPS receiver board combined with a GSM/GPRS transmitter and an RF transmitter in one wireless sensor node called Intelligent Communication Unit (ICU). Again, this system requires a connected network, which is not possible in a wilderness environment.

As far as we know, there is no work done yet in the area of building a security and privacy framework for a search and rescue system in a wilderness environment. The main difficulty in building this framework is the absence of any kind of infrastructure, be it a communication network or a public key infrastructure. Furthermore, the actual level of security and privacy needed in this environment varies based on circumstances as well as individuals' perception of danger. Our goal is make use of existing security principles and techniques, and adapt them meet the requirements of our system.

4 Motivation and Context Awareness

The amount of security and privacy support that a search and rescue system in a wilderness area should provide at any specific moment is strongly context dependent. There are two types of contexts involved here: situational and personal. Situational context refers to the level of danger as perceived by a person. Personal context refers to the level of comfort that a person has in divulging information about his/her movement. As a part of situational context, some situations may be perceived more dangerous than others, e.g. a sudden storm or a flash flood. As a part of personal context, some people may not care at all if their location and movement are being tracked. On the other hand, some other people may be extremely sensitive about their location and movement being tracked.

There are two important observations that motivate a multi-level security and privacy framework for CenWits. First, individual sensitivity towards tracking location and movement information varies based on the situation. Second, a lower level of security and privacy support generally translates to faster propagation of one's location and movement information to the control center. This is because a lower level of security generally incurs lower overhead. As a result, more witness records may be exchanged between hikers during an encounter. This in turn implies that the chances of locating a person in a short period of time improve if a that person uses a lower level of security and privacy. Thus, a person is more likely to accept a lower level of security and privacy, if he/she feels that there is danger. Hence, context awareness plays a crucial role in building a security and privacy framework that is acceptable to a large number of users.

5 Security Framework

It is clear that a single security model cannot address the needs of everyone under all situations. A multi-level security framework is an appropriate framework for a search and rescue system. The basic idea is that the system provides several different levels of security and privacy support, and an individual chooses an appropriate level based on his/her comfort level. Furthermore, because situations may change over time, the system also provides appropriate support that allows a user to switch from one security level to another based on the current situation.

Our preliminary design consists of five levels of security and privacy. Figure 1 illustrates these five levels along with effect of situational and personal contexts. Each of these five levels are designed to protect the location and movement information of an individual against a different type of adversary. Level 0 refers to an absence of any security or privacy support. In this level, the identity and movement information of a person is propagated without any attempt to hide it or prevent it from tampering. A user will typically opt for this level when he/she is in extreme danger. This level ofcourse has absolutely no impact on the speed at which tracking and movement information is propagated to the control center. Messages are exchanged in clear, and no sender or receiver authentication is needed. Thus, there is no security overhead.

Level 1 security provides support for protecting individual information from outsiders. At this level, hikers can exchange their location and movement information with one another with the guarantee that this information is not leaked out to an outsider. However, such a protection is not provided from the insiders, i.e. an insider (another hiker) can track the identity, location and movement of another hiker. Essentially, level 1 security has been designed to provide protection from a *passive outsider adversary*. A passive adversary is one who can simply listen to the network communication without any capability to actively insert packets, modify packets, or launch an attack against the network infrastructure. An implementation of this level requires sharing a single symmetric key among all hikers, access points and control center. All information exchanged in the system is encrypted using this key and an appropriate symmetric key protocol, e.g. AES. Level 1 requires encrypting each message before being sent. This naturally incurs some performance overhead compared to level 0.

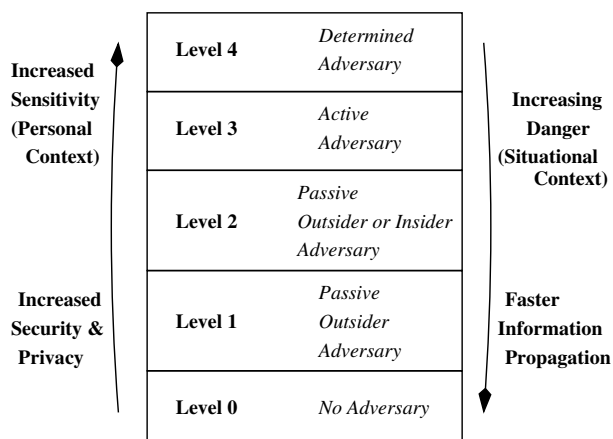


Fig. 1. A Multi-Level Security & Privacy Framework.

Level 2 security provides protection from a *passive adversary*. Thus, in addition to passive outsider adversaries, level 2 security provides support for protecting individual information from passive insider adversaries as well. A passive insider adversary is one who joins the system legitimately, and records all information it can receive from the network. However, like a passive outsider adversary, a passive insider adversary does not actively seek information from other insiders, e.g. by injecting spurious packets, or modifying packets. An implementation of this level requires that each hiker share a separate symmetric with the control center. This will ensure that a witness record generated by a hiker can only be read by the control center. Additional mechanisms are needed to detect duplicate witness records. This level incurs more security overhead than level

0, because in addition to encrypting each message, mechanisms are needed to detect duplicate witness records and manage keys for each hiker.

In addition to passive adversaries, level 3 security provides support for protecting individual information from active adversaries. An active adversary is one that actively seeks individual information by targeting individuals. Such an adversary may insert spurious packets or replay packets to gain information. This adversary may be an insider or an outsider. An implementation of this level requires each hiker share a separate key with the control center. In addition, it requires incorporating mechanisms to authenticate sender's identity, determine message integrity, and detect man-in-the-middle and replay attacks. As a result, level 3 incurs more performance overhead than level 2.

Finally, level 4 security support protects the entire system from *determined adversaries*, whose goal is to degrade or destroy the complete system, e.g. by launching a denial of service attack. Such adversaries are extremely powerful and may employ very expensive and sophisticated attack methods. Defending against such adversaries is very difficult, and likely to significantly increase the cost of the entire system. Furthermore, defending against such adversaries is likely to significantly slow down the propagation of witness information to the control center. At this time, we believe that the existence of determined adversaries is extremely unlikely in a search and rescue system in a wilderness environment. Hence, we are not planning to implement this level in CenWits at this moment.

6 Discussion and Current Status

Based on this multi-level security framework, a relatively carefree individual can set his sensor to level 0 or 1 security, while a security-sensitive person can set his sensor to level 2 or 3. It is very important that a user understands the implications of setting his/her sensor to a particular level. This means that if a user sets his/her sensor to level 3, the user understands that his/her location and movement information will move relatively slowly to the control center, and that means a search and rescue team will have relatively less accurate information about his/her location in case of emergency. Similarly, if a user sets his/her sensor to level 0, the user understands that his/her location and movement information can be tracked by relatively less sophisticated individuals. However, a search and rescue team will have a more accurate information about his/her location in case of emergency.

A user can change his/her security level setting at any time. So, if a user perceives that the current situation has turned dangerous at any time, e.g. if a heavy rainfall starts, he/she may switch to a lower sensor security level to ensure that his/her location and movement information is propagated faster. An interesting question is if this switching of security levels can be automated. If the system detects that the environment has become dangerous, e.g. there has been a severe thunderstorm warning, or an avalanche has occurred, it may automatically lower the security level chosen by a user. This capability is especially important in emergency situations when a person may not be in a position to manually

switch his/her sensor security level. This will require equipping the sensor device with appropriate sensing capabilities that can sense such dangerous situations. At present, we are in the final stages of completing our design of this multi-level security framework for CenWits based on these ideas.

References

1. 802.11-based tracking system. <http://www.pangonetworks.com/locator.htm>.
2. Brent geese 2002. <http://www.wwt.org.uk/brent/>.
3. The onstar system. <http://www.onstar.com>.
4. Personal locator beacons with GPS receiver and satellite transmitter. <http://www.aeromedix.com/>.
5. Personal tracking using GPS and GSM system. <http://www.ulocate.com/trimtrac.html>.
6. Rf based kid tracking system. <http://www.ion-kids.com/>.
7. F. Alessio. Performance measurements with motes technology. *MSWiM'04*, 2004.
8. P. Bahl and V. N. Padmanabhan. RADAR: An in-building RF-based user location and tracking system. *IEEE Infocom*, 2000.
9. K. Fall. A delay-tolerant network architecture for challenged internets. In *SIGCOMM*, 2003.
10. L. Gu and J. Stankovic. Radio triggered wake-up capability for sensor networks. In *Real-Time Applications Symposium*, 2004.
11. J. Hightower and G. Borriello. Location systems for ubiquitous computing. *IEEE Computer*, 2001.
12. W. Jaskowski, K. Jedrzejek, B. Nyczkowski, and S. Skowronek. Lifetch life saving system. *CSIDC*, 2004.
13. P. Juang, H. Oki, Y. Wang, M. Martonosi, L. Peh, and D. Rubenstein. Energy-efficient computing for wildlife tracking: design tradeoffs and early experiences with ZebraNet. In *ASPLOS*, 2002.
14. K. Kansal and M. Srivastava. Energy harvesting aware power management. In *Wireless Sensor Networks: A Systems Perspective*, 2005.
15. G. J. Pottie and W. J. Kaiser. Embedding the internet: wireless integrated network sensors. *Communications of the ACM*, 43(5), May 2000.
16. S. Roundy, P. K. Wright, and J. Rabaey. A study of low-level vibrations as a power source for wireless sensor networks. *Computer Communications*, 26(11), 2003.
17. C. Savarese, J. M. Rabaey, and J. Beutel. Locationing in distributed ad-hoc wireless sensor networks. *ICASSP*, 2001.
18. V. Shnayder, M. Hempstead, B. Chen, G. Allen, and M. Welsh. Simulating the power consumption of large-scale sensor network applications. In *Sensys*, 2004.
19. R. Want and A. Hopper. Active badges and personal interactive computing objects. *IEEE Transactions of Consumer Electronics*, 1992.
20. M. Welsh and G. Mainland. Programming sensor networks using abstract regions. *NSDI '04*, 2004.
21. J.-H. Huang, S. Amjad, and S. Mishra. CenWits: A Sensor-Based Loosely Coupled Search and Rescue System Using Witnesses. In *SenSys'05*, San Diego, CA, November 2005.

Towards Privacy Protection in a Middleware for Context-awareness

(Short Paper)

Linda Pareschi, Daniele Riboni, Sergio Mascetti, and Claudio Bettini

D.I.Co., University of Milan
via Comelico, 39, I-20135, Milan, Italy
{pareschi,riboni,mascetti,bettini}@dico.unimi.it

Abstract. Privacy is recognized as a fundamental issue for the provision of context-aware services. In this paper we present work in progress regarding the definition of a comprehensive framework for supporting context-aware services while protecting users' privacy. Our proposal is based on a combination of mechanisms for enforcing context-aware privacy policies and k -anonymity. Moreover, our proposed technique involves the use of stereotypes for generalizing precise identity information to the aim of protecting users' privacy.

1 Introduction

The recent proliferation of powerful mobile devices, wireless networks, and sensing technologies has enabled new classes of context-aware services, e.g., services that adapt themselves to the current situation of the user. However, since adaptation involves the communication to the service provider of user's private information such as her location, activity, and profile data, in order to be accepted by final users these services must be supported by mechanisms for preserving privacy. In the literature about privacy in location-based services (LBS) it has been shown that simply hiding users' explicit identifiers (e.g., user's name) may not be sufficient to guarantee privacy, since the user's identity can be possibly inferred from the other information sent to the service provider. To this aim, several approaches have been proposed for enforcing access control, or for guaranteeing anonymity. On the basis of the experience we have acquired in the last years while working on a framework for context-awareness [1], and on privacy in LBS [2], we argue that a satisfactory comprehensive solution for privacy in context-awareness is still missing. As a matter of fact, we believe that a solution based solely on access control is unsuitable for many services, since simply negating the access to a given context data (e.g., location) would determine the impossibility of providing the service at all (e.g., a LBS). On the other hand, work on anonymity for context-aware services has generally concentrated on solely location, while the set of data that can be useful for adaptation is much wider. In this paper we present work in progress regarding the definition of a comprehensive framework for privacy protection in context-awareness.




Fig. 1. Privacy threat and use-case scenario

External knowledge assumptions The privacy issue we are addressing, shown in Figure 1(a), consists in preventing the attacker from inferring the *sensitive association*, i.e. the association between the user identity and the *private information* (*PI*) of that user. According to the knowledge an attacker can acquire from external sources, data included in requests may increase his ability to reconstruct the sensitive association (these data are called *quasi-identifiers* (*QI*) [3]).

Techniques for privacy preservation strongly depend on which knowledge an hypothetical attacker can access. Therefore, in order to provide effective mechanisms for protecting the user privacy, external knowledge assumptions must be formalized with respect to the application logic of the required service and the privacy threat to be contrasted.

2 Requirements and proposed solution

In order to illustrate how we intend to address privacy issues, we consider the following running example:

Example 1. Consider a context-aware service providing points of interest (*POIs*) to mobile users, depending on data such as location, profile, and interests. Suppose that a hypothetical user Alina – a PhD student – is submitting a request to that service using her GPS-enabled smart phone, in order to find a bookshop.

The service presented in the above example is representative of a large number of mobile services, in which adaptation is performed considering not only location, but a wide set of context data.

Since we claim that formal knowledge assumptions are fundamental for defining a privacy preservation technique, we couple the running example described above with the following illustrative knowledge assumption:

- Γ_C : the attacker can acquire knowledge about context data (including spatio-temporal information and profile data), either directly from external sources, or inferring them from other public data.

Assuming Γ_C , each context data included in a requests can possibly act as QI .

According to this external knowledge assumption, in order to preserve Alina’s privacy it is necessary to hide either her precise location, and her profile and interests. Approaches based on *access control techniques* may determine the negation of that kind of data to the service provider. However, such a solution would not be suitable to the above mentioned service – and, in general, to any context-aware service– since spatio-temporal and context data are essential for selecting a set of resources that can be potentially interesting to the user.

A better solution would consist in the application of *obfuscation techniques*; i.e., in providing generalized data instead of their precise values. For instance (see Figure 1(b)), instead of communicating the exact information regarding Alina, it would be possible to provide the POI service with an area $A1$ including her precise location (e.g., the block where she currently is), her status (i.e., a PhD student), and her interests I . This solution would allow the service provider to select a set of POIs that are actually close to Alina’s interests and location. However, this approach does not prevent the attacker from identifying the issuer of the request. In fact, Alina could be identified if the attacker has the knowledge of Alina being the only PhD student in the area $A1$ having interests I .

Therefore, we have identified the following requirements for a privacy preservation middleware addressed to users of context-aware services:

- (a) A flexible mechanism for disclosing obfuscated context data, still guaranteeing a given level of privacy, is mandatory;
- (b) Personal data acting as QI must be generalized in a meaningful way in order to allow the service provider to adapt the service;
- (c) When location is not the only context data to be provided, the obfuscation process must determine the level of generalization of each context data in order to fulfill either privacy and adaptation requirements.

In order to protect the sensitive association between the user’s identity and her private information, obfuscation techniques can be applied either to QIs, to PIs, or to both QIs and PIs. However, as a first effort towards the definition of a comprehensive framework for privacy protection in context-awareness, in this work we focus our investigation on the first approach.

2.1 k -Anonymity for users of context-aware services

In order to address requirement (a), we adopt the *k-anonymity* technique already introduced in database systems ([4]) and applied to LBS [5, 2]. The intuition behind *k-anonymity* consists in making the user indistinguishable in a set of k potential issuers. According to Γ_C , any context data may act as QI ; therefore, in order to enforce *k-anonymity*, any context data must be considered in the generalization process.

Generalization depends on the semantics and representation of context data; for instance, if the data is represented by one (or more) numerical values (e.g., GPS coordinates), the exact data is generalized to an interval (or to an area). On the contrary, if the data is represented by a specific value v belonging to

a taxonomy T , it can be generalized to an ancestor of v in T . For example, given a possible stereotype hierarchy, the exact stereotype of Alina (i.e., *PhD student*) can be generalized to its parent *student*. Since the problem of optimal multidimensional anonymization is NP-hard, we plan to adopt an approximation algorithm for addressing multidimensional *k-anonymity*.

2.2 Identity generalization through stereotypes

In order to protect the user’s privacy, any data that uniquely identifies the user must be removed from the user request. Hence, the most commonly adopted solution consists in substituting the *user-id* (together with any possible user personal data) with an anonymous *pseudo-id*, which can be used by the application logic for performing tasks such as authentication and session management. However, since in such a solution any reference to the user personal data are lost, the service provider is no longer able to customize the service according to data such as the user age, gender, and formal education, which are considered very relevant for accurately personalizing services.

According to the requirement (b), we propose to extend the *pseudo-id* approach by the use of *stereotypes* [6] for generalizing the user’s identity. Stereotypes are useful abstractions for characterizing users’ demographic data, personalities, and goals. Our solution allows the application logic not only to perform authentication and session management, but also to customize the service according to relevant (generalized) personal data, while preserving anonymity.

2.3 Context-aware privacy policies

The generalization of context data obviously impacts on the quality of adaptation. Indeed, even if *refinement* techniques can be used for improving the service response, it could happen that, in order to achieve the desired anonymity level, context data become too general to provide the service at an acceptable quality level [7]. In order to address this issue, we allow users to declare privacy policies about the generalization of single context data, depending on the context itself:

Example 2. Alina declared the following privacy policies:

p1: If *activity*==*working* Then *anonymity-level*:=*high*

p2: If *activity*==*shopping* Then *anonymity-level*:=*low*

p3: If *activity*==*walking* Then *provide-accurate-location*

Alina asks the POI server for a bookshop. Since she is currently strolling for shopping, policies **p2** and **p3** hold: then, her desired anonymity level is set to a low value (corresponding to 4-anonymity), and the anonymizer provides a quite accurate location (in order to select resources that are actually close to her). Hence, in order to enforce Alina’s policies, the anonymizer generalizes Alina’s exact stereotype to *student* for obtaining a smaller area $A1$ containing 4 potential issuers (see Figure 1(b)), and communicates those data to the POI server.

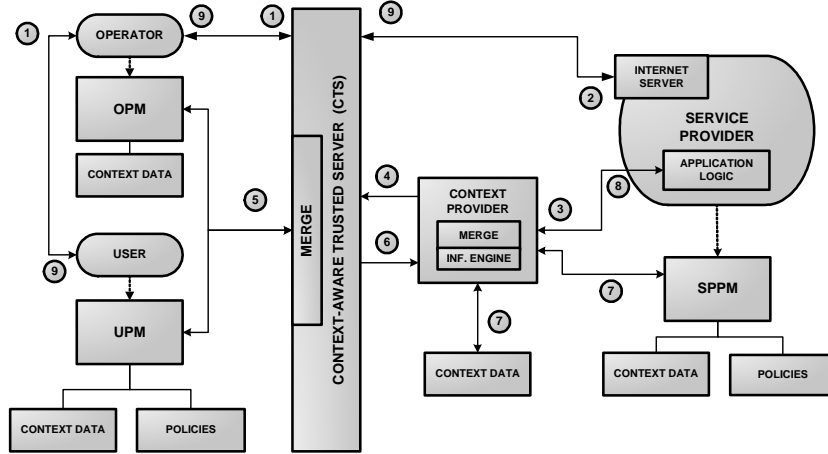


Fig. 2. Architecture overview

Moreover, for each context data the user can define the maximum granularity of generalization on the basis of context (e.g., *if I am walking, then generalize location to at most 500m*). Since it is not always possible to conciliate these constraints with the desired level of anonymity, a conflict resolution mechanism is adopted for choosing if either guaranteeing a lower anonymity level, generalizing the context data to a higher granularity, or not providing the service at all.

2.4 Architecture overview

In order to illustrate how privacy preserving techniques can be integrated into an architecture for context-awareness, in this section we consider an extension of the CARE middleware [1]. CARE supports the acquisition of context data from different sources, the reasoning with this data based on distributed policies, and the reconciliation of possibly conflicting information. Figure 2 depicts an extension of CARE to support privacy: the module devoted to apply our privacy preservation techniques (called CONTEXT-AWARE TRUSTED SERVER (CTS)) acts as an intermediary between the user trusted domain (left-hand side) and the rest of the world (right-hand side). The user trusted domain is constituted by the USER PROFILE MANAGER (UPM) – which manages user policies, and context data explicitly provided by the user or acquired from user-side sensors – and by the NETWORK OPERATOR PROFILE MANAGER (OPM), which manages context data such as users’ location and network status.

Numbers in Figure 2 represent the data flow upon a user request: each user request (1) is filtered by the CTS, which transforms the user ID into a *pseudoID* – which is used to identify the user request and to perform authentication – and removes any QI, before forwarding the request (2). Next (3), the service provider asks for the context data it needs for adapting the service to a central module called CONTEXT PROVIDER, which forwards the request to the CTS (4).

The CTS retrieves user’s privacy policies, and distributed context data from the user trusted domain, it merges context data solving possible conflicts, and generalizes them according to the privacy policies (5). Then (6), it sends those data – together with the (possibly generalized) user stereotype – to the CONTEXT PROVIDER, which merges them with context data retrieved from the SERVICE PROVIDER PROFILE MANAGER (SPPM) and external context sources (7), and evaluates service provider policies, thus obtaining the aggregated context data that are communicated to the application logic (8). Finally (9), the application logic adapts the service and communicates the service response to the CTS, which forwards it to the user. Note that, in order to avoid eavesdropping, any communication between the CTS, the service provider, and the CONTEXT PROVIDER is encrypted with public key cryptography.

3 Future work

This preliminary investigation has highlighted several important requirements that will be the object of our future research efforts. In the following we mention the most important aspects we are planning to consider: *i)* in order to allow users to define their own privacy preferences depending on context, a sufficiently expressive language for context-aware privacy preferences is needed; *ii)* in order to estimate the quality of service according to the chosen privacy policies, a mechanism for measuring the trade-off between adaptation quality and anonymization degree must be devised; *iii)* an efficient mechanism for multidimensional context data generalization must be defined; *iv)* effective privacy techniques must be applied to a *dynamic case*, i.e., when an attacker is able to reconstruct the sensitive association by means of requests issued by the same user in different time intervals.

References

1. Agostini, A., Bettini, C., Cesa-Bianchi, N., Maggiorini, D., Riboni, D., Ruberl, M., Sala, C., Vitali, D.: Towards Highly Adaptive Services for Mobile Computing. In: Proc. of IFIP TC8 Conf. on Mobile Information Systems, Springer (2004) 121–134
2. Bettini, C., Mascetti, S., Wang, X.S., Jajodia, S.: Anonymity in Location-based Services: towards a General Framework. In: Proc. of the 8th Int. Conf. on Mobile Data Management (MDM), IEEE Computer Society (2007)
3. Samarati, P.: Protecting Respondents’ Identities in Microdata Release. IEEE Trans. on Knowledge and Data Engineering **13**(6) (2001) 1010–1027
4. Sweeney, L.: k-Anonymity: a Model for Protecting Privacy. Int. J. Uncertain. Fuzziness Knowl.-Based Syst. **10**(5) (2002) 557–570
5. Mokbel, M.F., Chow, C.Y., Aref, W.G.: The New Casper: Query Processing for Location Services without Compromising Privacy. In: Proc. of the 32nd Int. Conf. on Very Large Data Bases (VLDB), VLDB Endowment (2006) 763–774
6. Rich, E.: User Modeling via Stereotypes. Cognitive Science **3**(4) (1979) 329–354
7. Aggarwal, C.C.: On k-Anonymity and the Curse of Dimensionality. In: Proc. of the 31st Int. Conf. on Very Large Data Bases (VLDB), ACM (2005) 901–909

Context-Aware Management Domains

Ricardo Neisse¹, Maarten Wegdam, Patrícia Dockhorn Costa, and Marten van Sinderen

CTIT, University of Twente, The Netherlands
{r.neisse, m.wegdam, p.dockhorncosta, m.j.vansinderen}@utwente.nl

Abstract. In this paper we extend the concept of management domains to a new concept called Context-Aware Management Domains (CAMDs). CAMDs enable context-aware management of policies allowing the grouping of entities based on context information. Since context is dynamic, so is the domain membership. As a consequence, the association of policies with the entities in the domain also becomes dynamic. In this paper we provide CAMD examples and an information model together with a discussing on our ongoing implementation for our target context-aware service platform.

1 Introduction

Context aware services adapt themselves to the current user's situation. An example of this is a tourist service which uses the current user location, activity, and preferences to personalize tourist advices. In order to support context awareness, service platforms have been designed to support context information acquisition, reasoning and distribution [1].

Typical context-aware service platforms have thousands or millions of entities (users, service providers, context providers, etc.) and different types of policies have to be managed. Policies are required, for instance, to control access to context information, to enforce user's privacy, and to manage trust relationships among the entities. Due to the complexity, dynamicity, and large number of entities, the specification of these policies can easily become unmanageable.

Standard policy management tools ease the policy management, however, the problem with these tools is that they provide either static management capabilities (e.g. management domains [2]), or, if there is some form of dynamic management, this is limited to one specific area (e.g., X-RBAC [3], [4], and [5] for access control and COMITY [6] for trust management). For this reason, these policy management tools do not fulfill the dynamic policy requirements of context-aware service platforms.

In a context-aware service platform, policies are defined based on the context of the entities. One example is a privacy policy stating that "Bob's identity should not be anonymized for nearby persons". In this case, "nearby persons" refers to a set of

¹ Supported by CNPq Scholarship – Brazil

