Design of a Multiagent-based E-Marketplace to Secure Service Trading on the Internet

Qin Li
Information Security Group
Royal Holloway
University of London
qin.li.2008@live.rhul.ac.uk

Keith M. Martin
Information Security Group
Royal Holloway
University of London
keith.martin@rhul.ac.uk

Jie Zhang
School of Computer Engineering
Nanyang Technological University Singapore
zhangj@ntu.edu.sg

ABSTRACT

Electronic marketplaces are not always easily regulated using traditional legal systems. As a result, suitable dispute prevention and resolution mechanisms for electronic marketplaces are of interest to design. In this paper, we design a multiagent-based e-marketplace where participants represented by autonomous software agents may be egocentric, strategic and even malicious. Our marketplace focuses on trading arbitrable and replicatable services, such as computational resources and data storage, over the Internet. We propose a novel dispute prevention and resolution mechanism that is theoretically proven to be able to induce good conduct for marketplace participants. Our marketplace also features cost-effectiveness, robustness and budget balance.

Categories and Subject Descriptors

I.2.11 [**ARTIFICIAL INTELLIGENCE**]: Distributed Artificial Intelligence – Intelligent agents, Multiagent systems

General Terms

Design, Theory

Keywords

Service Trading, Arbitrable and Replicatable Service, Incentive Mechanism, Dispute Resolution and Prevention

1. INTRODUCTION

An electronic marketplace for service trading on the Internet gathers users (both service providers and consumers) together and can provide a sound solution to benefit the users and society in general. Implementing an electronic marketplace as a multi-agent system is even more beneficial. Researchers in the field of artificial intelligence have been focusing on designing multiagent-based electronic marketplaces, where software agents act on behalf of their human users and organisations (service providers and consumers) to perform the tasks of service trading [9]. These include automating the processes of choosing appropriate trading partners, conducting trading transactions and post-purchase evaluation. With the assistance of software agents, the processes required by human users and organisations are offloaded and these processes can be considerably accelerated.

However, the individual agents in a multiagent-based electronic marketplace may be self-interested. They interact with each other to achieve their own goals and may therefore engage in deception. For example, provider agents may try to maximise their profit by delivering services of lower quality than they promised. Consumer agents may try to lie about the quality of the services they received. A fundamental requirement for a multiagent-based electronic marketplace is thus a highly efficient dispute prevention and resolution mechanism.

Some traditional approaches, such as state-enforced contractual guarantees, tend to be far less effective in multiagent-based marketplaces. There are several reasons for this [3, 12, 8, 4]. Firstly, the uptake of electronic marketplaces has tended to outpace the establishment of related legal regulation. Secondly, the one-time deal and multiple jurisdictional nature of online service trading presents challenges to conventional legal systems. These reasons result in traditional dispute prevention and resolution mechanisms becoming uncomfortably expensive and time consuming.

In this paper we propose a novel dispute resolution and prevention mechanism for multiagent-based marketplaces, based on some preliminary idea of [13]. We focus on services that meet the following properties. Firstly, the services exchanged in the multiagent-based electronic marketplace can be transformed into an "arbitrable" form at negligible cost. In the other words, the violation of a service provision agreement can easily be identified by a third party. We call such services *arbitrable*. Secondly, provision of the services can easily be replicated at negligible additional cost to service providers or benefit to consumers. We call such services *replicatable*. Both these conditions are rarely satisfied in a conventional marketplace trading tangible goods. However, for online services and digital products such as computational and data storage resources, both requirements can often be met. The first condition can normally be satisfied by applying a cryptographic non-repudiation of origin mechanism, such as a digital signature scheme [19]. The "non-tangible" nature of online services and digital products makes it feasible, at least in theory, to meet the second condition. For example, if a software vendor allows a consumer to download software that they have already successfully purchased, replicating this comes at little cost to the vendor and negligible benefit to the consumer.

On open networks such as the Internet, arbitrable and replicatable online services and digital products have become increasingly available. Some typical examples are computational and data storage resource exchange. The demand and supply of such services have been growing [10]. Grid and Cloud computing are both designed to facilitate such services. It is expected that the interest in electronic marketplaces for these types of service will only increase.

In the proposed multiagent-based marketplace, we address dispute prevention and resolution mechanism through the following approaches. We incorporate an autonomous *broker agent* and an *arbitrator agent*. The broker agent provides an intermediary service between providers and consumers. The arbitrator agent provides a dispute resolution and enforcement service. Based on the broker and arbitrator agents, we design an incentive mechanism for the multiagent-based marketplace. Lastly, we use a *reputation agent* to maintain the reputation of providers.

The electronic marketplace adopting our proposed mechanism has several important features. First, the marketplace induces *good conduct*. Participants have a strong incentive to comply with the policy of the marketplace. Second, the marketplace is *cost-effective*. It enables participants to find transaction partners that make the most cost-effective offer. Further, the cost of the operation of the marketplace to consumers and providers is within reasonable and acceptable limits. Another feature is *transaction security*, which ensures that consumers receive quality services from providers upon payment, and providers receive payment in full upon the provision of quality services to consumers. The fourth feature is *reputation robustness*. Provider agents' reputation is protected against malicious damage. Also, the marketplace provides a *budget balance* property. The operator of the electronic marketplace does not have incentive to make an unfair dispute resolution to favour any disputed party during arbitration.

The rest of the paper is organised as follows. Section 2 discusses some related work. Section 3 elaborates our proposed scheme. Section 4 conducts an analysis on the behaviour and strategy of participants of the marketplace. Section 5 discusses and justifies the properties of our proposed scheme. Section 6 concludes and identifies further issues.

2. RELATED WORK

There are some studies into the design of multiagent-based electronic marketplaces to promote honesty of provider and consumer agents. For example, trust-based incentive mechanisms, such as [6, 18], create incentives for providers to deliver satisfactory services and for consumers to truthfully report the results of the transactions with providers. However, our work takes advantage of a broker agent, arbitrator agent and reputation agent to provide a strong dispute resolution and prevention mechanism with additional features of cost-effectiveness, robustness and budget balance.

Similar to the notion of a broker agent, the use of a "mediator" has been studied and supported in the community of multiagent-based systems. For example, the benefit of adopting a mediator in a multiagent-based setting is discussed in [15, 1]. There are a few studies into arbitration services applied to electronic marketplaces. Milosevic et al. [14] demonstrate some benefits of an arbitrative enforcement service in an electronic marketplace. Daskalopulu et al. [7] present an approach for an artificial arbitrator undertaking dispute resolution in an electronic marketplace. However, our arbitrator agent design separately treats non-repudiatable misbehaviour and repudiatable misbehaviour.

There is a great deal of research dedicated to reputation systems used in electronic marketplaces. A number of studies, for example [5, 2, 17], demonstrate the effects of reputation mechanisms in electronic marketplaces. However, many reputation systems are vulnerable to threats. Our design of the multiagent-based marketplace provides an additional feature of robustness for the reputation agent against some common threats, which will be discussed in Section 5. In addition, different advanced reputation algorithms can be easily implemented into our marketplace.

3. DESCRIPTION OF THE MARKETPLACE

Our proposed marketplace provides consumer agents and provider agents with a trading platform for exchange of arbitrable and replicatable services. It consists of a broker agent, an arbitrator agent and a reputation agent (see Figure 1). The broker agent provides an intermediary service between provider agents and consumer agents. The arbitrator agent provides a service of dispute resolution and enforcement of resolution results. The reputation agent collects and provides reputation information about all provider agents. We require that the broker, arbitrator and reputation agents are trusted and interact honestly with each other. Although these components can be distributed, one convenient setting is to make them form a single entity, a *central server*, to perform all required functionalities (see Figure 1). We will hereafter adopt this setting for simplicity.

In this section we first provide an overview of the marketplace setting and then elaborate the operations of the reputation, broker and arbitrator agents.

3.1 Overview

There are some basic and common requirements for the marketplace setting:

- The identity of the central server is universally verifiable and every consumer and provider is represented by a unique identity.
- There is a cryptographic entity authentication mechanism, a cryptographic non-repudiation mechanism and a key management system in place to manage all the required cryptographic keys (including their distribution) so that every consumer agent and provider agent can be authenticated to the server.
- Services to be bought by consumer agents and sold by provider agents can be unambiguously specified.

- There is a method to evaluate the quality of service.
- There is a method for consumer agents to quantitatively predict the amount of services to buy (i.e. memory and bandwidth), and a method to verify the correctness of received services (i.e. the results of computational tasks performed).
- All participants, including consumers and providers, are rational decision makers[1], i.e. they do not blindly conform to, or divert from, stipulated rules, but choose their best action to maximise their payoff, according to their personal preferences [16].

A participant is admitted to the marketplace by the central server if the participant agrees with the rules of the system. Besides, during admission, the server has to ascertain that it is able to enforce payment for received services, fees and penalties resulting from misbehaviour. This guarantee can take various forms, such as a financial deposit. Upon admission, the server opens a financial account for a new participant with the balance equal to the amount of financial deposit. This account is reused for this participant for subsequent admissions. The server stops providing service to a participant if the balance of its account is insufficient to enforce payment, fees and penalties.

As shown in Figure 1, participants that are admitted to the system coordinate with the broker to perform transactions. Participants send requests to the broker and the broker provides *pair-off services* to find transaction partners for them. It also handles payments between paired-off participants. Participants coordinate with the arbitrator to resolve possible disputes which occur during transactions. Whenever a participant is aware of its partner's misbehaviour, it should immediately invoke arbitration services provided by the arbitrator agent. The details about the operations of the broker and arbitrator agents will be provided in Sections 3.3 and 3.4, respectively.

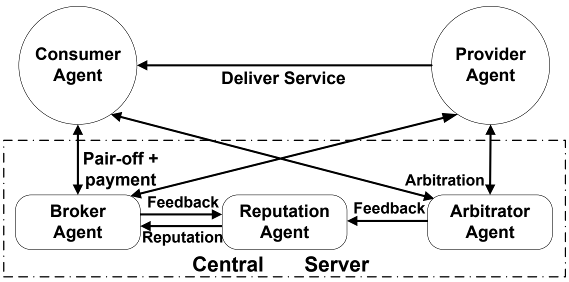

Figure 1: The Marketplace Structure

Within the central server, the reputation agent coordinates with the broker agent to exchange reputation information about provider agents. This coordination has two facets: the broker agent reports feedbacks concerning provider agents' performance to the reputation agent; the reputation agent provides the broker agent with the latest reputation information concerning provider agents. The arbitrator reports additional feedbacks concerning the performance of provider agents involved in disputes to the reputation agent.

[1]We will explain in Section 5 that this requirement can be relaxed to some extent.

3.2 Operation of the Reputation Agent

We define *misbehaviour* as behaviour which results in the expected outcome of the transaction being unachievable. We say that a participant *misbehaves* when it conducts misbehaviour, and a participant *behaves* when it does not conduct misbehaviour. For example, if a provider agent agrees to provide a consumer agent with a computational service, but sends an incorrect computational result, we say that the provider misbehaves. If the provider agent does not send any result to the consumer agent, it is also considered as misbehaving.

Our marketplace uses two types of feedback to describe the behaviour of a provider agent with respect to every transaction that the participant was previously involved in. If the provider agent is not found misbehaving, an r^+ feedback is reported for the provider agent. If it is found misbehaving, an $r^\times$ feedback is reported.

Given a provider agent i, the reputation agent defines X_i as the number of previous transactions in which participant i was not discovered misbehaving, which is the number of r^+ feedbacks reported about i. The reputation agent defines Y_i as the number of previous transactions in which i was found misbehaving, which is the number of $r^\times$ feedbacks reported about i. The reputation agent uses a reputation computation method to calculate the *reputation* of i. For example we apply the beta reputation computation approach [11] as follows to calculate the reputation of i:

$$R_i = \frac{X_i + 1}{X_i + Y_i + 2}.$$

Although our proposed marketplace still works without this reputation agent, incorporating it enhances the performance of the marketplace. Without the reputation agent, provider agents are still induced to engage in good conduct by the design of the broker and arbitrator agents. However, the reputation agent helps consumer agents to identify provider agents with good performance and to isolate and avoid those with poor performance. This also offers provider agents with an additional incentive to good conduct.

3.3 Operation of the Broker Agent

When a consumer agent wants to buy a service, it sends a request for purchase to the broker agent. It also sends a description of the service, its QoS requirements for the service, its requirement for the reputation value of its future transaction partner, an offer of price and, optionally, a priority list. When a provider agent wants to sell a service, it sends a request of sale to the broker agent. It also sends a description of the service, its capabilities, an expected price and, optionally, a priority list.

Upon receipt of a request, the broker agent firstly checks if the requesting agent has sufficient funds in its account to meet any possible payment, fees and penalties relating to this request. If not, the broker notifies the requesting agent and terminates this procedure. Otherwise, the broker agent performs a pair-off service so that the requested agent is partnered with the most suitable agent(s) from the other side. The pair-off procedure can be achieved in different ways, for example the broker agent makes a random decision amongst those that are equally qualified; the requesting agent makes its own decision upon the broker agent providing information of all qualified candidates from the other side; the broker agent negotiates with the two sides to reach a consensus between the two parties on their requirements and offers. If the pair-off is successful, the information concerning the new transaction will be added to an *unfinished transaction database*, which stores information concerning all unfinished transactions. This database has at least four fields, namely *transaction ID*, which uniquely identifies every transaction, *expiration time*, which specifies the expiration time of this transaction pre-agreed by the consumer and provider agents, *consumer agent ID* and *provider agent ID*, which specify the identities of the participants involved. Besides, the broker agent notifies the paired-off agents of the identities of their partners and the transaction ID.

Then the paired-off agents directly communicate with each other to conduct the transaction in the pre-agreed manner. Note that all messages exchanged between the paired-off agents should be non-repudiatable, i.e. they should be accompanied by a cryptographic commitment such as a digital signature. The receiver should validate the received commitment. During the transaction, if an agent is aware of its partner's misbehaviour, it should immediately invoke an arbitration service provided by the arbitrator agent. The operation of the arbitrator agent is detailed in Section 3.4.

Upon the expiration of the transaction, i.e. when the time reaches the expiration time recorded in the unfinished transaction database, if an arbitration has been invoked, regardless of the completion of the arbitration, then the broker terminates its procedure. If no arbitration has been invoked, then the broker agent debits the consumer agent and credits the provider agent by amount p, where p denotes their pre-agreed payment. Also the broker agent charges a service fee f_0, where $p > f_0 > 0$, from both the consumer and provider agents. Meanwhile, it reports to the reputation agent an r^+ feedback for the provider agent. Finally, the record of the transaction in the unfinished transactions database is deleted.

The procedure just described involves a set of actions that the agents are advised to comply with in order to successfully complete a transaction. However, it does not prevent misbehaviour, which directly results in an unsatisfactory outcome.

3.4 Operation of the Arbitrator Agent

Our approach to dealing with misbehaviour is based on the following idea. When misbehaviour occurs, we end it as quickly as possible and apply an appropriate penalty to the agent that misbehaves, as well as offer appropriate compensation to those that are affected. Our mechanism stipulates that whenever an agent is aware of its partner's misbehaviour, it should immediately invoke one of the two corresponding arbitration services that will be described in Section 3.4.1 and Section 3.4.2.

The choice between the two arbitration services should be made according to the characteristics of the misbehaviour. According to its repudiatability, we divide misbehaviour into two classes, *non-repudiatable misbehaviour* and *repudiatable misbehaviour*. Non-repudiatable misbehaviour occurs when the misbehaving agent is not able to later deny its misbehaviour. This is achieved by applying cryptographic mechanisms of non-repudiation of origin to all communication between providers and consumers. For example, a provider sending an incorrect result with its digital signature to the consumer is non-repudiatable misbehaviour. A consumer sending a signed task which requires more computational resources than it requested is considered as non-repudiatable misbehaviour. Repudiatable misbehaviour occurs when the agent that misbehaves is later able to deny its misbehaviour. A third party is not able to ascertain whether the accuser falsely censures the accused or the accused indeed misbehaved. For example, when an agent accuses its partner of not responding within a pre-agreed period of time, the arbitrator is not able to judge whether the agent is falsely accusing its partner or the partner indeed did not respond within the pre-agreed period of time. Hence, not responding within a pre-agreed period of time is a type of repudiatable misbehaviour.

3.4.1 Non-repudiatable Misbehaviour Arbitration

The operation of the arbitrator agent for non-repudiatable misbehaviour is as follows. An agent, either a consumer agent or a provider agent, sends a request to the arbitrator for arbitration, along with all related evidence. The arbitrator agent first checks if the received transaction ID, consumer ID and provider ID match a record in the unfinished transactions database. If so, the arbitration proceeds. If not, the agents requesting arbitration will be financially punished and the arbitration will terminate. The arbitrator agent verifies the validity of the evidence. If the provider agent is found misbehaving or false accusation, then the arbitrator agent punishes the provider agent by debiting an amount f_1, where $f_0 < f_1 < p + f_0$, and sends the reputation agent an $r^\times$ feedback for the provider agent. Meanwhile it charges the consumer agent no money. If the consumer agent is found misbehaving or making a false accusation, then the arbitrator agent punishes the consumer agent by debiting an amount $p + f_1$. Meanwhile it compensates the provider agent by an amount p and sends the reputation agent an r^+ feedback for the provider agent. The partnership between the two agents terminates. The record of the transaction in the unfinished transactions database is deleted. The financial punishment and compensation for the arbitration of non-repudiatable misbehaviour is summarised in Table 1.

Table 1: The Financial Punishment and Compensation for Non-repudiatable Misbehaviour Arbitration

Arbitration Result	Provider	Consumer
Provider misbehaving or false accusation	$-f_1$	0
Consumer misbehaving or false accusation	p	$-p - f_1$

3.4.2 Repudiatable Misbehaviour Arbitration

It is extremely challenging to arbitrate based on a report of repudiatable misbehaviour without any further investigation, since whether the accuser falsely censures the accused, or the accused indeed misbehaved, is unknown. Our mechanism stipulates a simple solution for such a situation. The solution is described as follows.

An agent, either a consumer or a provider, requests a repudiatable arbitration. The arbitrator agent first checks if the received transaction ID, consumer ID and provider ID match a record in the unfinished transactions database. If not, the agent requesting arbitration will be financially punished and the arbitration will terminate. If so, the arbitration proceeds and the arbitrator agent requests that both the consumer and provider agents repeat the communication between them for this transaction via the arbitrator agent. Then both agents communicate via the arbitrator.[2] If a misbehaviour occurs during the repeated communication, the arbitrator agent is able to detect it. If the provider agent is found misbehaving, then the arbitrator agent debits the provider agent by an amount $2f_1$, and sends an $r^\times$ feedback for the provider agent. Meanwhile, it charges no money from the consumer agent. If the consumer agent is found misbehaving, then the arbitrator agent debits the consumer agent by an amount $p + 2f_1$. Meanwhile, it credits the provider agent an amount p and sends an r^+ feedback for the provider agent. If no misbehaviour is detected at the end of the transaction, the arbitrator agent debits the consumer agent and credits the provider agent the pre-agreed payment p, since the transaction has now been completed. But it charges a service fee f_1, instead of f_0, from both the consumer and provider agents. The arbitrator agent sends an r^+ feedback for the provider agent. The partnership between the two agents is terminated. The record of the transaction in the unfinished transactions database is deleted. The financial punishment and compensation for repudiatable misbehaviour arbitration is summarised in Table 2.

Table 2: The Financial Punishment and Compensation for Repudiatable Misbehaviour Arbitration

Arbitration Result	Provider	Consumer
Provider misbehaving	$-2f_1$	0
Consumer misbehaving	p	$-p-2f_1$
No misbehaving	$p-f_1$	$-p-f_1$

4. PARTICIPANT BEHAVIOUR ANALYSIS

In this section, we theoretically analyse the strategies of participants. We show in each case that misbehaving is not the best strategy for any participant, independent of the strategy of the other participant.

Let $\mathcal{U}_i$ denote the overall utility, or benefit, that the agent i gains from a transaction. The utility of a consumer agent in one transaction is determined by three factors: service received, financial payment, time spent. Formally we define the utility function for a consumer agent i as follows:

$$\mathcal{U}_i = \mathcal{U}_i(s) + \mathcal{U}_i(b) + \mathcal{U}_i(t),$$

[2]Note that the repeated communication may not be the same as the original one.

where $\mathcal{U}_i(s)$, $\mathcal{U}_i(b)$ and $\mathcal{U}_i(t)$ respectively denote the utilities of the service received, financial payment and time spent. The utility of a provider agent in one transaction is determined by four factors: service provision, financial gain or loss (financial loss is possible if the agent is found misbehaving), time spent and reputation gain or loss. We formally define the utility function of a provider i as follows:

$$\mathcal{U}_i = \mathcal{U}_i(d) + \mathcal{U}_i(b) + \mathcal{U}_i(t) + \mathcal{U}_i(r),$$

where $\mathcal{U}_i(d)$, $\mathcal{U}_i(b)$, $\mathcal{U}_i(t)$ and $\mathcal{U}_i(r)$ respectively denote the utilities resulted from service provision, financial gain or loss, time spent, and reputation gain or loss.

We assume that individual utility functions have their own properties. For a consumer's utility on a received service $\mathcal{U}_i(s)$, we classify the received service s into two categories. We denote by s^+ and $s^\times$, respectively, a received service that does, or does not, meet the pre-agreed quality requirements. It is reasonable to assume $\mathcal{U}_i(s^\times) = 0$ and $\mathcal{U}_i(s^+) > 0$. We also assume that s is not a time critical service, i.e. $\mathcal{U}_i(s^+)$ remains the same when the service s^+ is received in one, or more than one transaction. For a provider's utility on its delivered service $\mathcal{U}_i(d)$, we also classify d into two categories. We denote by d^+ and $d^\times$, respectively, a service provided that does, or does not, meet the pre-agreed quality requirements. It is reasonable to assume that $\mathcal{U}_i(d) \leq 0$, $\mathcal{U}_i(d^+) < \mathcal{U}_i(d^\times)$, and $\mathcal{U}_i(d)$ is a non-increasing function against time.

For $\mathcal{U}_i(b)$, we denote by b the agent i's financial balance achieved from the transaction. It consists of the pre-agreed payment p and the two types of service fees f_0 and f_1. We assume that $\mathcal{U}_i(b)$ is a linearly increasing function.

For $\mathcal{U}_i(t)$, we denote by t the period of time that i spends on the transaction. We assume that $\mathcal{U}_i(t) \leq 0$ and $\mathcal{U}_i(t)$ is a linearly decreasing function. Besides, we denote by T the pre-agreed time for a transaction.

For $\mathcal{U}_i(r)$, we denote by r the feedback that is reported for the provider agent i for the transaction. We assume that no agent will attempt to gain an r^+ feedback at the cost of f_0, as the increase of reputation resulting from it is insignificant. Formally speaking, we assume that $\mathcal{U}_i(r = r^\times) < 0 < \mathcal{U}_i(r = r^+) < \mathcal{U}_i(b = f_0)$.

It is easily seen that if a consumer agent i sends a purchase request to the broker agent, then it will benefit from the transaction if the transaction is completed smoothly. Formally speaking

$$\mathcal{U}_i(s^+) + \mathcal{U}_i(b = -p - f_0) + \mathcal{U}_i(t = T) > 0. \quad (1)$$

Likewise, if a provider agent i sends a sale request, then it will benefit from the transaction if the transaction is completed smoothly. Formally speaking

$$\mathcal{U}_i(d^+) + \mathcal{U}_i(b = p - f_0) + \mathcal{U}_i(t = T) + \mathcal{U}_i(r = r^+) > 0. \quad (2)$$

4.1 Non-repudiatable Misbehaviour

In this section, we analyse participants' strategies with respect to non-repudiatable misbehaviour.

Lemma 1. *Given that a participant i has discovered its partner conducting some non-repudiatable misbehaviour and has acquired valid evidence, its best strategy is to request arbitration.*

Proof. Let $\mathcal{U}_i^Y$ and $\mathcal{U}_i^N$ denote the utilities that i gains if it does and does not request an arbitration, respectively. We prove this lemma by showing $\mathcal{U}_i^N < \mathcal{U}_i^Y$.

If i is a consumer agent, then i receives an unsatisfactory service $s^\times$ (as its partner has conducted some misbehaviour). If i requests an arbitration and provides the evidence, then the arbitration will result in no financial loss and gain for i (as the evidence is valid). The time that i spends on this transaction will be no longer than T (as once i has requested the arbitration, it stops spending time on this transaction.) If i does not request an arbitration, then it will be debited by amount $p + f_0$ (by the broker agent). The time that i spends on this transaction is T (as it has to involve in the transaction until it expires). Then $\mathcal{U}_i^Y$ and $\mathcal{U}_i^N$ can be formalised as follows:

$$\mathcal{U}_i^N = \mathcal{U}_i(s^\times) + \mathcal{U}_i(b = -p - f_0) + \mathcal{U}_i(t = T);$$
$$\mathcal{U}_i^Y = \mathcal{U}_i(s^\times) + \mathcal{U}_i(b = 0) + \mathcal{U}_i(t \leq T).$$

We have

$$\mathcal{U}_i^Y - \mathcal{U}_i^N = \mathcal{U}_i(b = p + f_0) + \mathcal{U}_i(t \leq T) - \mathcal{U}_i(t = T).$$

Because $\mathcal{U}_i(b = p + f_0) > 0$ and $\mathcal{U}_i(t \leq T) - \mathcal{U}_i(t = T) \geq 0$ (by the assumptions stated in the beginning of this section), we derive that $\mathcal{U}_i^N < \mathcal{U}_i^Y$.

Similarly, if i is a provider agent, we derive

$$\mathcal{U}_i^N = \mathcal{U}_i(d^\times) + \mathcal{U}_i(b = p - f_0) + \mathcal{U}_i(t = T) + \mathcal{U}_i(r = r^+);$$
$$\mathcal{U}_i^Y = \mathcal{U}_i(d^\times) + \mathcal{U}_i(b = p) + \mathcal{U}_i(t \leq T) + \mathcal{U}_i^Y(r = r^+).$$

It is also evident that $\mathcal{U}_i^N < \mathcal{U}_i^Y$. Hence the lemma holds for both consumer and provider agents. □

Theorem 1. *Conducting non-repudiatable misbehaviour is not the best strategy for a participant i.*

Proof. Similarly, let $\mathcal{U}_i^Y$ and $\mathcal{U}_i^N$ denote the utility that participant i gains from conducting and not conducting non-repudiatable misbehaviour, respectively. We will prove this lemma by showing $\mathcal{U}_i^Y < \mathcal{U}_i^N$. Note that Lemma 1 proves that the transaction partner of i will request arbitration if i conducts non-repudiatable misbehaivour.

If i is a consumer agent, then

$$\mathcal{U}_i^N = \mathcal{U}_i(s^+) + \mathcal{U}_i(b = -p - f_0) + \mathcal{U}_i(t = T);$$
$$\mathcal{U}_i^Y = \mathcal{U}_i(s^\times) + \mathcal{U}_i(b = -p - f_1) + \mathcal{U}_i(t \leq T).$$

It is evident that $\mathcal{U}_i^Y < \mathcal{U}_i^N$. If i is a provider agent, we derive

$$\mathcal{U}_i^N = \mathcal{U}_i(d^+) + \mathcal{U}_i(b = p - f_0) + \mathcal{U}_i(t = T) + \mathcal{U}_i(r = r^+);$$
$$\mathcal{U}_i^Y = \mathcal{U}_i(d^\times) + \mathcal{U}_i(b = -f_1) + \mathcal{U}_i(t \leq T) + \mathcal{U}_i(r = r^\times).$$

It is evident that $\mathcal{U}_i^Y < 0$. Besides $\mathcal{U}_i^N > 0$ by Inequality (2). Hence $\mathcal{U}_i^Y < \mathcal{U}_i^N$. Therefore this lemma holds for both consumer and provider agents. □

Theorem 2. *Given that a participant i has not discovered its partner conducting non-repudiatable misbehaviour, it is not the best strategy for i to request arbitration.*

Proof. Let $\mathcal{U}_i^N$ denote the utility that i gains if it does not request an arbitration. Let $\mathcal{U}_i^{Y_0}$ and $\mathcal{U}_i^{Y_1}$ denote the utilities that i gains if it requests an arbitration. The difference between $\mathcal{U}_i^{Y_0}$ and $\mathcal{U}_i^{Y_1}$ lies as follows. If i is a consumer agent, then $\mathcal{U}_i^{Y_0}$ and $\mathcal{U}_i^{Y_1}$ denote the utilities of i in the cases where i has received a satisfactory or unsatisfactory service at the moment of requesting the arbitration, respectively. If i is a provider agent, then $\mathcal{U}_i^{Y_0}$ and $\mathcal{U}_i^{Y_1}$ denote the utilities of i in the cases where i has provided a satisfactory or unsatisfactory service at the moment of requesting the arbitration, respectively. We prove this lemma by showing $\mathcal{U}_i^{Y_0} < \mathcal{U}_i^N$ and $\mathcal{U}_i^{Y_1} < \mathcal{U}_i^N$.

If i is a consumer agent, then

$$\mathcal{U}_i^N = \mathcal{U}_i(s^+) + \mathcal{U}_i(b = -p - f_0) + \mathcal{U}_i(t = T);$$
$$\mathcal{U}_i^{Y_0} = \mathcal{U}_i(s^+) + \mathcal{U}_i(b = -p - f_1) + \mathcal{U}_i(t = T);$$
$$\mathcal{U}_i^{Y_1} = \mathcal{U}_i(s^\times) + \mathcal{U}_i(b = -p - f_1) + \mathcal{U}_i(t \leq T).$$

It is evident that $\mathcal{U}_i^{Y_0} < \mathcal{U}_i^N$ and $\mathcal{U}_i^{Y_1} < \mathcal{U}_i^N$.

If i is a provider agent, then

$$\mathcal{U}_i^N = \mathcal{U}_i(d^+) + \mathcal{U}_i(b = p - f_0) + \mathcal{U}_i(t = T) + \mathcal{U}_i(r = r^+);$$
$$\mathcal{U}_i^{Y_0} = \mathcal{U}_i(d^+) + \mathcal{U}_i(b = -f_1) + \mathcal{U}_i(t = T) + \mathcal{U}_i(r = r^\times);$$
$$\mathcal{U}_i^{Y_1} = \mathcal{U}_i(d^\times) + \mathcal{U}_i(b = -f_1) + \mathcal{U}_i(t \leq T) + \mathcal{U}_i(r = r^\times).$$

It is evident that $\mathcal{U}_i^{Y_0} < \mathcal{U}_i^N$ and $\mathcal{U}_i^{Y_1} < \mathcal{U}_i^N$. Therefore the lemma holds for both consumer and provider agents. □

To summarise, it is not the best strategy for a participant to conduct non-repudiatable misbehaviour or falsely accuse its partner of conducting non-repudiatable misbehaviour.

4.2 Repudiatable Misbehaviour

In this section, we analyse participants' strategies with respect to repudiatable misbehaviour.

Lemma 2. *Given that a consumer agent i has conducted a repudiatable misbehaviour but its partner does not request arbitration, then it is the best strategy for i itself to request arbitration.*

Proof. Let $\mathcal{U}_i^Y$ and $\mathcal{U}_i^N$ denote the utilities that i gains if it itself does or does not request arbitration, respectively. We prove the lemma by showing $\mathcal{U}_i^N < \mathcal{U}_i^Y$. We derive

$$\mathcal{U}_i^N = \mathcal{U}_i(s^\times) + \mathcal{U}_i(b = -p - f_0) + \mathcal{U}_i(t = T);$$
$$\mathcal{U}_i^Y = \mathcal{U}_i(s^+) + \mathcal{U}_i(b = -p - f_1) + \mathcal{U}_i(t = 2T).$$

We have

$$\begin{aligned}\mathcal{U}_i^Y - \mathcal{U}_i^N &= \mathcal{U}_i(s^+) + \mathcal{U}_i(b = f_0 - f_1) + \mathcal{U}_i(t = T)\\ &> \mathcal{U}_i(s^+) + \mathcal{U}_i(b = -p - f_0) + \mathcal{U}_i(t = T)\\ &> 0,\end{aligned}$$

by the assumptions and inequality (1) stated in the beginning of this section. □

LEMMA 3. *Given that a provider agent i is behaving and has discovered its partner conducting some repudiatable misbehaviour, its best strategy is to request arbitration.*

PROOF. Let $\mathcal{U}_i^Y$ denote the utility that i gains if it requests arbitration and $\mathcal{U}_i^N$ denote the utility if i does not do so. We prove the lemma by showing $\mathcal{U}_i^N \le \mathcal{U}_i^Y$.

We will prove it in two steps. Firstly we will consider the situation where the consumer will not behave again before the transaction expires. Secondly we will consider the situation where the consumer will do so.

Let us consider the former case. We derive that

$$\mathcal{U}_i^N = \mathcal{U}_i(d^\times) + \mathcal{U}_i(b = p - f_0) + \mathcal{U}_i(t = T) + \mathcal{U}_i(r = r^+);$$
$$\mathcal{U}_i^Y = \mathcal{U}_i(d^\times) + \mathcal{U}_i(b = p) + \mathcal{U}_i(t \le 2T) + \mathcal{U}_i(r = r^+).$$

It is evident that $\mathcal{U}_i^N < \mathcal{U}_i^Y$.

Now we consider the latter case. We prove in Lemma 2 that the consumer will trigger the arbitration if the provider does not do so. Therefore we derive that

$$\mathcal{U}_i^N = \mathcal{U}_i(d^+) + \mathcal{U}_i(b = p - f_1) + \mathcal{U}_i(t = 2T) + \mathcal{U}_i(r = r^+);$$
$$\mathcal{U}_i^Y = \mathcal{U}_i(d^+) + \mathcal{U}_i(b = p - f_1) + \mathcal{U}_i(t = 2T) + \mathcal{U}_i(r = r^+).$$

It is easily seen that $\mathcal{U}_i^N = \mathcal{U}_i^Y$. Therefore, in both cases, $\mathcal{U}_i^N \le \mathcal{U}_i^Y$. □

LEMMA 4. *Given that a consumer agent i is behaving and has discovered its partner conducting some repudiatable misbehaviour, its best strategy is to request arbitration.*

PROOF. Let $\mathcal{U}_i^N$ denote the utility that i gains if it does not request arbitration. Let $\mathcal{U}_i^{Y_0}$ and $\mathcal{U}_i^{Y_1}$ denote i's utilities if it requests arbitration and during the arbitration the provider behaves or continues misbehaving, respectively. We prove the lemma by showing $\mathcal{U}_i^N < \mathcal{U}_i^{Y_0}$ and $\mathcal{U}_i^N < \mathcal{U}_i^{Y_1}$. We know that

$$\mathcal{U}_i^N = \mathcal{U}_i(s^\times) + \mathcal{U}_i(b = -p - f_0) + \mathcal{U}_i(t = T);$$
$$\mathcal{U}_i^{Y_0} = \mathcal{U}_i(s^+) + \mathcal{U}_i(b = -p - f_1) + \mathcal{U}_i(t \le 2T);$$
$$\mathcal{U}_i^{Y_1} = \mathcal{U}_i(s^\times) + \mathcal{U}_i(b = 0) + \mathcal{U}_i(t \le 2T).$$

We have

$$\begin{aligned}\mathcal{U}_i^{Y_0} - \mathcal{U}_i^N &= \mathcal{U}_i(s^+) + \mathcal{U}_i(b = f_0 - f_1) + \mathcal{U}_i(t \le T)\\ &> \mathcal{U}_i(s^+) + \mathcal{U}_i(b = -p - f_0) + \mathcal{U}_i(t = T)\\ &> 0,\end{aligned}$$

by the assumptions and Inequality (1) stated in the beginning of this section. Hence $\mathcal{U}_i^N < \mathcal{U}_i^{Y_0}$. Besides, it is evident that $\mathcal{U}_i^N < \mathcal{U}_i^{Y_1}$. □

THEOREM 3. *Conducting repudiatable misbehaviour is not the best strategy for a participant i.*

PROOF. Let $\mathcal{U}_i^N$ denote the utility of i if it behaves. Let $\mathcal{U}_i^{Y_0}$ and $\mathcal{U}_i^{Y_1}$ denote i's utilities if i initially conducts repudiatable misbehaviour and during the arbitration it behaves, or continues misbehaving, respectively. Note that Lemma 3 and Lemma 4 prove that if i conducts repudiatable misbehaviour then its partner will request arbitration. We prove this lemma by showing that $\mathcal{U}_i^{Y_0} < \mathcal{U}_i^N$ and $\mathcal{U}_i^{Y_1} < \mathcal{U}_i^N$.

If i is a consumer agent, then

$$\mathcal{U}_i^N = \mathcal{U}_i(s^+) + \mathcal{U}_i(b = -p - f_0) + \mathcal{U}_i(t = T);$$
$$\mathcal{U}_i^{Y_0} = \mathcal{U}_i(s^+) + \mathcal{U}_i(b = -p - f_1) + \mathcal{U}_i(t = 2T);$$
$$\mathcal{U}_i^{Y_1} = \mathcal{U}_i(s^\times) + \mathcal{U}_i(b = -p - 2f_1) + \mathcal{U}_i(t = 2T).$$

It is evident that $\mathcal{U}_i^{Y_0} < \mathcal{U}_i^N$ and $\mathcal{U}_i^{Y_1} < \mathcal{U}_i^N$. If i is a provider agent, then

$$\mathcal{U}_i^N = \mathcal{U}_i(d^+) + \mathcal{U}_i(b = p - f_0) + \mathcal{U}_i(t = T) + \mathcal{U}_i(r = r^+);$$
$$\mathcal{U}_i^{Y_0} = \mathcal{U}_i(d^+) + \mathcal{U}_i(b = p - f_1) + \mathcal{U}_i(t = 2T) + \mathcal{U}_i(r = r^+);$$
$$\mathcal{U}_i^{Y_1} = \mathcal{U}_i(d^\times) + \mathcal{U}_i(b = -2f_1) + \mathcal{U}_i(t = 2T) + \mathcal{U}_i(r = r^\times).$$

It is also evident that $\mathcal{U}_i^{Y_0} < \mathcal{U}_i^N$ and $\mathcal{U}_i^{Y_1} < \mathcal{U}_i^N$. □

THEOREM 4. *Given that a participant i is behaving and has not discovered its partner conducting any repudiatable misbehaviour, it is not the best strategy for the participant to request arbitration.*

PROOF. Let $\mathcal{U}_i^N$ denote i's utility if i does not request an arbitration. Let $\mathcal{U}_i^{Y_0}$ and $\mathcal{U}_i^{Y_1}$ denote i's utilities if it requests arbitration and behaves, or misbehaves, during the arbitration, respectively. By the similar argument as in Theorem 3, we derive that $\mathcal{U}_i^{Y_0} < \mathcal{U}_i^N$ and $\mathcal{U}_i^{Y_1} < \mathcal{U}_i^N$. □

To summarise, it is not the best strategy for a participant to conduct repudiatable misbehaviour or falsely accuse its partner of conducting repudiatable misbehaviour.

5. PROPERTIES OF OUR MECHANISM

In this section, we justify the features of our mechanism that we claimed in Section 1.

5.1 Induction of Good Conduct

We have shown, in Section 4, that conducting misbehaviour and false accusation is not the best strategy for any participant in our model. Further, the reputation information of provider agents provides them with additional incentive to good conduct. Therefore, the best strategy for a participant is to engage in good conduct, i.e. complying with the policy of the marketplace.

5.2 Cost-effectiveness

With assistance of the broker agent, it is easier for a participant to find transactional partners that make the most cost-effective offer. With respect to the cost of the operation of the marketplace, several aspects need to be considered. One is the computational cost. Each participant has to apply a non-repudiation mechanism on every message that is sent to its partner. Upon receiving a message from its partner, the participant has to verify the correctness of the attached cryptographic digest. Another aspect is the communication cost. Each participant has to interact with the server at the beginning of a transaction and during the resolution of a dispute. Since it is software agents that conduct transactions in the proposed multiagent-based marketplace, it is sensible to assume that each participant has reasonably abundant computational power and communication bandwidth. Hence, these computational and communication burdens can easily be borne. The third aspect is the time cost. There is a time delay between a participant sending any request to the server and receiving a response. This time delay can be alleviated by improving the performance of the central server. Such improvement can be driven by the competition among multiple marketplaces applying our scheme. The fourth aspect is the financial cost. Participants have to pay fees to the server on a per transaction basis. This cost also can be reduced to a reasonable and acceptable range driven by the competition among multiple marketplaces. Thus the various costs of the operation of the marketplace are easily borne or should establish themselves at reasonable and acceptable levels.

5.3 Transaction Security

From the perspective of consumers, if they do not misbehave then there are only two possible outcomes from a transaction. They either receive a satisfactory service at a financial cost of $p + f_0$ or $p + f_1$, or they receive an unsatisfactory service with no financial loss. In the latter case, the consumer can complete the task by repeatedly requesting new transactions until it receives a quality service. From the perspective of provider agents, if they do not misbehave then they receive an amount of $p - f_0$, $p - f_1$ or p. Hence, transaction security is achieved for consumer and providers.

5.4 Reputation Robustness

In this section, we discuss the robustness of the reputation mechanism incorporated into the marketplace. Such robustness is in fact provided by our marketplace design. As a result, the reputation assets of provider agents maintained by the marketplace is protected against all known major threats, which are summarised in [11, 17]. We categorise these threats into a few groups and show how our proposed marketplace makes the reputation mechanism immune to them, as follows.

Low incentive for reporting feedback. In our scheme, this can be interpreted as consumer agents having a low incentive to report their observation of their partners' behaviour. However, in this scheme, feedbacks are reported by the broker and arbitrator agents on the behalf of consumer agents. Every transaction will be reported a feedback according to the behaviour of the provider agent involved. Hence, the reputation agent will receive feedbacks regarding all existing transactions from the broker and arbitrator agents.

Dishonest feedback. This is an attack where consumers report untruthful feedbacks. In our mechanism, it can be interpreted as consumers falsely requesting arbitration when they do not find misbehaviour from their partners, or not requesting arbitration when they do discover misbehaviour from their partners. We have shown, in Section 4, that such behaviour is not the best strategy for them.

Ballot stuffing and bad-mouthing. These are attacks where multiple participants, possibly including both consumer and provider agents, collude together to boost or damage the reputation of a particular provider agent. In our mechanism, with respect to ballot stuffing, colluding provider and consumer agents have to each bear service fees f_0 in order to get an r^+ feedback reported for the provider agent. If f_0 is properly configured during implementation, ballot stuffing will be too expensive to conduct. With respect to bad-mouthing, as long as a provider behaves, its $r^\times$ feedback will never get increased, no matter how many consumer agents accuse it of misbehaviour. Hence, bad-mouthing is prevented.

Reporting feedback for non-existent transactions. Our mechanism only allow a feedback to be reported by the broker or arbitrator agent on the behalf of a consumer agent in the event where a transaction exists. Hence no feedback will be reported for non-existent transactions.

5.5 Budget Balance

With respect to an arbitration for non-repudiatable misbehaviour, the arbitrator agent gains the same amount of service fee f_1 regardless of its arbitration result. With respect to an arbitration for repudiatable misbehaviour, the arbitrator agent gains the same service fee $2f_1$ regardless of its arbitration decision. Therefore, the arbitrator agent has no incentive to make an unfair arbitration to favour any disputed participant.

5.6 Relaxation of Participants' Rationality

Although we assume that participants are rational, this assumption can be relaxed to an extent without affecting the robustness of the mechanism. In more specific term, our proposed scheme is robust as long as participants are rational about requesting an arbitration, regardless of their rationality about conducting misbehaviour. For example, our marketplace still works even if provider agents intentionally provide unsatisfactory services (this is an irrational behaviour as it reduces the utilities of the misbehaving provider agents) as long as consumer agents are rational (requesting an arbitration when they are aware of their transaction partner's misbehaviour).

6. CONCLUSIONS AND FUTURE WORK

We have proposed a multiagent-based marketplace to secure the trading of arbitrable and replicatable services. By incorporating a broker, an arbitrator and a reputation agent, it provides an incentive and a reputation mechanism for dispute prevention and resolution. It features good conduct induction, cost-effectiveness, transaction security, reputation robustness, budget balance and relaxation of participants' rationality.

Some issues arise from our proposed marketplace which are worth addressing in future work. Successful operation of a multiagent-based marketplace based on our model relies on the components of the central server. The security of the central server components becomes crucial. It would be interesting to explore how to reduce risks of attacks on these components, perhaps by distributing their functionalities.

We have only applied a generic reputation computation algorithm. It may be of interest to design a customised reputation computing method providing more accurate and richer reputation information.

7. REFERENCES

[1] I. Ashlagi, D. Monderer, and M. Tennenholtz. Mediators in position auctions. *Games and Economic Behavior*, 67(1):2–21, 2009.

[2] S. Ba and P. A. Pavlou. Evidence of the effect of trust building technology in electronic markets: Price premiums and buyer behavior. *MIS quarterly*, 26(3):243–268, 2002.

[3] S. Ba, A. B. Whinston, and H. Zhang. Building trust in online auction markets through an economic incentive mechanism. *Decision Support Systems*, 35(3):273–286, 2003.

[4] Y. Bakos and C. Dellarocas. Cooperation without enforcement? a comparative analysis of litigation and online reputation as quality assurance mechanisms. In *Proc. 23rd Internat. Conf. Inform. Systems*, pages 127–142, 2002.

[5] K.-Y. Chen, T. Hogg, and N. Wozny. Experimental study of market reputation mechanisms. In J. S. Breese, J. Feigenbaum, and M. I. Seltzer, editors, *ACM Conference on Electronic Commerce*, pages 234–235. ACM, 2004.

[6] R. Dash, S. Ramchurn, and N. Jennings. Trust-based mechanism design. In *Proceedings of the Third International Joint Conference on Autonomous Agents and Multiagent Systems*. IEEE Computer Society, 2004.

[7] A. Daskalopulu, T. Dimitrakos, and T. Maibaum. Evidence-based electronic contract performance monitoring. *Group decision and negotiation*, 11(6):469–485, 2002.

[8] C. Dellarocas. The digitization of word of mouth: Promise and challenges of online feedback mechanisms. *Management Science*, 49(10):1407–1424, 2003.

[9] R. Guttman, A. Moukas, and P. Maes. Agent-mediated electronic commerce: A survey. *The Knowledge Engineering Review*, 13(02):147–159, 2001.

[10] The Insight Research Corporation. *Grid Computing: A vertical Market Perspective 2006—2011*, 2006.

[11] A. Jøsang, R. Ismail, and C. Boyd. A survey of trust and reputation systems for online service provision. *Decision Support Systems*, 43(2):618–644, 2007.

[12] P. Kollock. The production of trust in online markets. *Advances in group processes*, 16(1):99–123, 1999.

[13] Q. Li and K. M. Martin. A secure marketplace for online services that induces good conduct. In *Short Paper Proceedings of the IFIP WG 11.11 International Conference on Trust Management*, 2010. ISSN:2079-2263.

[14] Z. Milosevic, A. Jøsang, T. Dimitrakos, and M. A. Patton. Discretionary enforcement of electronic contracts. In *Proc. of the Sixth International Enterprise Distributed Object Computing Conference, Lausanne, Switzerland*, pages 39–50, 2002.

[15] D. Monderer and M. Tennenholtz. Strong mediated equilibrium. *Artificial Intelligence*, 173(1):180–195, 2009.

[16] M. Osborne. *An Introduction to Game Theory*. Oxford University Press, 2004.

[17] P. Resnick and R. Zeckhauser. Trust among strangers in Internet transactions: Empirical analysis of eBay's reputation system. In M. R. Baye, editor, *The Economics of the Internet and E-Commerce*, volume 11 of *Advances in Applied Microeconomics*, pages 127–157. Elsevier Science, 2002.

[18] J. Zhang and R. Cohen. Design of a mechanism for promoting honesty in e-marketplaces. In *AAAI Conference on Artificial Intelligence (AAAI)*, pages 1495–1500. AAAI Press, 2007.

[19] J. Zhou. *Non-repudiation in Electronic Commerce*. Artech House, 2001.