The Workshops of the Thirtieth AAAI Conference on Artificial Intelligence
Incentives and Trust in Electronic Communities: Technical Report WS-16-09

Modeling Trust Evaluating Agents: Towards a Comprehensive Trust Management for Multi-Agent Systems

Abdullah Aref and Thomas Tran
School of Electrical Engineering and Computer Science
Faculty of Engineering, University of Ottawa
800 King Edward Ave. Ottawa, Ontario, K1N 6N5, Canada

Abstract

In multiagent systems, if interactions are based on trust, trustworthy trustees will have a greater impact on the results of interactions. Consequently, building a high trust may be an advantage for rational trustees. This work describes a trust establishment model that goes beyond trust evaluation to outline actions to direct trustees (instead of trusters). The model uses the number of transactions performed by trusters. A trustee will adjust its performance, depending on the average number of transactions carried out by that truster, relative to the mean number of transactions performed by all trusters interacting with this trustee. The proposed model does not depend on direct feedback, nor does it rely on current reputation of trustees in the community. Simulation results indicate that trustees empowered with the proposed model can be selected more by trusters.

Introduction

Trust management is considered a fundamental research interest in multi-agent systems, it includes trust establishment, engagement, evaluation and use (Sen 2013). If agents' interactions are based on trust, trustworthy trustees (TEs) will have a greater impact on interactions results. Trust often has to be acquired at a cost. Such cost may be compensated if improved trustworthiness leads to further profitable interactions. In such situations, building a high trust may be an advantage for rational TEs. Rational reasoning including trust factors will trade off the cost of building and keeping trust in the community with the future anticipated gains from holding the trust acquired. Though previous research has suggested and evaluated several trust and reputation techniques that evaluate the trustworthiness of TEs, slight consideration has been paid to trust establishment (Sen 2013), which is the principal objective of this work.

It is argued that reputation mechanisms in multi-agent systems can be used not only to enable trusters (TRs) to make better trust evaluations, but also to provide an incentive for good behavior among TEs (Burnett, Norman, and Sycara 2011; Castelfranchi, Falcone, and Marzo 2006). Reputation as a capital that benefit TEs is discussed in (Castelfranchi, Falcone, and Marzo 2006) and argued that TEs that have become trusted have better chances of being chosen as interactions partners and can raise the minimum 'price' they can obtain for their transactions. The use of trust-gain as an incentive mechanism for honesty in e-marketplace environments as in (Zhang and Cohen 2007) and (Burnett, Norman, and Sycara 2011) can be considered a starting point for the novel direction of developing models for TEs to establish trust.

In this work, we use the terminology from (Sen 2013), where the author distinguished between trust establishment, engagement, and evaluation. What Sen (Sen 2013) called engagement is also referred to in the literature as bootstrapping and cold start. In the literature, unfortunately, the term " trust establishment" is used to refer to each of the terms establishment, engagement and evaluation. Many researchers in the domain of ad-hoc networks use the term trust "establishment" to refer to what Sen defined as "evaluation" such as (Saini and Gautam 2011). Many researchers in the domain of service-oriented computing, such as (Malik and Bouguettaya 2009), use the term "establishment" to refer to what Sen defined as "engagement". With this in mind, few work in the literature addresses trust "establishment" as defined by Sen and used in our work. Furthermore, existing trust establishment models for MASs (Aref and Tran 2015b), (Tran, Cohen, and Langlois 2014) allow TEs to adjust their behavior based on direct feedback from TRs. Unlike existing work, TEIF is proposed for situations where direct feedback is not available or not preferred. The proposed model in this work uses the retention of TRs to model TRs behaviors to help TEs engendering the trust of TRs within the environment.

Related Work

It the argued that most research to date from the multi-agent systems trust modeling community has focused on algorithms for TRs to model the trustworthiness of TEs, in order to decide about which TE(s) to select (Aref and Tran 2015b). On the other hand, with few exceptions, trust modeling community ignored engagement and establishment decisions (Sen 2013). Recent surveys such as (Yu et al. 2013; Pinyol and Sabater-Mir 2013) provide more insight on existing work in the field of MAS trust modeling from the perspective of TRs. Modeling trust from TRs' view is not directly related to this work.

Bunett et al. (Burnett, Norman, and Sycara 2011) described a reputational incentive model based on the notion that untrustworthy TEs will have less chance of being selected, and so must work for less rewards than trustworthy

TEs in order to remain competitive. However, in doing this, those TEs obtain less reward(s) from interactions, while having to expend the same efforts performing the task. The potential losses or gains associated with reputational changes can be used as incentive for TE(s) to select a particular Performance level when interacting with TR(s).

Tran et al. (Tran, Cohen, and Langlois 2014) outlined a framework for allowing agents in multi-agent systems to engender trust from the other agents in the environment in the context of e-marketplaces, where TEs attempt to classify TRs into three non-overlapping categories based on product price and quality. Quality and price were used as the only dimensions, and changing one of them will affect the other. By the model, TRs are classified as price-sensitive TRs who are more interested in a low price than a high quality, quality-sensitive buyers who are interested in a high quality more than in a low price, and balanced buyers who consider price and quality equally important. TEs modify their behavior when predictions prove to be inaccurate. The primary aim of TEs is to acquire high levels of trust, of value for future interactions with TRs.

Aref and Tran (Aref and Tran 2015b) described a multi-criteria trust establishment model inspired by the research on customer satisfaction in the field of marketing. TEs attempt to improve their performance based on the importance of the satisfaction criterion and their level of demand, also referred to as the lack of satisfaction in (Aref and Tran 2015b). According to the model, it is important that TEs enhance their Performance for highly significant and extremely demanded features. On the other hand, TEs can decide to reduce their Performance for unimportant, and un-demanded feature(s) while the corresponding TR(s) may still be satisfied. The model considers the remaining features are either un-demanded and important or demanded and unimportant. Therefore, TEs can choose to make little if any, enhancement for those features. In addition to assuming the availability of direct feedback, the work did not describe the price that TEs have to pay for satisfying TRs.

Framework

Agent architecture

Based on the agent's architecture described in (Sen 2013), we assume that each agent has an embedded trust management module. This module stores models of other agents and interfaces both with the communication module and the decision selection mechanism. The subcomponents of the trust management module, in compliance with (Sen 2013), are:

- Evaluation: This component is responsible for evaluating the trustworthiness of other agents using different information sources, such as direct experience and witness testimonies such as Regret (Sabater and Sierra 2001) and DTMAS (Aref and Tran 2015a).
- Establishment: This component is responsible for determining the proper actions to establish the agent to be trustworthy to others, such as the work of Tran et al. (Tran, Cohen, and Langlois 2014).
- Engagement: This component is responsible for allowing rational agents to decide to interact and engage others with the aim of estimating their trustworthiness. In the literature, this ingredient is usually referred to as trust bootstrapping and cold start problem. Bootstrapping Trust Evaluations Through Stereotypes (Burnett, Norman, and Sycara 2010) and the work in (Malik and Bouguettaya 2009) are models that belong mainly to this component.
- Use: This component is responsible for determining how to select prospective sequences of actions meant on the trust models of other agents that have been learned. The trust decision-making model described in (Burnett, Norman, and Sycara 2011) is a model that belongs mainly to this component.

Agents and tasks

We assume a society of agents, A = $\{a_1, a_2, ...\}$, a set of possible tasks S =$\{ s_1, ..., s_k\}$, a set of TRs X =$\{ x_1, ..., x_n\}$ and a set of TEs Y =$\{ y_1, ..., y_m\}$ such that $X \cup Y \subseteq A$. The nature of tasks in S is application dependent. A TR $x \in X$ that desires to see some task accomplished, considers depending on a TE $y \in Y$ to perform the task on its behalf (Burnett, Norman, and Sycara 2011). Any TR x can request to collaborate with any other TE $y \in Y$ in order to achieve a task $s \in S$. TR x may request any task s zero or more times.

In response to the request t made by TR x to do task s, TE y proposes to deliver a utility gain for x by a transaction. TR x then gains some benefits from the interaction. This benefit is referred to as the perceived Utility Gain (UG^x) (Lerman and Galstyan 2001), which may or may not be the same as $ProposedUG^y_x(s,t)$. The outcomes of each task can be either satisfactory or dis-satisfactory (Burnett, Norman, and Sycara 2011).

Trust Evaluations

As our focus in this work is on trust establishment, we do not discuss how trust evaluations are formed or how they are used. Instead, we assume the existence of a trust evaluation model and a decision-making model. TR x models the trustworthiness of all TEs in the society using function

$$trust^x : Y \rightarrow (minTrust, maxTrust) \tag{1}$$

which is called the trust function of x, minTrust and maxTrust are minimum and maximum trust values respectively.

TEIF: A Trust Establishment Model Using Indirect Feedback

Overview

It is possible that TRs are not willing to provide direct feedback information to TEs, for many reasons, such as unjustified cost or being self-centered. It is argued in (Sen et al. 1994) that agents can act independently and autonomously when they do share problem-solving knowledge. Such agents are not affected by communication delays or misbehaviors of others.

In this work, we present TEIF; a trust establishment model based on indirect feedback from TRs where TEs use the retention of TRs to model TRs' behavior. Using TEIF, a TE models the behavior of a particular TR as follows:

- If the retention rate of the TR is less than the average retention rate by all TRs, this indicates that the TR is not so happy with the UG provided by the TE. In response, by TEIF, the TE should put some effort to encourage the TR to interact with itself later. For this purpose, the TE categorizes TRs into three non-overlapping sets, and responds with different enhancement to TRs in various categories. When the retention of the TR is way above the average, then the TE may carefully attempt to increase its profit. However, if the retention rate is around the average, then the TE retains the level of Performance , generally, unchanged.
- Furthermore, the TE attempts to attract a TR for the first (few) time(s) by proposing a relatively high UG in response to a bid request.

TEIF Components

Private Retention Index (PRI) In the absence of direct feedback from TRs, a TE could use the percent of transactions performed by a particular TR, relative to the highest number of transactions completed by any individual TR so far. However, this can be tiny and may become misleading. Alternatively, the TE may depend on what we call the Retention Indicator R_i to model the retention of TR$_i$. R_i is the average retention of a particular TR relative to the average retention per TR among those TRs having interacted with the TE previously. To simplify using this indicator in calculations, we normalize it by the largest known value of R to the TE, and we refer to the result as Private Retention Index (PRI)

Assuming that rational TRs do not frequently interact with untrusted TEs, this index can be used as an indicator of TR's satisfaction. This number cannot be greater than one or less than zero.

$$R_i = \frac{NT_i}{\frac{\sum_{j=1}^{N_{TR}} NT_j}{N_{TR}}} \tag{2}$$

$$PRI_i = \frac{R_i}{max_{TR}(R)} \tag{3}$$

- NT_i is the number of transactions performed by TR$_i$ with the TE.
- N_{TR} is the number of TRs having interacted so far with the TE.
- $max_{TR}(R)$ is the maximum Retention Indicator of individual TR interacted with the TE so far.

The PRI is meant to indicate the willingness of a TR to interact with the TE. We agree that the PRI says nothing about the relation of the TR with other TEs, i.e. competitors, as well as other TRs in the community. However, such information might not always be available due to selfishness, privacy, or lack of authorized providers, among other possible reasons.

Using TEIF, the TE can categorize TRs into four non-overlapping sets based on the value of the retention index. Those sets are:

- The set of engaged TRs: Is the set of frequently returning TRs, characterized by having PRI greater than or equal to Engagement Threshold (ET). The TE believes it is well trusted among TRs in this set. Therefore, the TE carefully attempts to make profits, in other words, reduce the provided UG. By "carefully" we mean small, gradual changes.
- The set of discoverers: Is the set of low frequently returning TRs characterized by having PRI less than Discovering Threshold (DT). The TE believes that TRs in this set have no preference of it over other, and such TRs are not serious in interacting with itself. Therefore, the TE does not put much effort in attempting to attract them.
- The set of users: Those with PRI greater than or equal to Using Threshold (UT) but less than ET. The TE believes that TRs in this set have average willingness to interact with it, but they have not decided to be loyal to it. Such TRs may easily lose interest in future interaction with the TE. Therefore, the TE puts relatively high effort attempting to attract TRs in this set.
- The set of choosers: Those with PRI greater than or equal to Discovering Threshold (DT) but less than UT. The TE believes that TRs in this set have not decided yet to prefer it over others, but they are potential future interactors. Moreover, there is a risk of those TRs losing interest in future interaction with the TE. Therefore, the TE put neutral effort in attempting to attract TRs in this set.

To address different sets, we introduce a Category Scaling Factor $CSF_k > 0$ corresponding to each category, where $k \in \{engaged, discoverers, choosers, users\}$. It is clear that the value of those CSF_k are application dependent. As a general guidelines, we propose that

- $CSF_{users} > 1$, $CSF_{users} \geq CSF_{choosers}$
- $CSF_{choosers} \approx 1$
- $CSF_{discoverers} < 1$, $CSF_{discoverers} \leq CSF_{choosers}$
- $CSF_{engaged} < 1$, $CSF_{engaged} \leq CSF_{choosers}$

General Retention Index (GRI) To indicate the general trend of retention among TRs, we define GRI as:

$$GRI_t = \overline{Ret_t} - \overline{Ret_{t-1}} \tag{4}$$

- GRI_t is the General Retention Index (GRI) at time instance t
- $\overline{Ret_t}$ is the average retention rate calculated at time instance t.

$$\overline{Ret} = \frac{\sum_{j=1}^{N_{TR}} NT_j}{N_{TR}} \tag{5}$$

- $\overline{Ret_{t-1}}$ is the average retention rate calculated at time instance $t - 1$

A reduction in the average retention rate, i.e. $GRI_t < 0$ can indicate a general trend of dissatisfaction among TRs, rather than individual behavior of a particular TR. In this case, the TE should scale up the provided UG for TRs. On the other hand, an increase in the average retention rate, i.e. $GRI_t > 0$, can indicate a general trend of satisfaction among TRs. In this case, the TE expects that a low retention rate of a TR is an individual behavior rather than a general dissatisfaction. As a result, the TE can still enhance the provided UG for that TR. However, the TE can scale down the provided UG. To address different cases, we introduce the General Scaling Factor GSF. The value of the scaling factor should not go below zero. In addition to its role in attracting TRs by promising higher UGs for low retention rate TRs, the use of the scaling factor may help TEs control the offered UG in order to enhance profit.

Each TE calculates the average retention rate per known TRs, i.e., those previously interacted with the TE, after each transaction or at the end of a time interval. If the newly calculated average is larger than the last calculated one, the TE can expect a general trend of satisfaction among TRs'. Therefore, it is a chance for the TE to make profit and the value of the scaling factor GSF will be set to High General Retention (HGR), a positive number less than 1. This, in effect, will scale down the offered UG. On the other hand, if the newly calculated average is less than the last calculated one, the TE can expect a general trend of dissatisfaction among TRs'. Therefore, significant enhancement may be necessary and the value of the scaling factor GSF will be set a value larger than HGR, we will refer to this value as Low General Retention (LGR), a positive number larger than 1 and $LGR > HGR$. This in effect will affect the speed of enhancing offered UG and the gained profit. In the absence of indicators of activity in the community, when the newly calculated average is close to the one calculated previously, the TE can expect a temporary stability in the TRs' behavior only minor enhancement can be used, if any, rather than a significant enhancement. Therefore, we recommend that the TE set the value of the GSF to a value close to 1, larger than HGR, and less than or equal LGR. We will refer to this value as the Neutral General Retention (NGR).

We use a general scaling factor GSF_l corresponding to each category, where $l \in \{HGR, NGR, LGR\}$. It is clear that the value of those GSF_k are application dependent. As a general guideline, we propose that

- $GSF_{LGR} \approx 3 * GSF_{HGR}$
- $GSF_{NGR} \approx 2 * GSF_{HGR}$
- $GSF_{HGR} \approx 0.5$

The Improvement Index The improvement efforts necessary for a TR depends on the category of the TR, the PRI_i and the general scaling factor as well. PRI_i indicates the frequency of retention of TR_i. TRs of the same category k, where $k \in \{engaged, discoverers, choosers, users\}$, can have different levels of retentions, and consequently, require other levels of improvements. The improvement index can show the improvement margins for a specific TR, and it is assessed according to the following equation:

$$I_i = GSF_l * CSF_k * (1 - PRI_i) \quad (6)$$

- CSF_k where $k \in \{engaged, discoverers, choosers, users\}$.
- GSF_l , where $l \in \{HGR, NGR, LGR\}$.
- If I_i calculated by eq. 4 is greater than 1, I_i will be reset to 1.

$ProposedUG_x^y(s,t)$ When a TE propose to deliver a value for TR by a transaction, the proposed value is scaled by the improvement index of the particular TR; that is

$$ProposedUG_x^y(s,t) = I_i * MaxImprovment + MinUG \quad (7)$$

- $MinUG$ is the minimum possible UG that the TE may deliver.
- $MaxUG$ is the maximum possible UG that the TE may deliver.
- $MaxImprovement$ is $MaxUG - MinUG$

Performance Analysis

It is often challenging to obtain suitable real world data sets for the comprehensive evaluation of trust models since the effectiveness of various trust models needs to be assessed under different environmental conditions (Yu et al. 2013). Therefore, in trust modeling for MASs research field, most of the existing trust evaluation models are evaluated using simulation or synthetic data (Yu et al. 2013). The situation is even harder for the case of trust establishment, as we are not aware of any simulation testbed that address trust establishment as used in this work to be used for the evaluation of TEIF.

Performance Measures

We agree with (Sen 2013) that trust often has to be acquired at a cost, and such cost may be compensated if improved trustworthiness leads to further profitable interactions. In such situations, building a high trust may be an advantage for rational TEs. By rational reasoning, including trust factors will trade off the cost of building and keeping trust in the community with the future envisioned gains from holding the trust acquired. Therefore, to study the Performance of TEIF, we use the following measures:

- Percent of overall transactions: A primary objective of a TE planning to enhance its trustworthiness estimation value, as seen by TRs, is to become selected by TRs for future transactions. The larger the number of transactions the TE performs with TRs, the closer it is in achieving this objective. However, the absolute number may be misleading as a TE may achieve 9 transactions out of 10 possible transactions is more successful than the one who achieves 9 out of 100. Therefore, we believe that it is preferable to use the percent of transactions performed with the TE relative to the overall number of transactions performed by all TEs in the system as a measure of achieving this objective. In this work, the percent of overall transactions measure is defined as the percent of transactions performed with TEs equipped with TEIF, out of all transactions that took place in the community. If TRs depend on a trust evaluation model to select TEs to interact with, then the higher percent of interactions will occur with the more trusted TEs. This measure indicates the benefits that TEs may achieve by adopting TEIF.

- Average of the percent of delivered UG: It is arguable that an honest TE can achieve a higher percent of overall transactions if it is committed to delivering the highest possible UG to TRs. However providing greater UG usually incurs extra cost. Therefore, a rational TE will attempt to achieve the greatest percent of overall transactions while keeping the provided UG as low as possible. We believe that the percent of the delivered UG relative to the highest possible value is more appropriate compared to the absolute UG value. In this work, the average of the percent of delivered UG per TE is defined as the summation of all delivered values percentages in all transactions divided by the number of transactions involving that TE. The overall average is the average of values obtained by individual TEs. This measure indicates the efforts needed to achieve the enhancement in the percent of overall transactions.

Simulation Environment

We use simulation to examine the Performance of the proposed model for a distributed MAS environment using the discrete-event MAS simulation toolkit MASON (Luke et al. 2005) with TEs providing services, and TRs consuming services. We assume that the Performance of individual TE in a particular service is independent of that in another service. Therefore, without loss of generality, and in order to reduce the complexity of the simulation environment, it is assumed that there is only one type of services in the simulated system. All TEs offer the same service with, possibly, different Performance s. In order to study the Performance of TEIF, we compare the proposed model with the reputational incentive model of (Burnett, Norman, and Sycara 2011). Network communication effects are not considered in this simulation. Each agent can reach each other agent. The simulation step is used as the time value for interactions. Transactions that take place in the same simulation step are considered simultaneous. Locating TEs and other agents are not part of the proposed model, and agents locate each other through the system. TRs request all or part of the TEs to bid. TEIF-empowered TEs, calculate the PRIs of the requesting TRs, determine the proper category for TRs and select the corresponding value of CSF. TRs then select partner TEs to interact with based on the expected UG to be gained from the transaction. Such value is calculated as

$$EV = trust * BV \tag{8}$$

- EV is the expected value.
- BV is the bid value declared by the TE.
- $trust$ is trust evaluation of the TE.

For trust evaluation, as the evaluation is not the part of our model, TRs use a simple probabilistic trust evaluation

$$trust = 0.5 * directTrust + 0.5 * indirectTrust \tag{9}$$

- Direct trust is the percent of good transactions performed so far with the TE, and the default value is 0.5.
- Indirect Trust represents the reputation of TE in the community, and calculated as the average direct trust value of those who previously interacted with the TE.
- Witnesses are assumed to be honest, as our main objective is to evaluate the the proposed trust establishment model on TEs rather than evaluating a the trust evaluation model(s) used by TRs, which are supposed to differentiate honest form dishonest witnesses.

Having selected a TE, the TR then interacts with the selected TE. A TE can serve many users at a time. A TR does not always use the service in every round. The probability it requests the service, called its activity level and set to 50%.

Demanding levels of TRs can vary from non-demanding to highly demanding, we used three levels of demanding behavior, a set of highly demanding that use a high trust threshold between 0.8 and 1.0. The base demanding TRs use a trust threshold between 0.25 and 0.5. Regular demanding TRs use a trust threshold of 0.5.

After each transaction, the TR updates the credibility of the TE that has participated in the transaction. As we aim to compare the Performance of TEs equipped with TEIF and those equipped with the reputational incentive model of (Burnett, Norman, and Sycara 2011), all TEs are assumed to be honest and the only difference among among TEs, other than the name, is the trust establishment model. This way we can relate the difference in Performance to the model of trust establishment used. TEs are randomly set to either use TEIF or the reputational incentive model (Burnett, Norman, and Sycara 2011) at creation time, and they do not change that.

After each transaction, the involved TE updates the retention rate of the corresponding TR. At the end of each simulation step, TEs update the average retention rates and the value of the GSF.

Table 1 presents the number of agents and other parameters used in the proposed model and those employed in the environment.

Parameter	Value
Total number of TEs	100
Total Number of TRs	100
High Demand TRs	30
Low Demand TRs	40
Normal Demand TRs	30
GSF_{HGR}	0.5
GSF_{LGR}	1.5
GSF_{NGR}	0.9
CSF_{users}	1.5
$CSF_{choosers}$	1.0
$CSF_{discoverers}$	0.5
$CSF_{engaged}$	0.75
ET	0.75
DT	0.25
UT	0.4

Table 1: Values of Used Parameters

Experimental results

Percent of overall transactions In order to examine whether the proposed model helps TEs attract TRs to perform transactions with themselves, we measure the percent of the overall transaction carried out by the group of TEIF-equipped TEs and those implementing the reputational incentive model described in (Burnett, Norman, and Sycara 2011). Figure 1 presents the Performance of TEIF, under this measure, compared to the Performance of the other model. The charted value is calculated as the averaged value for ten different runs of the experiment. For each run, the value is calculated at the end of of each simulation step.

The figure shows that using the proposed model, TEs can achieve a higher percent of overall transactions, compared to those using the other model. At the beginning of the simulation, TEs that use TEIF respond to service requests with high UGs. As TEIF directed TEs begin modeling the retention of TRs, the delivered UGs is reduced and, therefore, the percent of overall transactions is reduced accordingly, before heading toward stability.

Figure 1: Percent of Transactions

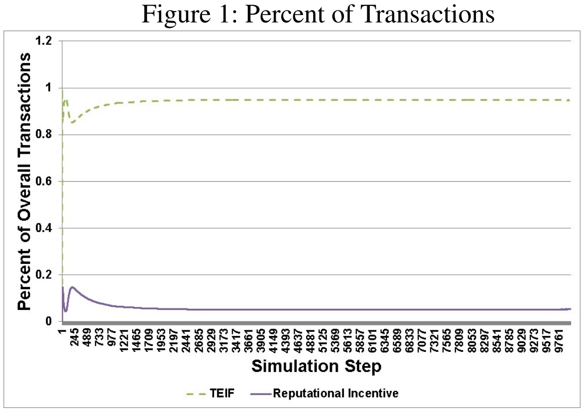

Average of the percent of delivered UG We examine the effect of the proposed model on the average provided UGs. We used the same experimental settings used to examine the percent of overall transactions. Figure 2 presents the average of percent of delivered UG per TE for those using TEIF compared to those using the reputational incentive (Burnett, Norman, and Sycara 2011). The charted value is calculated as the averaged value for ten different runs of the experiment. For each run, the average of the delivered UG is calculated at the end of of each simulation step. At the end of the experiments, the percent of the mean value relative to the maximum possible UG is calculated and charted.

The figure shows that using the proposed model, and TEs deliver a higher UG. In the beginning, TEIF enabled TEs offer UGs that is more than 1.6 of that provided by the other group of TEs. However, this ratio goes to about 1.26 after 500 simulation steps, and ends up as 1.06 by the end of the simulation. At the beginning of the simulation, TEs that use TEIF respond to service requests with high UGs, in order to attract TRs. As TEIF directed TEs begin modeling the retention of TRs, the delivered UGs is reduced and, therefore, the percent of overall transactions is reduced accordingly. As a result, those TEs tend to increase the delivered UG to avoid missing further transactions before heading toward stability. One factor that may affect the delivery UG, is the high competition between TEIF enabled TEs themselves. Analyzing the effect of this factor is left as a future work.

Figure 2: Percent of delivered UG

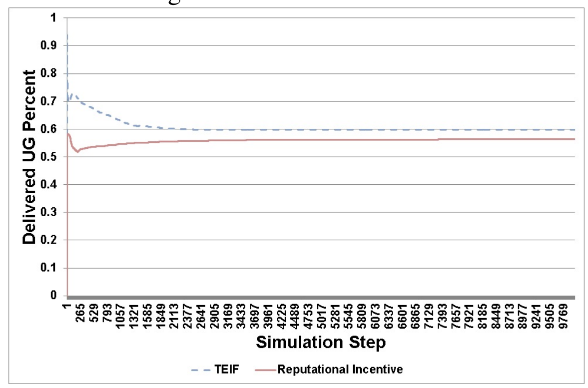

Conclusions and Future Work

In this work, we presented a trust establishment model for MASs using indirect feedback. The presented model allows TEs to adjust their behavior based on both: the behavior of the partner TR and the general conduct of the community of TRs. The aim of TEs is to enhance their trustworthiness estimation with the hope to be selected for future interactions. Simulation results indicate that trustworthy TEs can improve their portion of transactions if they adjust their delivered UG. Even though the increase in UG needs to be high in the beginning, it can be very little later on.

Currently, TEIF assumes a single service in the system. We would like to extend the model to address the case when individual TEs may provide different services, and more importantly when those services have different values. Dynamically determining parameter values and analyzing the effect of competition between TEIF-enabled TEs is left as future work.

References

Aref, A., and Tran, T. 2015a. A decentralized trustworthiness estimation model for open, multiagent systems (dtmas). *Journal of Trust Management* 2(1):3.

Aref, A., and Tran, T. 2015b. A trust establishment model in multi-agent systems.

Burnett, C.; Norman, T. J.; and Sycara, K. 2010. Bootstrapping trust evaluations through stereotypes. In *Proceedings of the 9th International Conference on Autonomous Agents and Multiagent Systems: Volume 1 - Volume 1*, AAMAS '10, 241–248. Richland, SC: International Foundation for Autonomous Agents and Multiagent Systems.

Burnett, C.; Norman, T. J.; and Sycara, K. P. 2011. Trust decision-making in multi-agent systems. In *IJCAI*, 115–120.

Castelfranchi, C.; Falcone, R.; and Marzo, F. 2006. Being trusted in a social network: Trust as relational capital. In Stlen, K.; Winsborough, W.; Martinelli, F.; and Massacci, F., eds., *Trust Management*, volume 3986 of *Lecture Notes in Computer Science*. Springer Berlin Heidelberg. 19–32.

Lerman, K., and Galstyan, A. 2001. A general methodology for mathematical analysis of multi-agent systems. *ISI-TR-529, USC Information Sciences Institute, Marina del Rey, CA*.

Luke, S.; Cioffi-Revilla, C.; Panait, L.; Sullivan, K.; and Balan, G. 2005. Mason: A multiagent simulation environment. *Simulation* 81(7):517–527.

Malik, Z., and Bouguettaya, A. 2009. *Trust Management for Service-Oriented Environments*. Springer Publishing Company, Incorporated, 1st edition.

Pinyol, I., and Sabater-Mir, J. 2013. Computational trust and reputation models for open multi-agent systems: a review. *Artificial Intelligence Review* 40(1):1–25.

Sabater, J., and Sierra, C. 2001. Regret: A reputation model for gregarious societies. In *Fourth workshop on deception fraud and trust in agent societies*, volume 70.

Saini, R., and Gautam, R. K. 2011. Establishment of dynamic trust among nodes in mobile ad-hoc network. In *Computational Intelligence and Communication Networks (CICN), 2011 International Conference on*, 346–349. IEEE.

Sen, S.; Sekaran, M.; Hale, J.; et al. 1994. Learning to coordinate without sharing information. In *AAAI*, 426–431.

Sen, S. 2013. A comprehensive approach to trust management. In *Proceedings of the 12th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2013)*.

Tran, T.; Cohen, R.; and Langlois, E. 2014. Establishing trust in multiagent environments: realizing the comprehensive trust management dream. In *Proceedings of the 17th International Workshop on Trust in Agent Societies,*.

Yu, H.; Shen, Z.; Leung, C.; Miao, C.; and Lesser, V. 2013. A survey of multi-agent trust management systems. *Access, IEEE* 1:35–50.

Zhang, J., and Cohen, R. 2007. Design of a mechanism for promoting honesty in e-marketplaces. In *Proceedings of the Twenty-Second AAAI Conference on Artificial Intelligence, July 22-26, 2007, Vancouver, British Columbia, Canada*, 1495–1500.