Cooperative Interactions: An Exchange Values Model

Maíra R. Rodrigues* and Michael Luck

School of Electronics and Computer Science, University of Southampton, Southampton, UK
{mrm03r,mml}@ecs.soton.ac.uk

Abstract. In non-economic cooperative applications with resource constraints, explicitly motivating cooperation is important so that autonomous service providers have incentives to cooperate. When participants of such applications have different skills and expectations over services, it may be that an agent receives less than expected from a cooperation. A decision-making strategy over interactions in this context must consider not only the motivation to cooperate, but also which interactions to perform to cope with resource limitations. In this paper, we present a computational approach for modelling non-economic cooperative interactions based on the theory of *exchange values*. Here, exchange values are used to motivate cooperative interactions, and to allow agents to identify successful and unsuccessful cooperations with others, in order to limit service provision and to improve the number of successful interactions. We also present a scenario in which agents participate in a cooperative application in the bioinformatics domain, and show how agents can improve their interactions using the proposed approach.

1 Introduction

According to [1], it can be useful in certain contexts to view a society as a market, in which individuals exchange goods, services and ideas to achieve their goals. While markets typically involve monetary exchanges, many do not necessarily involve economic capital. For example, in computer-supported scientific communities like bioinformatics, different types of information and tools can be exchanged in a cooperative way in order to improve individual or global results [2].

Although examples of systems that support cooperative bioinformatics applications already exist [3,4,5], there are still issues to be addressed to allow effective cooperation between participants [6]. In particular, resources need to be managed sensibly because they are provided free of charge. This is because services in bioinformatics generally require the processing of large amounts of data, so that responding to a request involves significant computational resources. In addition, the kind of automated experiments that arise in this domain tend to generate more service requests than if they were performed manually. Thus, both the increasing number of requests that automated experiments can generate, and the large amount of computation resources that these requests need, place a heavy overload on service providers, which can limit the number of requests that can

* The first author is supported by Coordenação de Aperfeiçoamento de Pessoal de Nível Superior (CAPES) of the Brazilian Ministry of Education.

P. Noriega et al. (Eds.): COIN 2006 Workshops, LNAI 4386, pp. 356–371, 2007.

be processed. These issues provide the context in which we consider non-economic cooperative interactions.

Explicitly motivating cooperation is also important for non-economic exchanges. Considering a system with self-interested entities, a service provider has an incentive to cooperate if it receives a service in return from the requester, either immediately or in the near future. Cooperation with immediate reciprocation is easier to model and to check whether the cooperative interaction is genuine (mutual), since it involves concrete actions that can be clearly observed by both entities. However, immediate reciprocation is not always possible since the provider may not need any service by the time the interaction takes place, or the requester may only be available to provide a service in return in the near future.

Cooperative situations in which reciprocation is not immediate raise interesting issues for modeling cooperative agents. First, there is no guarantee that the requester agent will reciprocate in the future. Second, the provider must receive some *value* in return from the requester, so the provider is motivated to cooperate even in the lack of a concrete, immediate return. Third, if we consider that in a multi-agent system agents might have different perspectives of the same service due, for example, to individual preferences and the relevance of each service to goals, a provider can evaluate a service it gets in return in the future as being of less quality than the service it provided in the past; that is, it may be that an agent receives less than expected from a cooperation.

In this context, the focus of this paper is on cooperative applications in which interactions are based on *deferred reciprocity*, and deal with the issues of modeling cooperative agents that act in such applications. We propose a model based on *exchange values* [7] to motivate such cooperative interactions, and to allow agents to identify successful and unsuccessful cooperations with others, in order to limit service provision. We also analyse the different reasons why interactions succeed as a result of personal influences, such as personal goals and preferences, on the evaluation of services, and how agents can react to improve the quality of their cooperative interactions.

Here, exchange values represent the agents' individual evaluation of provided and received services, and are associated with their interactions, indicating the effort, cost, or satisfaction to each agent. We believe that this approach is suitable for addressing the issue of motivating interactions in cooperative, non-economic applications, especially in cases of deferred reciprocation, in the sense that exchange values provide a system of credits and debts that motivates interactions by giving expectations of future gains (for example, a credit that is gained by performing a service can be charged in the future). Moreover, exchange values allow agents to analyse the outcome of interactions in terms of whether services they receive compensate for services they provide, and use this information to decide about future interactions, which is important when there are different types of services available and there is a need to limit service provision.

The key contribution of this paper is a computational approach for motivating and modelling non-economic cooperative interactions, based on the theory of exchange values.

The paper begins with an introduction to the theory of exchange values, including its advantages for modelling cooperative autonomous agents, followed by an analysis of the different cooperative interactions that can occur between autonomous agents in a non-economic cooperative scenario. We end by presenting an experimental simulation of this scenario and discussing obtained results.

2 Exchange Values

The application of exchange values for modelling interactions between agents was first proposed in [8], in which the mechanism for reasoning over interactions is based on Piaget's theory [7].

In previous work [9], exchange values have been used for addressing the problem of partner selection in dynamic and resource-constrained environments. Here, agents reduce their effort in finding available interaction partners by identifying those that are more likely to accept requests. However, the analysis of interactions in terms of whether the services that agents receive give some compensation for those that they provide, which is relevant to the kind of applications we focus on this paper, is not addressed.

In response, this paper proposes the use of exchange values in agent interactions not only to motivate service provision, but also to allow agents to analyse the outcome of their interactions in terms of gains and losses. We argue that this approach has the following advantages for cooperative and non-economic applications:

– autonomous agents can reason about continuing or stopping cooperative relations, instead of assuming indefinite cooperation, which is only possible with benevolent agents; and
– agents are motivated to provide good quality services to ensure that they are valorised by requesters.

Piaget's theory of exchange values, and the dynamics of exchange values — the way they vary according to gains and losses of values during interactions — are introduced in the next sections.

2.1 Piaget's Theory of Exchange Values

Social exchange is the particular interaction in which an individual performs an action on behalf of another and vice-versa. The *theory of exchange values* was proposed by Piaget [7] as an analysis of human social exchanges and the reasons for their persistence or discontinuity. More specifically, Piaget argues that in all social interactions in which one individual acts on behalf of another there is an *exchange of values* between them. These values result from each individual's evaluation of the provided or received action over a common scale of values.

According to Piaget, every action and reaction of two interacting individuals towards each other has an influence on their values: if the action is useful and beneficial it increases their values; if the action is harmful and disadvantageous it decreases their values; and if the action is neutral their values remain the same. Thus, when two individuals interact and one provides a service that is valuable to the other, three situations can happen as described below.

1. The individual that received the service can *pay back* the provider by giving an object or providing another service in return (*immediate exchange*). This is the

case, for example, if a researcher who is submitting a paper to a conference receives comments on his paper from a colleague, and returns comments on his colleague's paper which is being submitted to the same conference.
2. The receiver just *valorises* the provider by expressing gratitude or approval, instead of giving something immediately in return (*deferred exchange*). Using the same example as above, the researcher who receives the comments on his paper expresses gratitude and *valorises* his colleague's action.
3. The receiver neither returns a service to nor valorises the provider (*no exchange*).

All reactions of the receiver have also an effect on the provider's values. In the immediate exchange, the receiver returns a material action to the provider (comments on a paper), which constitutes an *actual*, concrete value for the latter (the comments are valuable to improve the quality of the paper). In the deferred exchange, the receiver returns an abstract action to the provider (a word or a gesture of approval, gratitude, etc), which constitutes a *virtual* value for the latter, in the sense that his valorisation gives him reputation, respect and authority, which are values he can use to get some benefit in future interactions (next time he is writing a paper, he can then ask the receiver for comments). The third reaction, however, was disadvantageous for the provider since the receiver did not reciprocate the action in any way, and the latter is ultimately devalorised by the provider as ungrateful or unjust.

The values that are exchanged between individuals are clear when concrete objects are involved in the exchange, like the immediate exchange described above. However, when the exchange involves virtual values, like in the *deferred exchange*, a more detailed analysis is needed. When two individuals α (the provider) and β (the receiver) interact, the performed action is for α an *actual renouncement*, since it requires the expenditure of time and effort, while for β it is an *actual satisfaction* or gain. Now, if β immediately performs an action in return to α, β has an actual renouncement and α has an actual satisfaction. At the end of this interaction, both individuals had an actual satisfaction (they received a concrete action). If there is a valorisation of α by β instead, as a reaction for the received action, this valorisation is for α a reward, a *virtual credit* that it can draw upon in the future, and for the receiver the valorisation constitutes a promise, a *virtual debt*, in the sense that the receiver feels obliged to return the favour to the provider in the future. At the end of this interaction, both individuals acquire virtual values, instead of only actual values as in the immediate cooperation. When the exchange is not immediate, the exchange values in an interaction between two individuals α and β are, therefore, the following as shown in Figure 1:

– the *renouncement* of the provider α on performing a service to β (r_α),
– the *satisfaction* β with the received service (s_β),
– the acknowledged *debt* of β as a consequence of his satisfaction (t_β), and
– the *valorisation* of α by β (v_α).

In the future, α can make use of this credit v_α and ask β to perform a service on its behalf; that is, α can realise its virtual values in actual values, as illustrated in Figure 2. Nothing forces β to accept the request, but it returns the favour to α not only because it feels gratitude and recognises its debt t_β, but also because it is a way to persevere with interactions with α when these are successful (otherwise, α will devalorise β as

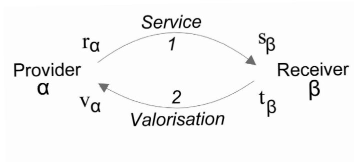

Fig. 1. Exchange of values between interacting individuals α and β: acquisition of virtual values

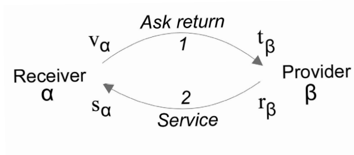

Fig. 2. Exchange of values between interacting individuals α and β: realisation of virtual values

ungrateful and will not interact with it again). On the other hand, if β's expectations were not fulfilled by α, β may not reciprocate the action since it is not interested in continuing the relationship with α.

If β agrees to perform an action for α in return, this action is an actual renouncement for β (r_β) since it requires investment of time and effort, and a satisfaction for α (s_α). After the realisation of virtual values, the exchange is *complete*: α provided a service to β and, later, β provided a service to α.

2.2 Dynamics of Exchange Values

The dynamics of exchange values, or the way they are accumulated and spent by individuals, is based on the premise that in every interaction in which an action is provided or received, something is lost and something is gained: the provider α *renounces* its time and resources for providing an action (r_α) but gains a credit as a result of his *valorisation* (v_α), and the receiver β gains *satisfaction* with the benefits of the received action (s_β) but loses in acquiring a *debt* with the provider in return (t_β). What follows is that, if the gains and losses of participants according to their individual evaluations are equivalent (if $r_\alpha = s_\beta$, $s_\beta = t_\beta$, $t_\beta = v_\alpha$ and, consequently, $v_\alpha = r_\alpha$), the interaction between them is said to be in *equilibrium* regarding the acquisition of virtual values (provided that the participants' evaluations are estimated over the same scale of values). We represent the equilibrium situation for interactions in which virtual values are acquired in the form of the equation below:

$$(r_\alpha = s_\beta) \wedge (s_\beta = t_\beta) \wedge (t_\beta = v_\alpha) \Rightarrow (v_\alpha = r_\alpha) \quad (1)$$

After the values of credit and debt are acquired and α requests an action from β in return, the exchange values change as follows: α loses its credit (v_α) but gains satisfaction with the received action (s_α), while β gains by paying its debt (t_β) but loses effort for performing the action (r_β). Again, if losses and gains are equivalent, the interaction between α and β is said to be in equilibrium regarding the realisation of values. The equilibrium situation for interactions in which virtual values are realised is represented in the form of the equation below:

$$(v_\alpha = t_\beta) \wedge (t_\beta = r_\beta) \wedge (r_\beta = s_\alpha) \Rightarrow (s_\alpha = v_\alpha) \tag{2}$$

In summary, the dynamics of exchange values is based on what is gain and what is lost in every interaction. However, since individual evaluations are subjective and individual interests can influence the valorisation of actions, other situations are possible in addition to the *general equilibrium* of exchange values. These situations are described in the next section.

3 Analysing Cooperative Interactions

3.1 Scenario

In the bioinformatics domain, different types of information are available, such as genome sequences and annotations, protein sequences, structure and interaction maps, as well as metabolic pathways and phylogenetic trees [10]. Each of these information sources can be associated with a particular organism, tissue, cell, disease, etc, resulting in a large variety of information that can be exchanged in a cooperative manner. Thus, since all types of information in bioinformatics is related in some way, the cooperation between users to exchange this information may facilitate some of the tasks they perform, such as validation of new discoveries against existing data, or even help them to gain insights on what path to follow towards a new discovery.

Bioinformatics data and services vary in terms of quality or performance, for example, genome annotations made by one group can be more accurate and described in more detail, while annotations from another group can be poor in detail and ambiguous. The same can happen for services: some may have a better response time and more precise algorithms, while others may take longer to perform or give poor quality results.

Despite sharing similar services or information, users can have very different goals and preferences towards service quality. Developing an automated cooperative system with all this variety is a challenge, mainly because it has consequences for the way in which agents evaluate the services and information they receive from others, and the way providers evaluate the service they provide.

3.2 Cooperative Situations

In our approach, exchange values represent the gains and losses of agents in each interaction. Thus, we define the success of an interaction in terms of the *balance* of the agent's values: a *successful interaction* for a generic agent is when its losses and gains are equivalent, or when its gain is greater than its loss, and an *unsuccessful interaction* is when the agent's gain is less than its loss.

It is important to note that, based on the outcome of interactions, an agent can decide whether it should continue with the cooperation (in cases of successful interactions), drop the cooperation, or review the quality of its services to adjust its evaluation to that of the partner in an attempt to improve its valorisation (in cases of unsuccessful interactions).

When two agents α and β interact, and α provides a service to β, their individual evaluations of the service involved can coincide (as in Equations 1 and 2) or can be different, since individual evaluations are subjective and individual interests such as goals and preferences can influence the evaluation of services. In the following we analyse how different individual perspectives over service evaluations can lead to successful or unsuccessful interactions, and discuss possible causes of each situation and how they affect the cooperation between agents.

Successful Interactions. A successful interaction is the one in which the balance of the exchange values of an agent is in equilibrium (its gains and losses coincide) or is positive (its gains are greater than losses). Examples of successful interactions are presented below.

1. *The provider's effort and the receiver's satisfaction coincide*. This is represented by the relation $r_\alpha = s_\beta \Rightarrow r_\alpha = v_\alpha$, which means their values are in equilibrium and the provider is valorised by the receiver in correspondence to its effort. Possible causes for the equilibrium of exchange values are that:
   (a) the service result was as expected for both provider α and requester β; or
   (b) α's and β's expectations are at the same level.
2. *The provider's effort is less than the receiver's satisfaction*. This is represented by the relation $r_\alpha < s_\beta \Rightarrow r_\alpha < v_\alpha$, which indicates that the receiver valorises the provider more than the latter valorises its efforts. This means that the interaction was successful for the provider, since it did not have to make much effort to satisfy the other agent. Possible causes for the high-valorisation of α's service are that:
   (a) α is over-skilled and can perform the service with a small effort;
   (b) β has a low expectation for the service result and thus over-values even a fairly poor service; or
   (c) even though the service is trivial, it helped β to achieve an important goal and thus had a higher value for the latter.
3. *The receiver valorises the provider more than its real satisfaction.* This is represented by the relation $s_\beta < t_\beta \Rightarrow r_\alpha < v_\alpha$, which means the interaction is successful and the provider is valorised more than it valorises its efforts. The cause for the high-valorisation of the provider can be the following:
   (a) the requester β wants to persuade provider α to continue the cooperation by valorising the service higher than its real satisfaction either because β is in a lower social position than α, α is the only provider for the service that β needs, or α is a busy provider and may have to limit its provision in the future.
4. *The debt is recognised by β and paid by spending similar effort*. This is represented by the relation $t_\beta = r_\beta \Rightarrow v_\alpha = s_\beta$, which means α's satisfaction with the service provided by β was equivalent to the credit it charged. This interaction is in equilibrium regarding the values gained and lost and the cooperation between α and β is therefore successful (since there was reciprocal cooperation).

   The causes for the equilibrium can be that β is compelled to return the favour to α not only because it recognises its debt t_β, but also because reciprocating is a way to persevere with cooperation with α when this cooperation is successful.

Otherwise, α will devalorise β as a bad cooperative partner and will not interact with it again.

In all situations described above — when there is either high-evaluation of β by α , equilibrium of α and β's evaluations, or reciprocation by spending similar effort — agents are successful in their interaction and, as a consequence, the cooperation between them tends to continue.

Unsuccessful Interactions. An unsuccessful interaction is the one in which the balance of the exchange values of an agent is negative (its losses are greater than gains). Examples of unsuccessful interactions are described as follows.

1. *The provider's effort is greater than the receiver's satisfaction*. This is represented by the relation $r_\alpha > s_\beta \Rightarrow r_\alpha > v_\alpha$, which indicates that if the satisfaction of β is less than the effort of α, the valorisation of α is less than its effort. This means that the interaction was not beneficial for α, since its *valorisation* did not compensate for its efforts. The disequilibrium of evaluations happened because β was not satisfied with the service it received, for which possible causes are:

   (a) that α's service had poor quality and did not meet β's expectations; or
   (b) that β had very high expectations and thus under-valued even a good service.

   In the bioinformatics example, this situation can occur when β requests from α the annotation data related to a specific tissue, but the data is only partially annotated and with poor descriptions. In this case, β's evaluation of the received service is less than expected. The consequence of this situation can be that:

   - α can either continue cooperating with β and improve the quality of its service to get a better evaluation;
   - α can cease its cooperation with β if the latter is not being fair in its valorisation;
   - β can continue cooperating with α and lower its expectations on service evaluation; or
   - β can cease the cooperation with α if the latter is providing a poor service and search for another partner to cooperate with.

2. *The receiver valorises the provider less than its real satisfaction.* This is represented by the relation $s_\beta > t_\beta \Rightarrow r_\alpha > v_\alpha$, which means that even though β was satisfied with the received service, its does not valorise α accordingly and, as a consequence, the latter's efforts are greater than its valorisation (this interaction is unsuccessful for α). The possible causes for the under-valuation of the provider in this case are that:

   (a) β has authority over α and thus does not feel obliged to reciprocate the service;
   (b) β wants to exploit α and to have as few debts as possible; or
   (c) β is busy and does not want to compromise its future time with α.

   In the bioinformatics scenario, this situation can arise, for example, when a new agent provides a service to an existing agent in the collaborative community, and even if the new agent provides a good service, the existing agent valorises it less. The effects of the under-valuation in α and β 's cooperative relation is that:

- α stops the cooperation with β if it thinks the interaction is not beneficial for it;
- α maintains the cooperation despite the under-valorisation if it needs β's service for achieving a goal in the future (either because there is no other agent to provide the service, or because β provides the service with best quality);
- β tends to keep the cooperation since it is beneficial for it (the debt it acquires as a result of α's valorisation is less than its gain from the received service).

When agent α is successful in its action, and achieves $v_\alpha = r_\alpha$or even $v_\alpha > r_\alpha$, α's valorisation constitutes a credit for it. In the future, α can make use of this credit v_α and ask β to perform a service on its behalf. Nothing forces β to reciprocate and accept the request, since there is no formal or legal commitment between them, and thus the following situations are possible.

3 *The debt is recognised by β but paid by spending less effort than the worth of the credit*. This is represented by the relation $t_\beta > r_\beta \Rightarrow v_\alpha > s_\alpha$, which means that α's satisfaction with the service provided by β was less than the credit it charged (the service α provided to β in the past was highly valued than the service it received from β in return). Thus, the interaction was not beneficial for α since it lost more than it gained. A possible cause for the disequilibrium is that β wants to exploit α by asking more than it is willing to return, or because β does not have enough skills to perform a good service. As a consequence, α tends to stop the cooperation with β if it has other possible partners to cooperate with.

4 *The debt is not recognised by β*. In this case, the interaction does not take place. Possible causes are that β just wants to take advantage of interactions and does not share or provide services to others, or that β's expectations were not fulfilled by α in the previous interactions, and thus β may not reciprocate the action since it is not interested in continuing interacting with α. It is clear in this case that there is no more cooperation between both agents.

In summary, when agents have different perspectives over service evaluations (such as different levels of expectation towards service results), or are influenced by personal interests (like providing a service with less effort in return for another received previously to reduce losses), interactions may not always be successful in terms of gains and losses of exchange values. Therefore, in a resource constrained environment in which agents take autonomous decisions about interactions, it is important that they avoid repeating unsuccessful interactions and try to maintain successful ones instead. To show this approach to decision-making over interactions, we set an experimental testbed with a similar scenario, which is described in the next sections.

4 Experiment

4.1 Scenario

Cooperative applications in bioinformatics are characterised by different types of available services, which can vary in terms of performance and quality (since providers have different skills to perform services). Because service providers are resource-bounded, they must limit the number of services they provide. In addition, service users have different perspectives over service quality.

In this context, agents cooperate by providing services to and requesting services from each other. Since services are free, providers receive from requesters a *valorisation*, based on the requesters evaluation of the received service, as an incentive for cooperation that the former uses as credit for asking something in return in the future. However, since agents have different perspectives of the same service — providers have different skills and users have different preferences — it may be that an agent receives less than expected from a cooperation (so that its interactions are unsuccessful, as described in Section 3.2).

We want to determine whether the analysis of cooperative situations can reduce the number of unsuccessful interactions, but without decreasing the number of interactions between participants, which would be expected if agents increase the restrictions on desirable interactions.

In seeking to determine that, we observe the *difference between the number of unsuccessful and successful interactions* for agents using two different decision-making strategies: *simple reciprocation*, and *analysis of cooperative situations*. Our hypothesis is that by analysing cooperative situations the agents can improve the number of successful interactions in the society, without significatively reducing the *number of achieved interactions*.

4.2 Strategies

Both strategies for selecting among alternative interaction partners are based on the analysis of exchange values. The difference is that, in the simple reciprocation, agents take into account only the credits and debts they have with other agents, as described in previous work [9], while in the analysis of cooperative situations, agents also take into account the balance of each interaction in which they participate to decide whether to cooperate and to which partner to send a request.

Consider the set of possible providers as $P = \{p_1, .., p_n\}$ for a service sr_i, and the set of received requests as $Q = \{q_1, .., q_m\}$. For *simple reciprocation*, the decision-making strategy is as follows.

1. For an agent α requesting a service:
   (a) remove from P agents that did not pay their debts in the past;
   (b) let P_o be an ordered sequence of the elements in P according to the credits α has with each agent in P, with higher credits first (since providers with debt are more likely to cooperate);
   (c) send a request to the first agent in P_o;
   (d) while the request is refused, send the request to the next agent in P_o.
2. For an agent β providing a service:
   (a) remove from Q requests from agents that did not pay their debts previously, since they are likely not to reciprocate in the future;

(b) let Q_o be an ordered sequence of the elements in Q according to the debts β has with each agent in Q, with higher debts first (so that β can reciprocate for services received previously);
(c) accept requests in order, until reaching maximum capacity, and refuse the remaining requests.

In the *analysis of cooperative situations* strategy, decision-making complements items 1(b) and 2(b) by avoiding repetition of unsuccessful interactions, as described in Section 3.2. This strategy is presented below:

1. For agent α:
   (b) reorder sequence P_o by moving the candidate providers with which α had unsuccessful interactions previously to the end of the sequence, such as cases in which α's satisfaction is smaller than its debt, $s_\alpha < t_\alpha$ (indicating that α received a service under its expectations).
2. For agent β:
   (b) deny requests from agents from which β received an unfair evaluation in previous interactions causing its renouncement to be greater than its credit, $r_\beta > v_\beta$; and reorder the requesters in Q_o with equivalent debt with those with which β had higher satisfaction from previously received services first.

Thus, by analysing the balance of their exchange values in previous interactions, requesters try to avoid sending requests to agents that provided a poor service in the past (resulting in $s_\alpha < t_\alpha$), and providers try to avoid continue cooperating with agents that did not reciprocate in the past and also those that are under-evaluating the service they are providing (resulting in $r_\beta > v_\beta$).

4.3 Simulation Configuration

To simulate cooperative interactions between agents with different perspectives over service evaluations, we require: first, that agents have different skills to provide services, and second, that agents have different expectations towards received services. Thus, every provider has an associated skill from the set $k = \{0.5, 1, 2\}$, where 0.5 means the provider has *low* skills; 1, *medium* skills; and 2, *high* skills. Also, every service an agent needs is associated with an expected quality of result $exp = \{0.5, 1, 2\}$, where 0.5 means the requester has *low* expectations; 1, *medium* expectations; and 2, *high* expectations.

Each service sr_i is associated with an execution effort eff_{sr_i} (for the purpose of comparison, we assume that all agents invest the same effort to perform the same service), and the relation between the effort to perform the service and the skill of the provider determines the service result, represented by res_{sr_i}. The service result is informed to the requester after execution and calculated as follows:

$$res_{sr_i} = eff_{sr_i} \times k_{sr_i}$$

A simulation in our experiment consists of a number of iterations. In each iteration, all agents perform an action: they can request a service or provide a service. They use their decision-making strategy to decide whether to provide a service and to whom to send a request. An *interaction* occurs when the request is accepted and a service is received. After every interaction, the agents determine their individual evaluation of the service. The requester (α) evaluates service sr_i as follows:

$$Eval_{\alpha}(sr_i) = \frac{res_{sr_i}}{exp_{sr_i}}$$

The provider (β) evaluates the performed service sr_i as follows:

$$Eval_{\beta}(sr_i) = \frac{eff_{sr_i}}{k_{sr_i}}$$

Based on these evaluations, α and β determine their exchange values, as described in the next section.

In real applications, evaluations might be determined in a straightforward way through an objective evaluation process. More specifically, the provider's evaluation ($Eval_{\beta}$) could be based on a *cost* function (which would include an effort measure related, for example, to processing time or memory usage), and the receiver's evaluation ($Eval_{\alpha}$) based on a *utility* function (which would consider the actual service result and an expected result).

In every simulation run, we fix the number of total requests the agents can send, and record the *total number of interactions* that occurred, and the *number of unsuccessful interactions*.

4.4 Determining Exchange Values

After agent α receives a service sr_i from another agent β, α determines its satisfaction (s) and debt (t) values. The satisfaction is determined by α based on its evaluation of the received service. We assume that the debt acknowledged by α is always 1 to represent the situations in which the result of the service provided by β achieved α's expectations ($\frac{res_{sr_i}}{exp_{sr_i}} = 1$). Therefore, α's satisfaction and debt values are:

- $V_{\alpha,\beta}(s) = Eval_{\alpha}(sr_i)$
- $V_{\alpha,\beta}(t) = 1$

Regarding the balance of α's exchange values, if α's evaluation of the service is less than expected, $V_{\alpha,\beta}(s) < V_{\alpha,\beta}(t)$, the interaction is considered unsuccessful for α. After determining its exchange values, α communicates its satisfaction to β to represent its valorisation of the latter.

After agent β provides a service to α, it determines its renouncement (r) and credit (v) values. The renouncement of β to perform sr_i is determined according to its evaluation of the service, and the valorisation of β by α is stored by β as a credit. Thus, β's renouncement and credit values are:

- $V_{\beta,\alpha}(r) = Eval_{\beta}(sr_i)$
- $V_{\beta,\alpha}(v) = V_{\alpha,\beta}(s)$

If β's renouncement is greater than the credit it gained, $V_{\beta,\alpha}(r) > V_{\beta,\alpha}(v)$, the interaction is considered unsuccessful for β.

Table 1. Population variation for Experiment 1

characteristic	percent in population
providers skills	low(20%), medium(60%), high(20%)
requester expectations	low(25%), medium(60%), high(15%)

Table 2. Results

Strategy	(U) Interactions (%)	Total interactions
SR	26.4	769
ACS	6.9	756

4.5 Results

All simulations used a society with 30 agents, which provide and request services from a set of 4 available services. We fixed the provider capacity at a maximum of 2 simultaneous services, and the total number of sent requests at 800. Services that are needed and provided are distributed equality over the agent population, so that no provider is busier than any other. The variation of the population of agents in terms of skills and expectations follows the proportions in Table 1.

We performed two different experiments: in the first one, we vary the agents' characteristics (skills and expectations) but keep the proportions in Table 1; and in the second one, we keep the proportions for requester expectations and vary the proportion of low-skilled providers.

In the first experiment, we performed 50 simulations for both decision-making strategies, varying the agents' characteristics but keeping the proportions in Table 1, to provide variation to the agent population from one simulation to another, guaranteeing that both decision-making strategies were simulated over the same agent population sample. In each simulation, we recorded the *total number of interactions* that occurred between agents, and the *number of unsuccessful interactions* (U) for each strategy. The average results for the 50 simulations are shown in Table 2.

According to the results, when using the simple reciprocation strategy (SR), of the 800 sent requests, an average of 769 actually resulted in an interaction and, from this total, 26% of the interactions were unsuccessful. However, when using the analysis of cooperative situations strategy (ACS), an average of 756 requests actually resulted in an interaction and, of this total, *only* 6.9% of the interactions were unsuccessful.

We observe that by analysing the balance of the exchange values when deciding over cooperative interactions, agents can significantly reduce the number of unsuccessful interactions. The decrease in the number of total interactions is justified by the increase in the number of constraints that agents make on desirable interactions, when they analyse the balance of exchange values. However, interactions were reduced by only 1.6% in contrast with the significant reduction in the number of unsuccessful interactions.

In the second experiment, we test the behaviour of both strategies when we increase the proportion of low-skilled providers in the agent population. To that end, we keep the proportions for requester expectations in Table 1, fix the proportion of high-skilled providers at 10%, and vary the proportion of low-skilled providers from 5% to 40%

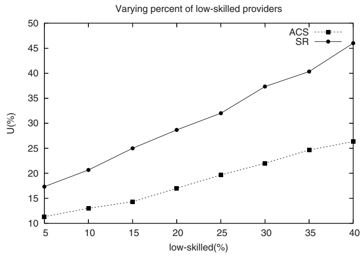

Fig. 3. Percentage of unsuccessful interactions for different proportions of low-skilled providers in the population

(the proportion of medium-skilled providers decreases proportionately with the increase of low-skilled providers). The results are shown in Figure 3.

According to the results, when there are more providers performing poor quality services (40%), the number of unsuccessful interactions can reach a critical amount (almost 50%) if agents rely on simple reciprocation. However, by analysing cooperative situations agents can reduce this amount to less than 30%.

We can see from the results that, when there is a significant number of poor quality services, it becomes more critical to identify unsuccessful interactions, so that successful ones can be maintained instead.

5 Related Work

Related work on exchange values [11] has proposed a centralised approach to coordinate the balance of interactions in terms of exchange values, in which a social equilibrium supervisor uses exchange values and social rules to coordinate exchanges between agents. An extension of this approach to address coordination of social exchanges in spite of different personality traits of agents is presented in [12].

The valorisation of agents that act cooperatively to motivate cooperation in a multi-agent system is used in [13] in the form of *brownie points*. Agents that perform tasks for the group gain brownie points and agents that defect from group tasks lose brownie points. Although it is applied to cooperative group work and our model focuses on one-to-one cooperations, the brownie point approach is similar to the valorisation through exchange values proposed in this paper in the sense that the credits earned by agents are a motivation for providing services with no monetary return, which is the case of the cooperative applications we focus on this paper. However, brownie points represent an agent's self-valorisation (the agent rewards itself for cooperating in a team) and not the valorisation an agent receives from others in retribution for the provided service.

The analysis of reciprocal interactions as a criterion for deciding whether to cooperate with other agents is proposed by [14,15] in the form of an expected utility-based decision-making. According to this approach, agents agree to cooperate if the cost of helping the requester agent is smaller than the expected benefit of receiving help from the requester and other agents in the future. By considering expected future help in the providers's utility function, agents are motivated to cooperate with each other since the probability of receiving resources increases with the number of times they helped others. Similar to this approach, the decision-making based on exchange values also provides a motivation for agents to cooperate with others and to maintain reciprocal relations as a guarantee for future benefits. However, the utility based decision-making does not consider the success of the cooperative relations in terms of the balance between the effort spent by the provider and the satisfaction of the receiver, which is important to consider if the environment has agents with different preferences and perspectives, or even agents that reciprocate but by providing low quality services.

Alternatively, cooperation and reciprocity supported by norms and organisations is proposed by [16], in which knowledge sharing between agents is supported by social contracts and rules inside a virtual organisation framework. Agents that request information from a provider have to offer a service in return that is of interest of the provider by means of a formal contract, which is monitored by an organisation-related agent that checks whether the agreement is followed. Although this approach uses reciprocity to motivate cooperation, it differs from our approach in the way reciprocity is achieved. Instead of using formal social contracts to enforce reciprocal relations, the exchange values approach relies on the informal commitments represented by virtual credits and debts which influence the chance of future interactions. Also, their approach does not consider the evaluation of services neither the influence of the outcome of interactions for the maintenance of cooperative relations.

6 Conclusion and Future Work

We have presented a computational approach for modelling non-economic, autonomous cooperative interactions based on the theory of *exchange values*. Here, exchange values *motivate* agent interactions and their maintenance through a system of credits and debts. Moreover, the credits and debts acquired by agents during interactions are based on their individual evaluation of the service being performed or received, and the balance of these evaluations indicates whether an interaction was successful for each agent involved. We argue that, by analysing the *balance* of exchange values in past interactions, autonomous agents can identify situations in which the cooperation with other agents is unsuccessful and decide whether to maintain this cooperation.

We presented a scenario in which agents participate in a cooperative application in the bioinformatics domain by requesting and providing services to each other with bounded resources, and with different perspectives over service evaluation. Agents have two different cooperative strategies, one based on simple reciprocation of credits and debts, and another based on the analysis of cooperative situations. We use an experimental testbed to compare the number of unsuccessful interactions that occur when agents use both approaches, and show that agents can *reduce the number of unsuccessful interactions* by analysing cooperative situations and still keep a high number of achieved interactions.

Future work aims to combine the analysis of exchange values and their balance with the analysis of service dependencies between agents to improve the decision-making strategy, as proposed in [8]. We also aim at developing a qualitative representation for exchange values to which quantitative evaluations are mapped, to facilitate their application to different scenarios (with different evaluation scales).

References

1. Burt, R.: The network structure of social capital. In: Research in Organizational Behavior, vol. 22, JAI press, Greenwich, CT (2000)
2. Stein, L.: Creating a bioinformatics nation. Nature 417, 119–120 (2002)
3. Gao, H.T., Hayes, J.H., Cai, H.: Integrating biological research through web services. IEEE Computer 38, 26–31 (2005)
4. Goble, C., Wroe, C., Stevens, R.: the myGrid consortium: The mygrid project: services architecture and demonstrator. In: UK e-Science All Hands Meeting 2003, pp. 595–603 (2003)
5. Overbeek, R., Disz, T., Stevens, R.: The SEED: A peer-to-peer environment for genome annotation. Communications of the ACM 47(11), 47–50 (2004)
6. Foster, I.: Service-oriented science. Science 308(5723), 814–817 (2005)
7. Piaget, J.: Sociological Studies. Routlege, London (1973)
8. Rodrigues, M.R., da Rocha Costa, A.C., Bordini, R.H.: A system of exchange values to support social interactions in artificial societies. In: Second International Joint Conference on Autonomous Agents and Multiagent Systems, Melbourne, pp. 81–88 (2003)
9. Rodrigues, M.R., Luck, M.: Analysing partner selection through exchange values. In: Sichman, J.S., Antunes, L. (eds.) MABS 2005. LNCS (LNAI), vol. 3891, pp. 24–40. Springer, Heidelberg (2006)
10. Campbell, A.M., Heyer, L.J.: Discovering Genomics Proteomics and Bioinformatics. Benjamin Cummings, San Francisco, CA (2002)
11. Dimuro, G.P., Costa, A.C.R.: Qualitative Markov decision processes and the coordination of social exchanges in multi-agent systems. In: Gmytrasiewicz, P., Parsons, S. (eds.) Workshop on Game Theoretic and Decision Theoretic Agents at IJCAI Conference, Edinburgh (2005)
12. Dimuro, G.P., Costa, A.C.R., Gonąlves, L.V., Hübner, A.: Centralized Regulation of Social Exchanges between Personality-based Agents. In: Boissier, O., Padget, J., Dignum, V., Lindemann, G., Matson, E., Ossowski, S., Sichman, J.S., Vázquez-Salceda, J. (eds.) Coordination, Organizations, Institutions, and Norms in Multi-Agent Systems, Springer-Verlag, Berlin (2006)
13. Glass, A., Grosz, B.: Socially conscious decision-making. In: Proceedings of the fourth international conference on Autonomous agents, pp. 217–224. ACM Press, New York (2000)
14. Sen, S., Dutta, P.S., Saha, S.: Emergence and stability of collaborations among rational agents. In: Klusch, M., Omicini, A., Ossowski, S., Laamanen, H. (eds.) CIA 2003. LNCS (LNAI), vol. 2782, pp. 192–205. Springer, Heidelberg (2003)
15. Banerjee, D., Saha, S., Dasgupta, P., Sen, S.: Reciprocal resource sharing in p2p environments. In: Fourth International Joint Conference on Autonomous Agents and Multiagent Systems, Utrecht, Netherlands (2005)
16. Dignum, V., Dignum, F.: Knowledge market: Agent-mediated knowledge sharing. In: Third International/Central and Eastern European Conference on Multi-Agent Systems, Prague (2003)