
Initial trust formation in Virtual Organisations

**Tatyana Ryutov*, Clifford Neuman, Li Zhou
and Noria Foukia**

Information Sciences Institute,
University of Southern California,
4676 Admiralty Way, Suite 1001,
Marina del Rey, CA 90292, USA
Fax: (310) 823-6714 E-mail: tryutov@isi.edu
E-mail: zhoulia2000@hotmail.com
E-mail: nfoukia@infoscience.otago.ac.nz
Website: <http://clifford.neuman.name>
*Corresponding author

Abstract: We present a conceptual framework that introduces key concepts capturing initial trust establishment in on-demand Virtual Organisations (VO). This framework can be used to develop trust formation protocols and policies. The framework can serve as a basis for implementing an automated system that facilitates the establishment of a VO, considerably reduces the effort for setting up a VO and consequently reduces the VO's time to operation. A novel aspect of the initial trust establishment described in this paper is the consideration of the mutual trust effects of the participants' behaviour during the trust negotiation process.

Keywords: virtual organisation; VO; trust formation.

Reference to this paper should be made as follows: Ryutov, T., Neuman, C., Zhou, L. and Foukia, N. (2007) 'Initial trust formation in Virtual Organisations', *Int. J. Internet Technology and Secured Transactions*, Vol. x, No. x, pp.xxx-xxx.

Biographical notes: Tatyana Ryutov received an MS Degree in Applied Mathematics from Moscow State University, Russia, in 1991, and MS and PhD Degrees in Computer Science from the University of Southern California (USC), in 1999 and 2002, respectively. She joined USC/ISI in 1996 working as a Graduate Research Assistant, and focused on the development and implementation of the access control framework for distributed systems that supports active policies adaptive to network threat conditions. Currently, she is working as a Computer Scientist at the Information Sciences Institute of the USC.

Clifford Neuman received an SB from MIT and MS and PhD Degrees from the University of Washington. He is Director of the USC Centre for Computer Systems Security, a researcher at the Information Sciences Institute (ISI), and a faculty member in the Computer Science Department at the University of Southern California (USC). He is principal designer of the Kerberos authentication system, developed the NetCheque and NetCash electronic payment systems, and the Prospero Directory Service.

Li Zhou received the BS Degree in Computer Science from Peking University in 2001. He is a PhD student in the Computer Science Department, University of Southern California (USC). Currently, he is working in the Generic Operating System Technology (GOST) Group, USC Information Science Institute, as a Graduate Researching Assistant. His major research area is trust-based decision making of distributed system.

Noria Foukia received a BS and MS Degrees in Pure and Applied Mathematics in 1995 and 1996, from the University of Science Lyon-I, France, and a MS Degree in Network and Distributed Systems from the University of Nice-Sophia Antipolis, France, in 1998. She joined the Teleinformatics and Operating Systems groups of the University of Geneva, Switzerland as Research Assistant, working on charging and accounting for the internet and on network security, and received her PhD Degree in 2004. She did a post-doc at the ISI/USC, USA from 2004 to 2005. She is currently a Lecturer at the University of Otago, New Zealand. Her research interests include network security, privacy, trust management.

1 Introduction

Business networks are evolving into VOs consisting of independent service providers that combine resources to exploit a particular market opportunity. Establishing trust between VO participants is the basis for successful collaboration and joint performance. However, the dynamic, temporal and often competitive nature of VOs presents substantial challenges to trust formation (Ishaya and Mundy, 2004). Furthermore, the lack of face-to-face communication and often-insufficient information about VO participants inhibits traditional trust building techniques. Research on trust formation in communication environments such as VOs, grids, and the semantic web is in its early stages (Almenarez et al., 2004; Cardoso and Oliveira, 2004; Davulcu et al., 1999; Dellarocas et al., 2002; Grefen et al., 2000; Periorellis and Parastatidis, 2005; Rocha and Oliveira, 2001; Tagg and Quirchmayr, 2001).

This paper considers initial negotiation that involves two participants wishing to establish a VO. For example, consider two companies (that have never cooperated before) wishing to establish a VO for a commercial collaboration, such as computer assembly. One of them produces desktops, the other one manufactures monitors. Both parties are interested in maximising their share of the profits on the end product sales. In order to commit to a joint venture, the companies have to establish initial trust. The trust building process is comprised of two stages: goal verification and agreement negotiation. First each party needs to decide whether the trade-off between risk and expected return is acceptable. The parties must agree on a common goal (product assembly) and verify that incentives are in place to achieve this goal (the end product consists of resources that belong to each company). Next, the parties need to negotiate a mechanism to share the profit. During the negotiations, parties exchange offers (agreement proposals) and counter-offers until the agreement is reached (initial trust is established) or failure is reported. While the agreement enforcement is a vital part of the VO infrastructure, the discussion of the monitoring and enforcing mechanisms is outside of the scope of this paper. We do not consider trust evolution during the operation of the established VO.

Our framework can be used to develop trust formation protocols and policies. The framework can serve as a basis for implementing an automated system that facilitates the establishment of a VO, considerably reduces the effort for setting up a VO and consequently reduces the VO's time to operation.

During the formation of a VO, participants have a choice of service providers with whom they may cooperate. Each party must decide whether it is willing to engage in a business relationship with the other party. During the initial phase, when parties become aware of each other, trust is at its most fragile state: individual's tentative trust assumptions could be easily destroyed by one's early beliefs. In the period after the initial one – the continuing relationship – the parties interact and their judgements about each other become a function of the interactions themselves (McKnight et al., 1998).

Trust exists in the context where there are benevolent and malicious behaviour and risks involved in trusting the other party (Gray et al., 2003; Josang, 1996; Josang and Presti, 2004). In order to establish trust, the parties must exhibit interdependence: one party is not able to fulfill its function or achieve its objectives without action by the other party. In this situation, each party desires assurance that the other's actions will enable their interdependent objectives to be met. The initial trust formation process consists of the two phases:

- *goal verification*: the participants must agree on a mutually beneficial goal in order to justify a VO creation and confirm interdependence – a necessary precondition of trust
- *agreement negotiation* to determine and agree on what constitutes benevolent/malicious behaviour and how to reach the goal.

Although trust formation is essential, it is not sufficient for the prevention and detection of fraudulent transactions. Sanctions are used for enforcement of benevolent behaviour.

It is difficult to give an exact definition of trust due to its complex subjective nature. However, we believe that trust in VOs can be explicitly modelled using four constructs: expectations, obligations, agreements and suspicion levels.

In our model, trust reflects intention of an entity *a* to accept vulnerability/risks based on positive *expectations* of the intentions of entity *b* to fulfill commitments (*obligations*), demonstrate honesty in negotiations and refrain from excessive advantage. *Obligations* represent the participant's commitment to provide a service under certain terms and conditions to other participant. The concept of an *agreement* is used to make an explicit declaration of the expected and malicious behaviour defined in the participants' obligations, and sanctions in the case of non-conformance to the obligations. When agreement is reached, the *initial trust* is established. By the *initial trust* we mean the *decision* of both parties to accept the risks and willingness to participate in alliance. Trust may change in the course of VO interactions.

The framework represents trust in a manner that captures human intuition, such that positive outcomes of interactions preserve or amplify trust, while trust erodes with negative experiences. Distrust (or *suspicion level*) specifies a means of revoking previously established trust based on observed behaviour. Parties tend to confirm their early belief that the other party is trustworthy or not by evaluating the other's actions and attributing them to the benevolent/malicious behaviours (Falcone and Castelfranchi, 2004). The suspicion level is determined by monitoring the 'lack of

goodwill' that appears as signs of non-cooperation and the partner's proactively opportunistic behaviours during the trust establishment process.

A novel aspect of the initial trust establishment considered in this paper is monitoring the participant's behaviour during the agreement negotiation process and adjusting trust values based on the perceived behaviour. By modelling the agreement negotiation in this way, the relationship between the trust, initial participants' behaviour, expectations, obligations and agreements become more readily apparent for future VO work.

2 Key concepts

To model the trust establishing phase between two entities a and b , we consider the following structures:

- participants a and b that wish to create a VO
- suspicion levels sl_a and sl_b attributed to the participants a and b
- mutual goal G
- negotiation history $H_{a,b}$ maintained locally by the participants a and b
- agreement A between a and b .

The next sections briefly describe each structure.

2.1 Participants

We think of participants as consisting of resources R , attributes SV , actions X , and local policies I :

$$a = \langle R_a, SV_a, X_a, I_a \rangle, \quad \langle b = \langle R_b, SV_b, X_b, I_b \rangle. \quad (1)$$

Set $R : R = \{r_i\}$ represents humans, physical resources, services, information, products that a member contributes to VO according to a negotiated agreement.

Set $S : S = \{s_j\}$ represents requested attributes that a participant must prove. Set $SV : SV = \{sv_k\}$ represents *verified* (proved) attributes, such as identity, competence, technical knowledge, skills, that a participant must prove to the other party. In the implementation, these attributes are supported by digital certificates issued by trusted authorities that prove the possession of the attributes.

Set $X : X = \{x_l\}$ denotes actions that must be performed on the resources (e.g., ship product parts before a deadline). The actions can be conditioned on time, location or other context.

Set $I : I = \{i_m\}$ denotes the participant's local policies that express the participant's interests and set of private goals. The policies can, for example, represent the workload acceptance levels (together with their agreed prices); include both a minimum desired production output (under which a partner's participation may not be profitable anymore) and a maximum committed contribution to the VO.

VO participants cooperate to achieve specific objectives/goals not otherwise possible separately by performing a sequence of goal-directed actions (*Plan*). The goal G determines the intended objective of the VO – the purpose of the collaboration, expressed

as a set of rules that specify how the objective can be met. Agreeing on a common goal serves as a pre-condition of trust. Typically, the ultimate goal of VOs is to create value (*utility*) to shareholders through selling an end product F (either goods or services). For the interdependence to exist, the end product must represent a composition of services/resources (R_a, R_b) that belong to each participant (a and b):

$$G : G = \langle Utility(\Gamma), Plan(\Gamma) \rangle. \quad (2)$$

E represents expectations – a wish list comprising of a set of obligations that the other entity should undertake as part of the agreement.

O represents obligations an entity would be willing to enter into under an agreement with another participant. A member commits itself to provide a service under certain terms and conditions to another member. An obligation may represent demanded workload for each participant, resources to be contributed, required prices for each participant's contribution, profit distribution, etc.

Set $W : W = \{w_m\}$ denotes sanctions. Deviation from prescribed behaviours may be admitted and properly addressed through sanctions that are enforced by the agreement enforcement mechanisms. Obligations include at least one sanction in case of non-performance; otherwise, the obligations might be ineffective.

Obligations and expectations are expressed as a set of Boolean formulas over declarative statements making explicit the expected pattern of the participants' interactions.

Each statement describes whether the statement imposes an *action* that must be taken by another member in the federation and a set of sanctions w_i (w_i is in W) that will be carried out in case of non-performance, or whether the statement is a *predicate* that must be met, and exception activity taken if it is not met. The statements are described in terms of:

- resources/services (set R) that are expected from the other party or must be contributed to the party
- attributes (set S) that the participant or other VO members must prove
- specific actions (set X), which a member has to perform on its own, or other members' resources.

The first step is the goal verification process. During this step the participants must agree on a mutually beneficial goal by exchanging the goal verification messages GV . Each member a must verify that the goal G is aligned with the participants' private interests I_a and that incentives are in place to achieve the goal G . For the interdependence to exist, the end product must represent a composition of services/resources (R_a, R_b) that belong to each participant (a and b). To achieve this, the participants must (at least partially) reveal their utilities/capabilities R_a and R_b .

$$GV_{a,b} = \langle G, r_a \rangle \quad (3)$$

$$verify_goal(G, I_a, R_a, R_b) \rightarrow true \text{ or } false. \quad (4)$$

During the agreement negotiation phase, the participants have to agree on a mechanism to share the utility. A participant a enters a negotiation with a participant b with an agreement proposal $P_{a,b}$:

$$P_{a,b} = \langle G, E_{a,b}, O_{a,b} \{sv_k^a\}, \{s_l^b\} \rangle. \quad (5)$$

The agreement proposal consists of the following:

- expectations $E_{a,b}$ expressed as a set of Boolean formulas
- a possibly empty set of obligations $O_{a,b}$ expressed as a set of Boolean formulas
- a possibly empty set of verified attributes $\{sv_k^a\}$
- a possibly empty set of attributes $\{s_l^b\}$ requested by a from b .

We define the rules from the point of view of a . We omitted the similar set of rules defined from the point of view of b for brevity.

Set I represents policies that control the goal validation and agreement negotiation processes. These policies are conditioned on suspicion levels. The local policies I of a participant include the following rule sets:^a

$$fulfill(O_{a,b}, E_{a,b}, O_{b,a}, E_{b,a}) \rightarrow true \text{ or } false. \quad (6)$$

In the simplest case $O_{a,b}$ and $E_{a,b}$ fulfill $O_{b,a}$ and $E_{b,a}$ if they are equal. That is: $O_{a,b} = E_{b,a}$ and $E_{a,b} = O_{b,a}$.

In general, the obligations that fulfill expectations can be defined as a ‘superset’ of the expectations. For example, consider a user expectation “computational resources must be available no <4 h in any 24 h period”. Then any obligation that guarantees resource availability over 4 hours per day fulfils the expectation.

$$counter_of_fer(sl_b, P_{b,a}) \rightarrow P_{a,b}. \quad (7)$$

This rule takes the suspicion level sl_b attributed to b , and the agreement proposal $P_{b,a}$ sent by b as input and returns a counter offer $P_{a,b}$. The offer is created with proposed changes to the expectations and obligations specified in $P_{b,a}$. The changes depend on sl_b . If the suspicion level reaches some threshold, a may chose to end the negotiation based on the lack of trust in b .

$$update_sl(sl_b, H_{a,b}, I_a) \rightarrow sl'_b. \quad (8)$$

This rule takes the suspicion level sl_b attributed to the member b , a set of negotiation sessions between a and b , $H_{a,b}$, and a ’s local policies I_a as input and returns a new suspicion level sl'_b . Even though the participating parties may have no prior knowledge of one another, some characteristics (e.g., geographic location, communications mechanisms, etc.) relevant to the type of the agreement being negotiated contribute to the initial trust in the participant. These characteristics are specified in local policies I and are used for calculating the suspicion level. For example, when negotiating the QoS guarantees with a remote participant behind a slow connection, we might be more suspicious about unrealistically high QoS guarantees proposed by the participant, than when negotiating with a participant behind a high speed connection.

During the negotiation process, each participant observes the flow of offers and counter-offers and according to the rules expressed in local policies I , adjusts the suspicion level for the negotiating party if a suspicious behaviour is detected.

2.2 Suspicion Levels (SLs)

Initial trust is expressed as conformance of participants to normative negotiation behaviours described in the participants' policy I . SLs indicate perceived deviation from the expected behaviour. The values are derived from interaction histories H , allowing the estimation of likely future behaviours. Each suspicion level sl_a is attributed to a particular member a and is represented by a vector:

$$sl_a = \langle sl_a^i \rangle, \quad sl_{\min} \leq sl_a^i \leq sl_{\max}. \quad (9)$$

Each component sl_a^i is attributed to a particular type of suspicion. For example, sl_a^1 may assess likelihood of sensitive information leaks by asking other participants to reveal sensitive attributes that may be unrelated to the nature of agreement (e.g., asking for company's list of clients or internal organisational structure when negotiating a profit share). Investigating the behavioural aspects to be monitored and policies that guide the Suspicion Level adjustment are areas for further research.

sl_{\min} represents strong belief that the participant a is acting in accordance to the normative behaviour, sl_{\max} represents strong belief that the participant a is acting in contrary to the expected or desired behaviour. The SL increases with the occurring times of the suspicious event and decreases when a 'positive' event happens. The value by which the SL is increased (or decreased) depends on the confidence level that the repeated event indicates malicious (or positive) activity.

2.3 Negotiation history

Each participant maintains a history H of interactions with the negotiating party that consists of a set of negotiation rounds Rd^t ordered by the time of occurrence t :

$$H_{a,b} = \{Rd_{a,b}^t\}. \quad (10)$$

Each round Rd^t includes a proposal and a counter proposal received during that round.

$$Rd_{a,b}^t = \{P_{a,b}^t, P_{b,a}^t\} \quad (11)$$

where

$P_{a,b}^t$: proposal sent by member a to member b in the negotiation round $Rd_{a,b}^t$

$P_{b,a}^t$: proposal sent by member b to member a in the negotiation round $Rd_{a,b}^t$.

2.4 Agreements

A represents the established agreement that explicitly specifies legitimate interactions between VO members. Agreements are created during the agreement negotiation process where participants iteratively exchange agreement proposals. The process is guided by the participants' private interests expressed in local policies I .

$A = \langle G, O_{a,b}, O_{b,a} \rangle$ represents an agreement between members a and b . The sanctions are part of the agreement, since they are included within the obligations. The agreement is reached when:

$$\text{fulfill}(O_{a,b}, E_{a,b}, O_{b,a}, E_{b,a}) \rightarrow \text{true} \quad (12)$$

$$\text{fulfill}(O_{b,a}, E_{b,a}, O_{a,b}, E_{a,b}) \rightarrow \text{true}. \quad (13)$$

3 Trust formation example

Consider the example given in the Introduction section that illustrates an agreement negotiation between two enterprises a and b wishing to establish a VO for joint computer assembly. a produces desktops, b manufactures monitors. Each participant defines the minimum profit share it is willing to accept in its local policies I . This information is secret and is not revealed to the other party.

3.1 Goal verification

In this example the goal G :

$$G = \langle \text{Utility}(\Gamma), \text{Plan}(\Gamma) \rangle \quad (14)$$

$$\Gamma = \{r_{\text{desktop}}, r_{\text{monitor}}\}, \quad \text{Plan} = (x_{\text{ship}}, r_{\text{monitor}}) \quad (15)$$

is to get profit from selling a computer (end product) R that is composed of a desktop produced by a and a monitor that must be supplied by b . Assume that the plan to achieve the goal G consists of b shipping the monitor to a . Assume that a initiates a negotiation by sending a message to b that contains an invitation to create a VO to achieve the goal:

$$GV_{a,b} = \langle G, r_{\text{desktop}} \rangle. \quad (16)$$

When b receives this message, it first verifies the goal by making sure that G is aligned with b 's private interests expressed in policies I_b and that the goal is achievable.

If, for example, the end product has to include some software $r = \{r_{\text{desktop}}, r_{\text{monitor}}, r_{\text{software}}\}$, then the goal can not be achieved by the efforts of a and b alone. They would either have to find an additional partner to form a virtual team or change the goal. Assume that the goal verification succeeds:

$$\text{verify_goal}(G, I_b, r_{\text{desktop}}, r_{\text{monitor}}) \rightarrow \text{true} \quad (17)$$

b next sends the goal verification message to a , which performs similar goal verification checks.

$$GV_{a,b} = \langle G, r_{\text{monitor}} \rangle \quad (18)$$

$$\text{verify_goal}(G, I_a, r_{\text{monitor}}, r_{\text{desktop}}) \rightarrow \text{true}. \quad (19)$$

At this point the participants agreed on a mutually beneficial goal and confirmed the interdependence. During the agreement negotiation phase, the participants have to agree on a mechanism to share the profit.

3.2 Agreement negotiation

Round 1

Assume that a initiates a negotiation by sending a message to b that contains the first version of an agreement proposal:

$$P_{a,b}^1 = \langle G, E_{a,b}^1, O_{a,b}^1, \phi, s_b \rangle. \quad (20)$$

The proposal contains:

- a 's expectation $E_{a,b}^1 = (x_{ship}, r_{monitor}, w_{penalty})$ for b to ship a monitor and a sanction – specified penalty $w_{penalty}$ in the case of non-performance
- a 's obligation $O_{a,b}^1 = (x_{share}, r_{profit}, 20\%)$ to give b 20% of the profit
- the quality assurance requirement s^b that b needs to prove to a in order to satisfy the compatibility and quality requirements posed by a on b 's product.

The proposal contains an empty set of a 's verified credentials.

When b receives a 's proposal, it first calculates the suspicion level sl_a attributed to a . Since prior interactions between a and b are limited to just the goal verification process, the sl_a is set to a default value $sl_a^{default}$. The b 's local policies I_b are employed when calculating the suspicion level. For example, the policies may assign a higher suspicion level to foreign partners.

b generates its expectations $E_{b,a}$ and obligations $O_{b,a}$ according to the local policies I_b that indicate, for example, the minimum profit of 60%.

b checks whether the proposed obligations and expectations fulfill its own requirements. Since a proposed lower profit share than b expects, the *fulfill* function returns *false*. Assume that b accepts the requested sanction $w_{penalty}$.

Next b generates a counter-proposal $P_{a,b}^1$. The proposal indicates that:

- b is willing to supply the monitor if a commits to the profit share of 60%
- b agrees to pay a penalty $w_{penalty}$ if b does not perform on time
- a will be a subject to legal sanction $w_{sanction}$ if a does not provide the agreed profit share.

Next b sends the counter-offer and its certification sv^b to a .

$$update_sl(\phi, \phi, I_b) \rightarrow sl_a^{default} \quad (21)$$

$$fulfill(O_{a,b}^1, E_{a,b}^1, O_{b,a}, E_{b,a}) \rightarrow false, \quad (22)$$

where

$$O_{b,a} = (x_{share}, r_{profit}, 60\%, w_{penalty}) \quad (23)$$

$$E_{b,a} = (x_{ship}, r_{monitor}, w_{sanction}) \quad (24)$$

$$counter_of_fer(sl_a^{default}, P_{a,b}^1) \rightarrow P_{b,a}^1 \quad (25)$$

$$P_{b,a}^1 = \langle G, E_{b,a}^1, O_{b,a}^1, sv^b, \phi \rangle, \quad (26)$$

where

$$E_{b,a}^1 = (x_{ship}, r_{monitor}, w_{sanction}) \quad (27)$$

$$O_{b,a}^1 = (x_{share}, r_{profit}, 60\%, w_{penalty}). \quad (28)$$

Rounds 2–N

When a receives b 's counter-offer, a notices that b 's profit expectation is higher than the minimum profit specified in a 's local policies I_a (which indicates the minimum profit of 50%), and makes a counter offer setting b 's profit to a lower value. The negotiation proceeds in a similar fashion until the agreement is reached.

If b is acting suspiciously, for example, constantly exhibiting opportunistic behaviour by not lowering its profit expectation (or even rising the profit expectation each new round), a will increase the suspicion level sl_b . This behaviour is not necessarily malicious, it is natural for one to wish for a higher profit. However, a can infer that b is likely trying to probe the maximum profit share that a can offer, which is a sign of 'lack of good will'. Such behaviour may indicate potential future problems in the collaboration with b .

When sl_b reaches certain threshold (defined in the policies I_a), a may decide not to establish the agreement with b (by sending a negotiation failure message to b) and find a more cooperative partner for the collaboration.

4 Implementation considerations

In this section we briefly describe infrastructures supporting the framework. In previous work, we developed an *Adaptive Trust Negotiation and Access Control (ATNAC)* framework (Ryutov et al., 2005) to address issues of access control in open systems. We are applying these techniques to negotiation of agreement proposals, rather than using them solely to negotiate proofs of security attributes.

The ATNAC framework is based on two well-established systems GAA-API (Ryutov and Neuman, 2002; Ryutov et al., 2003) and TrustBuilder (Winslett et al., 2002). The GAA-API provides adaptive access control that captures dynamically changing system security requirements. The TrustBuilder system regulates when and how sensitive information is disclosed to other parties. The Analyzer maintains a separate SL for each requester based on the IP address and certificate-based identity (if available), and stores the information in a Suspicion Database. Analyzer dynamically calculates the SLs based on the information reported by the GAA-API and TrustBuilder.

This combination extends the capabilities of each system. In particular, the framework allows us to detect and thwart certain attacks on electronic transactions, to adapt information disclosure and resource access policies according to a level of suspicion.

We are extending the ATNAC to support the functions defined in Section 3. In particular, we are using the GAA-API to implement the *fulfill()* function that takes proposed and local expectations and obligations, and returns a decision whether the input fulfils the local requirements.

The TrustBuilder modules serve as a basis for implementing the *attributerelease()* and *counteroffer()* functions that control the sensitive attribute disclosure, building counter offers, and controlling the negotiation process according to the strategies expressed in the local policies. One of the challenges is mapping the local policies *I*, obligations, and expectations to the policy formats supported by the GAA-API (uses EACL format) and TrustBuilder (employs X.509v3 digital certificates and TPL policies).

The Analyzer module implements the *updatesl()* that monitors the interactions between members during the agreement negotiation, updates suspicion levels and detects suspicious behaviours and violations as they occur according to the policies *I*. Proper evidence is collected on both the actions and the lack of actions of the agents by observing the negotiation history *H*.

In the implementation, the SL may consist of several components that are related to different aspects of observed behaviour. For example, a participant repeatedly presents forged credentials or irrelevant credentials that were not requested by the other party, or a participant persistently tries to get the better shared profit by engaging in lengthy negotiations with little or no progress toward a mutually satisfactory outcome. In our current implementation, the SL is comprised of three components:

$$sl = \langle sl_{DoS}, sl_{IL}, sl_O \rangle. \quad (29)$$

sl_{DoS} indicates a probability of DoS attack on behalf of the requester. sl_{IL} is attributed to sensitive information leakage attempts. Finally, sl_O indicates other suspicious behaviours (e.g., misuses of a user's identity or impersonation attempts). All three values range from 0.0 to 1.0.

The Analyzer (local to the participant) increases the SL with the occurring times of the suspicious event and decreases the SL when a 'positive' event happens (e.g., successful agreement negotiation and/or successfully completed business transaction). The value by which the SL is increased depends on the confidence level that the repeated events indicate malicious activity. For example, the Analyzer may increase a particular component of SL by 0.25, 0.25, 0.5 on the first, second and third consecutive errors.

We will extend the SL with additional components (and will design and implement the corresponding *updatesl()* functions) attributed to the 'lack of good will' and other suspicious behaviours observed during the agreement negotiation process.

5 Related work

Zuo and Panda (2005) present a trust management model for maintaining trust levels within a VO. The levels can be dynamically updated to accommodate the dynamic nature of a VO. The model supports integration of partial trust values to evaluate the composite trust of a compound entity.

Cahill et al. (2003) present SECURE – a general framework for trust and risk driven decision making. SECURE project focuses primarily on building trust infrastructures for large ad hoc wireless networks. SECURE utilises the risk and trust to determine whether an interaction can occur.

Quirchmayr et al. (2002) describe an approach to modelling contract establishment in Virtual Enterprises based on the first order predicate logic formalism. VO's main concern is supporting enforceable contracts. This means that it is possible to determine whether the actions of parties are in accordance with the contract in effect. They do not consider notion of trust, negotiation techniques and policies. Damianou et al. (2001) developed annotation and tools for specifying, analysing and enforcing obligation and authorisation policies for managing large scale distributed systems.

PeerTrust (Nejdl et al., 2004) is a trust management system that uses a simple and expressive policy language based on distributed logic programs. PeerTrust agents perform automated trust negotiation to obtain access to sensitive resources. Bonatti and Samarati (2002) proposed a framework based on policy language and an interaction model for regulating access to network services. This trust establishment framework uses logical rules for accessing services and avoiding unnecessary disclosure of sensitive information.

Our work differs from the work cited above in the introduction of the concept of suspicion level and mutual trust effect of the participants behaviour during the trust negotiation process.

6 Conclusions and future work

We presented a framework supporting on-demand creation of trust relationships (or agreements) within the context of cross-institutional VOs. The trust formation process is based on:

- knowledge (third party information – recommendations, knowledge about entity's nature, such as competence, technical capabilities, and skills)
- behavioural aspects: signs of good will (not self-serving or opportunistic), cooperation and information sharing openness
- predictions of future behaviour based on observed interactions.

This framework can be used to implement a server that accepts VO creation proposals for a particular purpose from different organisations. The submission of a proposal initiates the goal verification and negotiation processes. As a result, either a VO agreement is created or the information about potential partners is returned.

Future work includes specification of the exact structure of local policies I , statements and sanctions (that comprise expectations and obligations); extension of the framework to support multi-party negotiations; and investigation of whether the local policies I need to be updated dynamically to accommodate new user requirements and obligations imposed by expectations of the new members.

Acknowledgements

This research was supported by funding the National Science Foundation under grants No. CCR-0325951 and ACI-0325409. The views and conclusions contained herein are those of the authors and should not be interpreted as representing the official policies or

endorsement of the funding agencies. Figures and descriptions were provided by the authors and are used with permission.

References

- Almenarez, F., Marin, A., Campo, C. and Garcia, C. (2004) 'A pervasive trust management model for dynamic open environments', *First Workshop on Pervasive Security, Privacy and Trust (PSPT 2004)*, August 26, Boston, MA, USA.
- Bonatti, P. and Samarati, P. (2002) 'A unified framework for regulating access and information release on the web', *Journal of Computer Security*, Vol. 10, No. 3, pp.241–271.
- Cahill, V., Shand, B., Gray, E., Bryce, C., Dimmock, N., Twigg, A., Bacon, J., English, C., Wagealla, W., Terzis, S., Nicon, P., di Marzo Serugendo, G., Seigneur, J.-M., Carbone, M., Krukow, K., Jensen, C., Chen, Y. and Nielsen, M. (2003) 'Using trust for secure collaboration in uncertain environments', *IEEE Pervasive Computing*, Vol. 2, No. 3, pp.52–61.
- Cardoso, H.L. and Oliveira, E. (2004) 'Virtual enterprise normative framework within electronic institutions', *ESAW'04 – 5th Int. Workshop on Engineering Societies in the Agents World*, Toulouse, Toulouse, France, October 20–22, pp.14–32.
- Damianou, N., Dulay, N., Lupu, E. and Sloman, M. (2001) 'The ponder policy specification language', *Proceedings of Workshop on Policies for Distributed Systems and Networks (POLICY 2001)*, Springer-Verlag, LNCS 1995, Bristol, UK, pp.18–38.
- Davulcu, H., Kifer, M., Pokorny, L.R., Ramakrishnan, C.R., Ramakrishnan, I.V. and Dawson, S. (1999) 'Modeling and analysis of interactions in virtual enterprises', *Proceedings of the Workshop on Research Issues in Data Engineering – Information Technology for Virtual Enterprises (RIDE-VE'99)*, March, Australia, pp.12–18.
- Dellarocas, C., Klein, M. and Rodriguez-Aguilar, J.A. (2002) 'An exception handling architecture for open electronic marketplaces of contract net software agents', *Proceedings of the 2nd ACM Conference on Electronic Commerce*, Minneapolis, Minnesota, USA, pp.225–232.
- Falcone, R. and Castelfranchi, C. (2004) 'Trust dynamics: how trust is influenced by direct experiences and by trust itself', *AAMAS*, Vol. 2, pp.740–747.
- Gray, E., Seigneur, J.-M., Chen, Y. and Jensen, C.D. (2003) 'Trust propagation in small worlds', in Nixon, P. and Terzis, S. (Eds.): *Proceedings of the First International Conference on Trust Management*, Springer-Verlag, April Vol. 2692 of LNCS, pp.239–254.
- Grefen, P., Aberer, K., Hoffner, Y. and Ludwig, H. (2000) 'Crossflow: cross-organizational workflow management in dynamic virtual enterprises', *Computer Systems Science and Engineering*, Vol. 15, No. 5, pp.277–290.
- Ishaya, T. and Mundy, D.P. (2004) 'Trust development and management in virtual communities', *iTrust*, pp.266–276.
- Josang, A. (1996) 'The right type of trust for distributed systems', *New Security Paradigms Workshop*, pp.119–131.
- Josang, A. and Presti, S.L. (2004) 'Analysing the relationship between risk and trust', in Jensen, C., Poslad, S. and Dimitrakos, T. (Eds.): *Proceedings of Second International Conference on Trust Management (iTrust 2004)*, Oxford, UK, Vol. LNCS 2995, pp.135–145.
- McKnight, D.H., Cummings, L.L. and Chervany, N.L. (1998) 'Initial trust formation in new organizational relationships', *Academy of Management Review*, Vol. 23, No. 3, pp.473–490.
- Nejdl, W., Olmedilla, D. and Winslett, M. (2004) 'PeerTrust: automated trust negotiation for peers on the semantic web', *Proceedings of the Workshop on Secure Data Management in a Connected World (SDM'04) in Conjunction with 30th International Conference on Very Large Data Bases*, Toronto, Canada, pp.118–132.
- Periorellis, P. and Parastatidis, S. (2005) 'Task-based access control for virtual organizations', *Lecture Notes in Computer Science*, Vol. 3409, pp.38–47.

- Quirchmayr, G., Milosevic, Z., Tagg, R., Cole, J. and Kulkarni, S. (2002) 'Establishment of virtual enterprise contracts', *Proceedings of DEXA'02 Conference*, September 2–6, France, pp.236–248.
- Rocha, A.P. and Oliveira, E. (2001) 'Electronic institutions as a framework for agents' negotiation and mutual commitment', *Lecture Notes in Computer Science*, Vol. 2258, pp.232–245, ISBN: 3-540-43030-X.
- Ryutov, T. and Neuman, C. (2002) 'The specification and enforcement of advanced security policies', *Proceedings of the Conference on Policies for Distributed Systems and Networks (POLICY 2002)*, June 5–7, Monterey, California, pp.128–138.
- Ryutov, T., Neuman, C., Kim, D. and Zhou, L. (2003) 'Integrated access control and intrusion detection for web servers', *IEEE Transactions on Parallel and Distributed Systems*, Vol. 14, No. 9, pp.841–850.
- Ryutov, T., Zhou, L., Neuman, C., Leithead, T. and Seamons, K. (2005) 'Adaptive trust negotiation and access control', *Proceedings of SACMAT, 10th ACM Symposium on Access Control Models and Technologies*, June, Stockholm, Sweden.
- Tagg, R. and Quirchmayr, G. (2001) 'Towards an interconnection model for evolution of shared workflows in a virtual enterprise', *Proceedings of Third Int. Conference on Information Integration and Web-Based Applications and Services*, Austria.
- Winslett, M., Yu, T., Seamons, K.E., Hess, A., Jacobson, J., Jarvis, R., Smith, B. and Yu, L. (2002) 'Negotiating trust on the web', *IEEE Internet Computing*, Vol. 6, No. 6, pp.30–37.
- Zuo, Y. and Panda, B. (2005) 'Component based trust management in the context of a virtual organization', *ACM Symposium on Applied Computing*, Santa Fe, New Mexico, USA, pp.1582–1588.

Note

^aIn this paper we define the rules from the point of view of *a*. We omitted the similar set of rules defined from the point of view of *b* for brevity.

Bibliography

- Blaze, M., Feigenbaum, J. and Lacy, J. (1996) 'Decentralized trust management', *Proceedings of the IEEE Symposium on Research in Security and Privacy*, Oakland CA, USA, pp.96–117.
- Griffel, F. (1998) 'Electronic contracting with COSMOS – how to establish, negotiate and execute electronic contracts on the internet', *Proceedings of the Second International Enterprise Distributed Object Computation Workshop*, San Diego, CA, USA.
- Macarena-Matos, L.M., Afsarmanesh, H. and Rabelo, R. (2003) 'Infrastructure developments for agile virtual enterprises', *International Journal of Computer Integrated Manufacturing*, ISSN 0951-192X, Vol. 16, Nos. 4–5, pp.235–254.
- Mezzetti, N. (2003) 'Towards a model for trust relationships in virtual enterprises', *14th International Workshop on Database and Expert Systems Applications (DEXA'03)*, Prague, Czech Republic.
- Nielsen, M. and Krukow, K. (2004) 'On the formal modeling of trust in reputation-based systems', *LNCS*, Springer-Verlag, ISSN: 0302-9743, Vol. 3113, pp.192–204.