

A Bayesian Reputation System for Virtual Organizations

Jochen Haller

SAP Research
Vincenz-Priessnitz-Str. 1,
76131 Karlsruhe, Germany
jochen.haller@sap.com

Abstract. Virtual Organizations (VOs) are an emerging business model in today's Internet economy. Increased specialization and focusing on an organization's core competencies requires such novel models to address business opportunities. In a VO, a set of sovereign, geographically dispersed organizations temporarily pool their resources to jointly address a business opportunity. The decision making process determining which potential partners are invited to join the VO is crucial with respect to entire VO's success. The possibility of a VO partner performing badly during the VO's operational phase or announcing bankruptcy endangers the investment taken in integrating their processes and infrastructure for the purpose of the VO. A reputation system can provide additional decision support besides the a priori knowledge from quotations and bidding to avoid events such as VO partner replacement by helping to choose reliable partners in the first place. To achieve this, reputation, an objective trust measure, is optimally aggregated from multiple independent trust sources that inherently characterize an organization's reliability. To allow for the desired predictions of an organization's future performance, a stochastic modeling approach is chosen. The paper will present a taxonomy of TIs for VO environments, a stochastic model to maintain and aggregate trust sources, so called Trust Indicators, and the inclusion of other subjective measures such as feedback.

1 Introduction

In today's business world, commercial relationships become increasingly flexible and are formed on demand whenever a business opportunity emerges. To cater for these changing demands, business relationships integrate Information and Communication Technologies (ICT) to (partially) automate certain decision processes, e.g. the swift discovery and selection of business partners. Virtual Organizations (VOs) are a prominent example of such emerging business models. A VO is defined as a temporary coalition of otherwise independent organizations or individuals, collaborating to achieve a common business goal, one party alone could not master. A VO follows a phased life cycle, first, a VO manager or system integrator observes a business opportunity and discovers potential VO partners (identification phase). Second, the VO manager negotiates with the potential partners, until the required set of VO members is selected, bringing the right set of specialized expertise in the VO. During this phase, the ICT infrastructure

in each partner's domain is set up and configured for collaboration in the VO and also contracts are established (formation phase). Third, the VO executes, each VO member runs his business processes contributing to the overall VO goal - to exploit the business opportunity - until it is reached (operational phase). In case of abnormalities or exceptions like misperforming VO members, adaptations can be made by the VO manager (evolution phase), but the VO stays within the operational phase. Finally, fourth, after having achieved the VO goal, common goods, VO results etc. are dispersed among the VO members according to the contractual agreements. Final processes, for instance billing, are performed (dissolution phase) [1, 2].

VO like structures can already be observed in different business domains. While on the one hand established domains, such as automotive engineering still retain rather static business relationships, a car manufacturer for instance tends to stay with well established and confirmedly reliable car part suppliers, recently founded domains on the other hand exhibit more agile buyer-supplier relationships. An example for the latter is the high-tech industry, such as chip manufacture, where chip prices change on a daily basis. When a buyer selects a supplier, this essentially means taking a trusting decision. If the buyer has no prior knowledge of a particular (set of) supplier(s), there is no absolute certainty about the supplier's future reliable behavior. In that case, an online reputation system can help to minimize the risk of selecting a bad supplier who frequently delivers late or not at all. In contrast to hard security measures that apply during regular operations between a buyer and supplier such as confidential communication channels or access control measures, reputation based decision support belongs to the class of soft security measures [3]. Soft security aims at complementing hard security rather than replacing it. In the context of this publication, we focus on centralized reputation models that cater best for reputation and trust requirements in VOs [4]. The reputation system is assumed to be owned or hosted by an explicitly trusted third party (TTP). Following up on the above statement about the buyer's expectation of a supplier, we define trust as the subjective probability by which the buyer expects a supplier to behave reliably [5]. While trust is a subjective probability, different buyers have different trust perceptions or expectations of the same buyer, reputation strives at providing an objective trust measure. It is defined as the business context specific aggregation of (subjective) trust values from multiple independent sources to support a supplier's decision making process with respect to an intended collaboration with prospective buyers [6].

In the following, this short paper will outline a snapshot of work currently conducted on a reputation system model that caters for the following, not necessarily only VO specific, requirements:

1. integrate a stochastic trust management model that takes the specifics of business relationships into account
2. in particular, tie trust to observable parameters (in the following termed as Trust Indicators, short TI), inherently characterizing the abilities of a business partner
3. deal with uncertain, incomplete business partner information in dynamic VO environments
4. take long running VOs into account that require all involved parties to query for reputation, not only a buyer investigating about suppliers

- integrate direct feedback to increase the quality of future reputation responses

Section 2 will introduce the reputation model, that integrates a stochastic trust management approach where trust is tied to a taxonomy of so-called trust indicators characterizing business partners. A feedback mechanism will also be outlined. Section 3 concludes.

2 Trust and Reputation Model

Many existing trust management approaches, especially in Peer-2-Peer (P2P) environments, root trust solely in feedback given by peers after having conducted a transaction [7, 8]. Feedback is a highly subjective, relative input that may vary between peers who participated in the same transaction. In contrast to these approaches, we believe that rooting trust in absolute, observable properties - the TIs - inherently characterizing an organization's reliability is a sounder approach for trust among VO members. Feedback is still an important input that improves the quality and accuracy of the reputation value aggregating the TI rooted trust values.

2.1 Taxonomy of TIs

This subsection provides an excerpt of an assembled classification/taxonomy of TIs and their individual modeling. The full taxonomy with more extensive examples will be published in a full paper soon. A TI models one aspect of trust in an organization, participating in VO. Since VOs can emerge in different business domains, e.g. high-tech or engineering industries, having a large set of members with specialized expertise, TIs can have heterogeneous origins and meaning. Therefore, it makes sense to classify TIs that share a similar origin and meaning. This approach has the benefit that the relevance of a TI class for a particular business domain and even individual VO members requesting reputation values can be easily determined.



Fig. 1. TI Taxonomy

Figure 1 shows the top-level TI classes of a taxonomy currently encompassing 146 different TIs. TIs are classified into:

- Financial TIs modeling the financial trust aspects of an organization, an example is the cash flow quote indicator [9].
- Organizational TIs modeling the organizational reliability and stability of an organization, e.g. with the employee fluctuation indicator [10].
- Operational TIs model an organization's operational reliability, for instance with the delivery delay indicator [11, 12].

- External TIs cater for trust relevant aspects external to an organization such as with the country bond spread index indicator that aggregates country risk [13].
- Third Party TIs allow for interfacing with other specialized, third party reputation or expert systems providing trust relevant indicators in form of recommendation.

It already becomes obvious from the TI top-level class descriptions that the taxonomy exercise is an interdisciplinary effort drawing, among others, from the fields of risk management, operations research and Key Performance Measurement. A related approach was conducted by Tan [14] who assembled a Trust Matrix entailing trust aspects for Electronic Commerce but who remained with his work on a higher and abstract level.

After the top-down description of the high level TI classes, we follow with the detailed modeling of individual TIs. Re-iterating, a TI is an observable property characterizing a trust aspect of an organization, therefore with an impact on its reputation. Taking the operational TI "delivery delay" as an example, the time difference passing between an agreed upon delivery date by buyer and supplier and the actual delivery date is observed. Most suppliers aim at minimizing delay, hence it can be expected that suppliers more often deliver slightly late with the number of suppliers decreasing with increased delivery delay. The goal of TI modeling is finally to predict future VO member behavior based on previously collected (TI) data. The fact that the availability of correctly observed data can not be guaranteed in productive systems motivates a TI model based on probability distributions. The described behavior of the "delivery delay" TI for instance suggests a model based on the Exponential distribution [6]. To cater for missing or incorrect data, observed data does not directly determine the TI distribution. Instead, a Bayes update, by applying the following Bayes theorem equation is used:

$$P(\theta|X) = \frac{P(X|\theta)P(\theta)}{P(X)} = \frac{P(X|\theta)P(\theta)}{\sum_S P(X|\theta)P(\theta)}$$

Observed data X contributes to the prior or empirical distribution $P(X)$. The distribution assumption, in case of delivery delay that it is exponentially distributed, determines the likelihood distribution $P(X|\theta)$ estimating the true parameter θ . Evaluating the equation leads to the posterior distribution $P(\theta|X)$ with the best fit of the distribution parameterized by the observed data to the distribution assumption. The posterior distribution is further employed for aggregating towards a reputation value. Assuming discretionary density functions, allowing for more efficient numerical evaluation, the right side of the equation is evaluated. The normalizing denominator's sum index S iterates over the equidistant intervals or states of the grid.

Figure 2 visualizes the described example distributions graphically, based on data sampled from a Enterprise Resource Planning system.

Besides these mathematical properties, the TI model also entails the following attributes:

Name N . Every TI is uniquely identified by a name N .

Domain D . A TI can be based on observations of a continuous or discrete variable x . The possible values of x are the domain of the TI.




Fig. 2. TI Delivery Delay

Update time period Δt_{upd} . Trust information is likely to arrive in different intervals.

The attribute Δt_{upd} defines a fixed time grid telling the reputation system, how often to update a TI.

Observation time period Δt_{obs} . The time period Δt_{obs} defines a maximal time window to look into the past. Beyond that, observations are regarded to carry no more significance.

Time weighting function ω . Among n observations x_i at times t_i , $i \in \{1, \dots, n\}$ within the time window, old ones are less likely to carry information about future values than newer ones. Each TI incorporates a monotonically increasing weighting function $\omega(t) > 0$ that can implement forgetting of older observations and put emphasis on newer ones.

Trust preference mapping π . In order to judge the level of trustworthiness displayed by a TI, we define an ordinal scale 1 to p_{max} , where 1 represents the lowest and p_{max} the highest level of trust indicated by the TI. To compare TIs, the scale is the same for all TIs. $p_{max} = const$. π defines a function $\pi: S \rightarrow \{1, \dots, p_{max}\}$ mapping the states S to the different levels of trust indicated by them. This mapping enables an expert to incorporate his knowledge on the particular TI domain.

It has to be noted that this information rich TI model can express complex properties of business transactions and relationships. Other reputation systems in related work already attempted stochastic models for business transactions or returned feedback, but these attempts could only model binary events or transactions, that is if a transaction ended positive or negative. A prominent example, the Beta Reputation System in [15] uses the Beta distribution that is parametrized by the amount of positive and negative previous outcomes of a particular transaction. Real business transactions involving delivery of chip components such as heatsinks, do not behave that atomically and may be long-running, requiring a more sophisticated model of transaction indicators and their attributes.

2.2 Stochastic Model

The set of TI instances characterizes one VO member organization. The TI's posterior distributions periodically obtained for each TI after Δt_{obs} must then be aggregated to a reputation value for this organization. To achieve this goal and still retain the predictive properties of the probabilistic approach, a Bayes Network (BN) is employed. A BN can be visualized as a directed, acyclic graph.




Fig. 3. Network Topology

Figure 3 shows the BN topology, a tree with three layers, used for the presented reputation system. The edges denote causal dependencies of the nodes and each node in the BN holds a Conditional Probability Table (CPT) with the probability values for its random variable(s) depending on the parent node's random variables. The bottom layer entails the TI information nodes periodically inputting newly observed data into the BN with the updated TI posterior distribution. Since each TI is modeled according to its own properties, TIs may have different domains, states etc., they are not directly comparable. Therefore, the middle layer preference nodes transform the probabilities corresponding to the each TI's states to a normalized scale that can be configured for each reputation system instance, e.g. a scale of $p_{max} = 5$ values (0 = very bad, 1 = bad, 2 = medium, 3 = good, 4 = very good). The same states are held by each middle layer node. The state ranges for each preference node state can be initially set from a pre-defined template or are a domain expert's input at system configuration time. This also applies for initially bootstrapping the BN. If no historic TI data is available, the TI posterior distributions may be bootstrapped from a template configuration set that the mediator nodes then yield a desired probability, e.g. 2 = medium in above scale. If historic TI data is available, e.g. from a reputation system instance running in the same or a similar business domain, this TI data can be transferred, anonymized if required, to the new system instance.

In this approach, the BN performs information learning with each TI update. The reputation distribution can then be obtained from the top layer root node, a conditional probability distribution depending on the preference nodes on the middle layer. For a

human being, the VO Manager representative, requesting a reputation value for a potential VO member, it makes sense to return the distribution's expectation value as the reputation value and variance to indicate the uncertainty inherent to the reputation value. The information conveyed by an entire probability distribution typically overloads a human being while a computer system further using the reputation distribution for control decisions may benefit from the additional level of detail.

2.3 Feedback

As mentioned in the beginning, the reputation system relying on observable TIs may still benefit from feedback given by participants of a business transaction who received decision support from the reputation system. It has to ensure, e.g. with an established protected session reference that this feedback uniquely relates to the correct transaction and stems from authenticated participants. After several experiments with a prototypical implementation of the presented reputation system, that provides reputation values to human requestors, a feedback mechanism accepting feedback in the interval $[0, 1]$ with an associated certainty value of the same interval was contrived, that can be interpreted as a density function as well. This density is then introduced to the BN as another probability condition to the root node. The feedback hereby applies directly to the overall reputation which more closely aligns to human feedback strategies who typically do not break down feedback to the level of individual TIs but rather follow a gut-feeling approach rating the overall business transaction.

3 Conclusion and Future Work

In conclusion, this short paper presented work in progress, a reputation system integrating a stochastic trust management approach. Revisiting the initially stated requirements, with the rich TI model, the system can take specifics of even complex and long running business transactions into account (1., 4.). Trust is hereby rooted in absolute, observable TIs that characterize an organization and its behavior in a business domain (2.). TIs are periodically updated, delivering fresh input for the reputation system. The TI model itself takes care of near dynamic, real-time data updates by modeling the optimal update frequency. The Bayes update performed when fresh data arrives smoothes out incorrect data, the posterior distribution represents the best fit to the modeled distribution assumption (3.). The reputation system is available for all VO members, takes ongoing collaboration into account and a feedback mechanism that influences the reputation distribution directly was presented (5.).

Currently ongoing and future work deals with a set of sound, business domain dependent, configuration sets for bootstrapping a reputation system instance without historic data. Such a set takes business domain specifics into account and aims at avoiding attacks on the system such as collusion and discrimination attacks or badmouthing. The bayesian trust management model itself will be improved by modeling dependencies among TIs themselves on the bottom BN layer. This may potentially lead to cycles in the BN that violates the claim that a directed BN has to be acyclic. Several approaches

from BN and graph theory are currently explored, the most promising so far is the junction tree approach, transforming a BN's graph with cycles into a acyclic junction tree [16, 17]. The downside is the high cost of such a transformation.

References

1. Strader, T., Lin, F., Shaw, M.: Information structure for electronic virtual organization management. *Decision Support Systems* (1998) 75–94
2. Wilson, M., Arenas, A., Chadwick, D., Dimitrakos, T., Doser, J., Giambiagi, P., Golby, D., Geuer-Pollman, C., Haller, J., Ketil, S., Mahler, T., Martino, L., Parent, X., Ristol, S., Sairamesh, J., Schubert, L., Tuptuk, N.: The trustcom framework v0.5. In: *TrustCoM workshop at the 6th IFIP Working Conference on VIRTUAL ENTERPRISES (PRO-VE'05)*. (2005)
3. Rasmusson, L., Janssen, S.: Simulated social control for secure internet commerce. In Meadows, C., ed.: *Proceedings of the 1996 New Security Paradigms Workshop, ACM* (1996)
4. Robinson, P., Karabulut, Y., Haller, J.: Dynamic virtual organization management for service oriented enterprise applications. In: *The First International Conference on Collaborative Computing: Networking, Applications and Worksharing (CollaborateCom 2005)*. (2005)
5. Gambetta, D. In: *Can We Trust Trust?* Basil Blackwell (1988) 213–237 Reprinted in electronic edition from Department of Sociology, University of Oxford, chapter 13, pp. 213-237.
6. Haller, J.: A stochastic approach for trust management. In: *International Workshop on Security and Trust in Decentralized/Distributed Data Structures (STD3S)*. (2006)
7. Regan, K., Cohen, R., Poupart, P.: The Advisor-POMDP: A principled approach to trust through reputation in electronic markets. In: *Proceedings of Privacy, Security and Trust (PST05)*. (2005)
8. Zacharia, G., Moukas, A., Maes, P.: Collaborative reputation mechanisms in electronic marketplaces. In: *Proceedings of the Thirty-second Annual Hawaii International Conference on System Sciences (HICSS-32)*. (1999)
9. Allen, S.: *Financial risk management: a practitioner's guide to managing market and credit risk*. Wiley, Chichester [u.a.] (2003)
10. Teitelbaum, D., Axtell, R.: *Firm Size Dynamics of Industries: Stochastic Growth Processes, Large Fluctuations, and the Population of Firms*. (2005)
11. King, J.L.: *Operational risk: measurement and modelling*. Wiley, Chichester [u.a.] (2001)
12. Cruz, M.G.: *Modeling, measuring and hedging operational risk*. Wiley, Chichester [u.a.] (2003)
13. Romeike, F.: *Modernes Risikomanagement: die Markt-, Kredit- und operationellen Risiken zukunftsorientiert steuern*. Wiley-VCH, Weinheim (2005)
14. Tan, Y.H.: A trust matrix model for electronic commerce. In: *Trust Management, First International Conference, iTrust 2003*. Volume 2692., Heraklion, Crete, Greece, Springer (2002) 33–45
15. Ismail, R., Josang, A.: The beta reputation system. In: *Proceedings of the 15th Bled Conference on Electronic Commerce*. (2002)
16. Murphy, K.P.: (The bayes net toolbox for MATLAB)
17. Heckerman, D., Chickering, D.M., Meek, C., Rounthwaite, R., Kadie, C.M.: Dependency networks for inference, collaborative filtering, and data visualization. *Journal of Machine Learning Research* **1** (2000) 49–75