UNIVERSITY OF SOUTHAMPTON

A Direct Reputation Model for Virtual Organisations Formation

by

Arturo Avila Rosas

A thesis submitted in partial fulfillment for the
degree of Master of Philosophy

in the
Faculty of Engineering, Science and Mathematics
School of Electronics and Computer Science

March 2005

This thesis represents the final form of the thesis as deposit in the University after examination.

UNIVERSITY OF SOUTHAMPTON

ABSTRACT

FACULTY OF ENGINEERING, SCIENCE AND MATHEMATICS
SCHOOL OF ELECTRONICS AND COMPUTER SCIENCE

Master of Philosophy

by Arturo Avila Rosas

Recently, a large number of new collaborative, networked organisations have emerged, having as motivator the explosive progress in computer networks and communication systems, but also as a reaction to market pressures that demand customised, high quality products and services at lower costs and, at the same time, shorter production and marketing times. Promising grater flexibility, resource optimisation and responsiveness in *competitive open environments*, Virtual Organisations are presented as an example of this trend that has pervaded not only business domains but other areas such as e-science. What distinguishes Virtual Organisations from other forms of organisation is the full mutual dependence of their members to achieve their goal and therefore the need for cooperation. However, open environments, where Virtual Organisations are embedded, involve organisations and individuals that do not necessarily share the same objectives and interests, which they might not know in advance, and where they might not trust each other, but should work together and help each other to achieve a common goal. One of the key omissions in the computational representation of Virtual Organisations relates to the need to take into account more subjective facets like the *reputation* of the individuals, that helps to cope with the heterogeneity, the autonomy and the diversity of interests among their members. This thesis addresses the introduction of reputation into Virtual Organisations, by providing a reputation model based on the adaptive evaluation of direct experiences to identify trustworthy members of a Virtual Organisation.

This thesis makes three main contributions. First, it provides the means for designers of VOs to identify situations where reputation is required. This is needed since members of a Virtual Organisation, which are also self-interested, must cooperate with each other, and reputation represents a way to identify the possible deceivers and the more attractive partners in the environment. Second, it provides a reputation model for Virtual Organisation partners that does not require individuals to interact extensively to detect possible abusive or honest partners. Third, it provides a comparative analysis of current reputation systems used in e-commerce and Multi-Agent Systems, and of the new model proposed.

Contents

List of Figures

List of Tables

Nomenclature

w	time window
t	discrete time step
ϵ	dimension
i	target or service provider
j	source or service consumer
imp	impression
Q	set of enabling qualities
K	context
QoS	Quality of Service
w_t	importance of impression
T	final time step of a period
DR	Direct Reputation
CT	Competence Trust
GT	Goodwill Trust
DT	Degree of Trust
$PerR$	Perceived Risk
RR	Relational Risk
PR	Performance Risk
a_t	action at t
$Q(\cdot,\cdot)$	value of an action function
r_i	ith reward
$\alpha(\cdot,\cdot)$	learning rate function
$\beta(\cdot,\cdot)$	similarity function
$E(\cdot,\cdot)$	Expectation function
SR	Social Reputation
wr	reliance parameter
$ReDr$	Reliability of Direct Reputation
$ReSr$	Reliability of Social Reputation
R	Risk

Acknowledgements

I would like to begin by thanking the two persons without whom it would not have been possible for me to submit this thesis: my supervisor Michael Luck, and my wife Rosy. I would also like to thank Melani Ashleigh whose comments on an early draft of this thesis helped me to better understand the subject of my research. I also thank all the researchers who shared with me their publications.

I thank the Instituto Mexicano del Petróleo (México) for sponsoring and supporting my academic career.

Many thanks to all persons who have helped me and whose friendship is very important to me: Axel, Romeo, Fabiola, Dionisio, Toño, Jorge and especially to Cora and Alejandro. Thanks to my schoolmates at the IAM Group and my mates at the Mexican Society of the University of Southampton.

Thanks to my family for their love, support and encouragement.

To my wife and my parents.

Chapter 1

Introduction

1.1 Motivation for Research

After seeing a video mobile phone advertisement, the executives in a company dedicated to design and sell equipment needed to build Internet technology-based networks decide project-manage a new product that exploits the enormous advantages of the video mobile media in the Internet. Executives are conscious of the company cannot afford to be second-rated at any aspect of their business in such a highly competitive business environment. On the other hand, they do not realistically expect to be first-rate at everything either. To retain their lead in that important and growing market, they decide to form strategic alliances with other companies. This undertaking requires to outsource different production processes including the mobile devices manufacture and Internet Services. To succeed in this new venture, one of the executives enters the web site of Fast Virtual Manufacturing *(www.virtual-manufacture.org), a virtual infrastructure that brings together several companies that are market leaders within their core capabilities, and assists with the dynamic creation of temporary inter-organisational* Virtual Organisations. Virtual Manufacturing *consists of a network of companies that facilitates project communication through computer technology. Through* Virtual Manufacturing *the executive is able to establish clear roles and responsibilities for all partners, including a framework for participating virtual collaborations. After a short electronic negotiation between potential partners, the companies that fulfil the new product requirements best are chosen and receive electronics contracts.*

The above example illustrates the immense potential for computer-supported assembly and operation of *Virtual Organisations*. The idea of the *Virtual Organisation* was popularised by Davidow and Malone in the *The Virtual Corporation* [25], and became one of the buzzwords of the 1990's. The growing presence of global marketplaces have motivated companies to continuously seek and adopt new ways to address the competitive pressure, and Virtual Organisations appeared as a new form of organisation to relieve

that pressure. Urgent needs such as shortening development and manufacturing cycles, reducing time to market and operational cost, increasing customer satisfaction, operating on a global scale and reach, and rapidly adapting to new market changes have led companies to automation, collaboration and distribution [54, 6] through the formation of Virtual Organisations. An interesting aspect is the increasing number of companies involved with cooperative strategies, sometimes even with direct competitors.

The Internet has opened the opportunity to different size companies, in many cases geographically dispersed, to participate and take advantage of Virtual Organisations. In this way, Virtual Organisations hold great promises for bringing together large numbers of companies and potential consumers across wide geographic regions on a scale that is unprecedented in the physical world.

Given the above, many companies see important advantages in forming Virtual Organisations. Among these advantages is the fact that companies have the potential to dramatically improve the efficiency of trading through the reduction of search and transaction costs. Moreover, companies can offer complex services or customised goods by cooperating in different product development phases and sharing critical business processes, resources, core competences, skills and know-how with each other [43, 46]. At the same time, they obtain less tangible benefits such as reduction of the risk of investments (or allowing someone else to assume the risk), or increase in market size or market share [5]. Thanks to this new organisational model, suppliers, manufactures and retailers are enable to act as a single Virtual Organisation; these integrated companies create a more competitive whole.

From this place, the field is ripe for the introduction of information systems that discover windows of opportunity, select the *best* proposals, automate logistics planning and facilitate the cooperation among multiple entities (companies or individuals) such as manufactures, part suppliers, shippers, and specialised contractors.

Traditionally, deciding what to outsource and to whom, ensuring that the task are done in the proper sequence and that the final product is ready within the time constraints, has been the job of a human decision-maker. In many current alliances, the decision-maker also has to keep track of any problem derived from providing of a service (delays, bad quality, overprice) that could jeopardise the completion of the tasks, and renegotiate with the suppliers or others as needed. Although Information Technologies (IT) have made possible, at certain extent, the automation of such activities, the *formation* of Virtual Organisations is still dependent on human decisions [8, 17].

Virtual Organisations' managers are reluctant to delegate completely the whole formation process. Decision involved in the formation of Virtual Organisations, such as the selection the most suitable set of partners and the allocation of tasks, are made by the human managers basically because they do not necessarily *trust* in the partners suggested to form the alliance. This lack of trust is understandable because of the nature

of the potential partners who are characterised by *self-interested* rather than *cooperative* entities and therefore it is not advisable to make naive assumptions about their *benevolent* character.

In this work we consider how current technologies used to manage Virtual Organisations can be enriched to consider the aspect of trust.

1.2 Basic Concepts in Virtual Organisations

A Virtual Organisation (VO) is defined as a *network of organizations or individuals that share skills and resources, and is conceived with the objective of delivering a product or service in response to a customer necessity, and whose cooperation is supported by computer networks* (adapted from [51]). In a VO partners are dispersed geographically and do not work under the same formal organisation [50]. In the e-science context, VO are also seen as the fundamental organising principle of the *Grid* [32], whose main objective is support of large scale scientific computing.

The general objective of Virtual Organisations (VOs) is obviously to satisfy a necessity that cannot be satisfied by a single organisation (company or individual) or optimise the satisfaction of a necessity.

1.2.1 Categories of VOs

The structure of VO and their number of partners might be either static, or dynamic. Moreover, the VO can last for only one service provision (the time needed to satisfy its purpose) up to some months or years. The above mentioned features (structure, number of partners and relationship duration) define two categories of VOs:

- Static Virtual Organisations (SVOs), and
- Dynamic Virtual Organisations (DVOs)

SVOs describe organisations that are built over time and last several years with the same relationships. Partners are linked together in a static and fixed way, that is, the shared competences and resources are tightly integrated. The relationships among the partners (for example relations of cooperation, collaboration, competence, business) are pre-defined, fixed and well integrated and customised among the partners. The network of organisations is fixed and pre-determinated and thus, the structure of the VO is static.

In contrast, DVOs link partners dynamically, on demand, and according to the opportunities detected. Due to the non-existence of fixed relationships, DVOs might change continuously to take full advantage of the open, global opportunities offered by the Internet and the global economy.

1.2.2 The life cycle of a VO

Different stages of development can be identified in a VO since a business opportunity is detected up to the achievement of all its goals. In order to illustrate the stages of a VO, consider the example above mentioned.

1. Executives of the company detect a business niche to produce a new electronic device. That new business opportunity is the main goal of the new VO. Executives must decide which are the attributes of the product or service to be delivered. Relevant attributes can be cost, quality, price, delivery time, as well as constraints on the participants in the organisation. In our example, executives may decide the characteristics of the electronic device, markets, the approximate price in the market, quantity of sales, functionality of the device, quality of the offered service, commercialisation, etc.

2. Then, executives divide properly the main goal into tasks to be assigned to its potential partners. This is a highly dynamic activity because tasks must fit the main goal and at the same time exploit rationally the resources and abilities of the potential partners.

3. Now, it is time to select the potential partners to form the VO and allocate the required tasks among them. Note that the activities of distribution of tasks and decomposition of the main goal cannot be assumed to be independent. On the one hand, the given allocation options (list of potential partners) restrict the reasonable decompositions; on the other hand, decomposition decisions clearly restrict the possible allocations. In our example, executives may decide to decompose the task of providing the electronic device in many activities, allowing different companies to manufacture and assemble small components or, leave to one company the whole process of manufacture and assembly.

4. The next stage is the *negotiation* of the terms for providing each task. Partners must come to a mutually acceptable course of action to accomplish their particular tasks; that is, to come to an agreement on the *conditions* such as the delivery time, the price, the quality, etc. The process by which such agreements are made is necessarily one of negotiation; this process is based on he exchange of proposals and counter proposals and it ends in the acceptance or the rejection of a deal.

5. After selecting partners and distributing tasks, executives compose the tasks to a whole process that reflects the total added value of the VO. In our example, executives may find one or more potential ways of offer the same service, they must decide which is the best option based on different criteria such as price, quality, etc. And finally, executives send electronic contract to their partners.

6. Partners require to cooperate to achieve their objectives by sharing information, resources, etc., and such cooperation is supported by infrastructures that assure safe communications, privacy and authentication. Nevertheless, partners require a highly level of trust because once the VO starts to operate, organisations may cede direct control of nearly all their operations creating a relation of dependence on their partners. Additionally, when it is necessary to add/or replace a partner, or change roles of partners, executives may alter the structure of the VO.

7. Finally, when each partner finishes its tasks and after the VO has achieved successfully all its goals, executives may decide, on one hand, to continue with the existing partnership making some modification to embark in a new venture, or in the other hand, dissolve the alliances.

From above, three stages in the life-cycle of VOs can be identified namely *formation*, *operation* and *dissolution*. The formation process is important because the success of a VO depends completely on the capabilities of the partners selected. Since all partners have to work cooperatively to reach a common goal, the *selection* of partners must consider not only the requirements of the VO (e.g. price, quality, delivery time, etc.) but also whether the potential partners could inspire and maintain others with confidence.

1.2.3 Formation of VOs

Proponents of VOs claim that unlike the off-line traditional organisations, partners can choose the best possible deal for every service and work with different partners every time. However such affirmations hinge on an important assumptions often unstated: that companies and their partners are *willing to cooperate* to achieve shared objectives in spite of their own individual interest.

In cooperative interactions, individuals are interested in maximising the *social welfare* rather than their own utility. Moreover, cooperation occurs only if two or more individuals have a common goal and depend on each other to achieve it (*mutual dependence*) and the beneficiaries are all those involved in achieving such a goal.

However, partner cooperation cannot be assumed to exist in all members of a VO for many reasons. First, individuals in a VO are likely to represent different organisations that have their own aims and objectives, and both particular and global objectives could not be compatible. Second, because VOs are characterised by short-term relationships that may entail no plans for future association, partners are tempted to abandon the alliance without performing their tasks.*Free-riders* are characterised by this sort of behaviour, and they either may limit or deny their contributions to an alliance or simply behave opportunistically. Third, organisations could refuse to cooperate with other

partners in the alliance who are unknown and whom they do not share past interaction history.

Hence, the formation of VOs is far from trivial. It is a balancing act between potentially conflicting interests and concerns of participating organisations. Some researchers [13, 23, 37] refer a lack of cooperation, the low level individual commitment and the self-interested behaviour of partners as causes for the relatively high rate of failure of VOs.

Moreover, when forming a VO, not only there is no certainty that the partners are actually committed to the common goal but also there are difficulties to identify who is likely to act opportunistically. In such circumstances, the interesting questions is what enables potential partners to garner enough *confidence* in partner cooperation so that they are not overwhelmed by the potential risks in VOs. Das claims that low level of confidence not only discourages the formation of a VO but also leads partners to view each other with suspicion, with a harmful effect on their cooperative relationship, if the VO is formed [23].

Confidence in partner cooperation is defined as an organisation's perceived level of certainty that its partner will pursue mutually compatible interests in the alliance, rather than act *opportunistically*. According to [23], the sense of confidence comes from two distinct sources: *trust* and *control*.

In summary, the formation of VO characterised through the following questions:

1. How do VOs decide who is selected as a partner?
2. How do partners decide how and when to interact with each other?
3. How do VOs promote cooperative behaviour?
4. What kind of mechanism could be used to commit one partner to adopt an agreement?

This thesis aims to answer those questions by considering the *reputation* of the potential partners as the main way to develop confidence in the VO that its partners will pursue mutually compatible interests in the VO rather than act *opportunistically*.

1.3 Information Technologies to set up VOs

The VO concept is so fundamental in today's economy, science and engineering that many Information Technologies have contributed to the formation and management of VOs. These Technologies include EDI (Electronic Data Interchange), novel industry standards for the Internet known as *Web Services*, *Grid* technologies, as well as *Intelligent Agents*.

1.3.1 Electronic Data Interchange EDI

The first attempts to realise inter-organisational systems have been done in the area of Electronic Document Interchange (EDI), where a set of organisations integrate their *business processes* only for electronic commerce purposes. In that case, different organisations have static business relationships, the communication mechanisms used were message passing, the business processes were internal modules or legacy systems maintained by the different partners, the network protocols used were secure virtual private networks, and the format of the EDI messages were proprietary following certain structure and regulations.

Although EDI was a significant progress towards the direction of inter-organisational business automation, the resulted solutions were very expensive and rather closed to be adopted by small firms. Certain problems regarding the standard format of EDI messages, the insecure open transport networks, and the rather restricted context of EDI, i.e. only focus on electronic commerce, made EDI not an attractive solution for VOs.

1.3.2 Internet

Whereas EDI supported electronic business by automating existing processes and enabling electronic document exchange between specific organisations, a number of other technologies approach inter-domain business relationships by integrating applications between business partners anywhere in the Internet. Internet Technologies can be found almost every where, this ubiquity property makes them attractive as a platform for constructing VOs. The Internet allows companies to search actively for potential members in a great number of databases and to transfer a huge amount of information. At the same time Internet is the basis for cooperation. Particularly, a recent set of standards known as *Web Services* has made possible the integration and cooperation of diverse service providers over Internet. In this context, a *service* is a software piece accessible through Internet that provides a specific capability, for example, the ability to book and buy fly tickets, bid in an electronic auction, the locating and operation of laboratory equipment, etc. With Web Services, *service* procurement is simply a matter of finding and selecting the right services and composing them into a solution.

Among these standards there is a common language for transport data called XML (Extensible Markup Language), a language for describing the functional aspects of the service known as WSDL (Web Service Description Language), and a directory of services and providers called UDDI (Universal Description Discovery and Integration) that make services technically and commercially accessible providing the basis finding and advertising such services, companies register their capabilities so others may find them and bind to their services over the Internet.

Although Web Services seem a promising technology for creating and operating VOs, there are some drawbacks that still need to be coped with, such as: security, selecting trusted partners, foster cooperative behaviour, etc.

1.3.3 Grid Technologies

Grid computing has emerged as a new field within the realm of distributed computing that focuses on large-scale resource sharing making distributed computational resources available as easily as electric power is through the electricity distribution grid. Foster [32] defines the Grid problem as controlled and coordinated resource sharing and resource use in dynamic, scalable VOs.

Grid Technologies basically consist of a *Grid Architecture* that is oriented towards the establishment, manage, and exploitation of dynamic, cross-organisational VO sharing relationships among any potential participants. The sharing is not primarily file exchange but rather direct access to computers, software, data, and other resources, as required by a range of collaborative problem-solving an resource-brokering strategies emerging in industry, science, and engineering. Secure sharing is one of the main concerns of the Grid Architecture, sharing is highly controlled, both resource providers and consumers define clearly what is shared, who is allowed to share, and the conditions under which sharing occurs.

The concept of sharing computing resources across VOs raises a large number of issues that traditional distributed computing does not address: how to structure flexible just-in-time sharing relationships; how to structure control over resources, taking care of local or global policies; and how to agree on quality of service, scheduling and co-allocation.

1.3.4 Intelligent Agents

Finally, another emerging technical area for the development of VOs is the concept of intelligent agents. According to Jennings et al. [39], an intelligent agent is a computer system that autonomously performs tasks in an open environment, that is reactive and proactive, and that is able to interact with other agents.

An agent does not need to be continually controlled by an user, but autonomously monitors its environment and takes action as it deems appropriate. This characteristic of agents makes them particularly suited for applications that can be decomposed into stand-alone process, each capable of doing useful things without continuous directions by someone.

The interest in agents arose through the notion of Multi-Agent Systems (MAS) in the early nineties, driven by the Distributed Artificial Intelligence (DAI) research community [39]. In MAS, each agent cooperates with other agents inside the community to

solve a particular complex problem. Therefore, a MAS may be defined as a set of agents that interact with each other and with the environment to solve a particular problem or to provide a service in a coordinated and distributed manner [39]. If these agents are located in different organisations and execute certain tasks on behalf of their users then a Multi-Agent Systems for VOs can be formed.

1.3.5 Discussion

As mentioned above, the concept of VOs poses significant challenges that can and have been addressed by different Information Technologies (ITs) to a certain extent. Indeed, ITs have supported the formation, operation and dissolution of VOs at two levels:

- Firstly, providing an extensible, open infrastructure to share resources between organisations, sometimes such resources are called *services*, in a flexible, secure and coordinated way. *Web Services* and the *Grid* are good examples of such technologies supporting the interoperability through protocols that define how organisations interact with one another in order to achieve and specific behaviour, and the structure of the information exchanged during this interaction.
- Secondly, providing mechanisms that are used for organisations to cooperate and coordinate their activities. Some of these mechanisms are based on *game theory* and economic models and have been adapted for MAS to support cooperation and coordination in VOs. In fact several recent projects have addressed the application of the agent paradigm and related market based negotiation mechanisms for the partner selection, service composition and task allocation during the formation and operation of VOs.

Cooperation in this context involves heterogeneous, autonomous, and not totally cooperative entities. Some related issues have gained particular importance, such as interoperability and information sharing, safe communications, privacy and authentication, and base support infrastructures. But cooperation in a networked environment is not purely a technological issue.

Whereas information technologies are definitely applied to improve the speed with which these VOs can be created and also for operating this organisational model on a large scale, social aspects that depend on the nature of the members of the VO relating to trust, improving collaborative behaviour, reward sharing, etc. have not been given adequate recognition as significant aspects in the field of VOs. These issues are important because they are related with the sharing relationships created during the operation of the VO. Organisations agree to share resources, core competences, skills and know-how with each other, but before engaging in a sharing relationship they must come up with cooperation agreements. Partner cooperation represents a slightly paradoxical situation:

individuals are supposed to pursue their own goals, but they are simultaneously required to restrain this natural pursuit in order to make VOs work.

Furthermore, relationships can vary dynamically over time, in terms of the resources involved, the nature of the access permitted, and the partners to whom access is permitted. And these relationships do not necessarily involve an explicitly named set of individuals, but rather may be defined implicitly by *policies* that govern access to shared resources.

The dynamic nature of relationships in a VO means that mechanisms for characterising and controlling the relationships that exist at a particular point in the time are needed.

In this work we present a mechanism based on the direct experiences between self-interested entities, representing different organisations, that can be used to build trust among the participants of a VO. This mechanism attempts to reconcile individual rationality with VO achievements. In particular, we will consider the history of past interactions known as *reputation* as a source of information to elicit trust and foster cooperative behaviour. Through reputation, it is possible for self-interested entities to determine the extent of trust that can be given to a partner and determine whether such a partner will be able to cooperate suitably within the VO.

1.4 Trust and Reputation in VOs

Because VOs essentially involve two or more organisations or agents cooperating to pursue a common objective, satisfactory cooperation is vital to their success. However, before an agent can make the decision to go ahead with the VO, it needs to be *confident* that the rest of the potential partners will act cooperatively. This feeling of confidence is known in this thesis as *trust*. A low level of trust leads partners to view each other with suspicion, with unavoidable deteriorations in their cooperative relationship. Or even worse, partners may flatly refuse to participate in a VO if they distrust each other, discouraging the formation of a VO. Clearly, trust is a basic ingredient in forming VOs because it determines not only whether potential partners are willing to cooperate, but also the degree of cooperation in the VO. There are several reasons for not having a benevolent character while forming a VO and blindly trusting in partner cooperation:

- the lack of a prior history of working together, which creates a sense of distance among the partners;
- the self-interested nature of the partners; and
- the lack of mechanisms to enforce cooperative behaviour.

Trust among partners is critical for VOs. Without trust, commitment to the goals of the organisation can waver, as members perceive the alliance as weak or disintegrating, fractured by misunderstanding or mistrust. Moreover, it is difficult to trust people you do not know well, who you have not observed in action over time, and who are not committed to the same goals.

Then, individuals must be able to gather as much information, through sensing capabilities, about their potential partners and the environment. This information will be used to take decisions about either to engage, or not, in a VO, or to select, or not a potential partner. Additionally, actions of partners in a VO can be constrained by the system within which the VO evolves. In this sense, it involves devising mechanisms that exercise a level of control on the interaction inside the VO so as to directly influence the behaviour of its partners.

Moreover, when interactions are constrained, risk might be eventually decreased and partners can make optimal decisions when they have completely evaluated and reasoned about the information collected. However, in a practical context, taking optimal decisions is impeded by physical and resource limitations. For example, limited computational and storage capabilities united with limited bandwidth and speed of communications channels clearly restrict the individuals' capabilities for sensing in real world applications. Thus it is not difficult to conclude that it is practically impossible to reach a state of perfect information about partners' capabilities, intentions, and properties. Individuals have to face considerable degrees of uncertainty while making decisions. This situation leads individuals to trust each other in order to decrease the risk associated with participating in VOs. Trust has been defined in a number of ways in different domains (see for a general description). However, we find the following definition, that is based on the definitions of Jonker [41] and Castelfranchi [11], most useful for our purposes:

Trust is: *the attitude of an agent towards an interaction partner about how likely it is to act honestly and cooperatively in situations where the achievement of a common goal depends on partner's actions; that results from reasoning about its capabilities, willingness, reliability, intentions, etc..*

There are generally two approaches to trust in VOs characterised by internal and external factors. Firstly, the internal factors are related with some form of evaluation (reasoning, experience, perception, etc.) of others competence, knowledge, reliability, goodwill of an individual. Hence, to allow individuals trust each other, they must able to reason about the reciprocative nature, reliability, or honesty of their partners. This ability is captured through *trust models.* Trust models aim to enable individuals to calculate the amount of trust they can place in their partners. In this way, a high amount of trust in a potential partner would entail its selection as a partner while a low amount of trust would result in its rejection. The knowledge that allows us to take such decisions

is acquired through inferences drawn from the outcomes of multiple direct interactions with different partners.

The external factors involve rules or mechanisms to establish protective conditions and guarantees. These mechanisms are called *control mechanisms* and aim to assure trust among partners and to prevent deception and fraud. They are based on disciplinary actions to punish the violators (contracts, fines, roles, institutions, etc.) and ensure that the participants will find no better option than behaving in trustworthy way by interacting honestly with each other.

Time is an irreplaceable prerequisite for trust-building in partner relationships. Sometimes, due to the highly dynamic life cycle of a VO, partner relationships could be characterised by service-oriented relationships that may entail no past history, nor plans for future association. In these temporary relationships, time is a vital but unusual component in the trust-building process. The lack of a prior history of working together as well as no face-to-face communication could create a sense of distance among the partners, weakening the identity as a group. In such situations reputation is a valuable source of information for building trust. Therefore, supplementary information about potential partners is required in such situations and reputation provides such information.

Reputation of an individual is *a perception regarding its intention and competences, which is held by other agents through the formation and dissemination of subjective evaluations which are based on experiences and observations of its past actions* (adapted from [44]).

Reputation is a basic ingredient in the trust models. Reputation about a potential partner involves information of its previous interactions, including its past behaviour, intentions, competences that are fundamental for both selecting a partner and creating the conditions to foster cooperation among the members of a VO.

Additionally, reputation provides valuable information to detect free-riders, and it acts as a mechanism of control in two ways, (1) as a warning level, to protect others from transgressors, and (2) as a punishment: individuals that suffer from bad reputation are likely to be excluded from VOs, and will not find partners for exchange. Reputation also affects the interactions in a community. According to Raub [58], if an individual anticipates that his current behaviour will affect not only the immediate consequences he will face in the actual situation but also the later behaviour of his partners and thus his own future consequences, then he has an incentive for trade-off between short-run effects of present decisions and their long-run effects on reputation. This trade-off can significantly affect the individual's behaviour. Hence, the reputation of partners can be used for instilling the motivation to cooperate in the members of a VO. This motivation also establishes trust among the members of the VO. The reputation of a potential partner also motivates it to act in a cooperative fashion and deter any abusive

behaviour. If a partner becomes indifferent to its reputation and continues to act in an opportunistic way, it is rejected from the VO.

As can be seen, reputation pervades the formation process of VOs in may forms. Reputation is an important factor in the formation of VOs because of the following reasons:

1. it is a determinant in the selection of potential partners because it allows the restriction of access to free-riders or partners of lower status (i.e. those with low level capabilities or those who are untrustworthy);
2. it provides useful information about the capabilities of potential partners;
3. it is a signal of the trustworthiness of the potential partners;
4. it induces *good* behaviour; and
5. it establishes the basis for trust building.

We will further detail the relevant work in the literature given the above concepts of trust and reputation in Chapter 3 and describe our model of reputation-based trust in Chapter 4 which cope with the requirements listed above.

1.5 Aims of the Research

This research aims to develop a model of reputation based on direct interactions to select partners for creating a VO. It aims to support the formation of VOs and at the same time fosters cooperative behaviour during the operation of VOs. Special attention is given to the *selection partner* process which will be supported by an adaptive reputation mechanism.

A central argument of this work is that *reputation*, which is understood to be a process for gathering and transmitting the evaluation of direct interactions, underlies trust building. Our focus is on a means for establishing trust rapidly in the VO. By using reputation in the process of creation of a VO, partners trust to the extent determined by the reputation of the others, avoiding deceptions created by false expectations.

Thus, the objectives of this research are to:

- describe what are the problems to be faced in the formation of highly dynamic VOs due to the lack of mechanism for dealing with the uncertain behaviour of their self-interested members and fostering cooperation; and shown how reputation provides the means for coping with them;

- develop a model of reputation that enable the formation of VOs, and allow estimate the risk of interacting with a certain partner by evaluating its behaviour in a reduced number of interactions; and
- experimentally demonstrate both the validity of the model developed and the fulfilment of the requirements of the VO formation task; in particular, perform a comparative analysis of the reputation model proposed and existent reputation models to demonstrate the accuracy of values of reputation calculated and their efficiency for detecting abusive behaviour in the context of VOs.

To achieve these objectives, a model of reputation has been developed. The model presented builds upon reinforcement learning techniques. In particular, the following contributions have been made to the state-of-the-art. Firstly, we have delineated and computationally modelled the evaluation of dynamic interactions in VOs. Secondly, using reinforcement learning, we defined a measure of reputation and an update mechanism that is derived from direct interactions and the information gathered from other members in a community. Thirdly, we have shown how reputation can guide the selection of interaction partners, an issue that needs to be agreed during the negotiation, and how reputation can build trust among the members of a VO.

1.6 Thesis Overview

Our research on the formation of VOs based on reputation information is presented in six chapters, including this one that provides the motivation context for the work, and they are organised as follows:

Chapter 2 introduces the concept and characteristics of VOs. Next, the two broad categories of VOs, namely Static and Dynamic VOs, are presented. Then, the problems associated with the formation of VOs are introduced and the problem of selecting partners is explained in detail. Thus, the chapter explores the issues of creating and maintaining trust in VO. Finally, reputation is introduced as a way to cope with the problems of VO formation.

Chapter 3 critically reviews the literature related to reputation in the context of electronic communities. Next, the requirements for reputation systems in general and those specific for VOs are discussed. Finally, the chapter discusses the weaknesses and omissions of current reputation systems in relation to the requirements of VOs.

Chapter 4 is a key chapter in this work, and it describes a model of direct reputation. This model is based on reinforcement learning techniques.

Chapter 5 provides a comparative experimental analysis to demonstrate the validity of the proposed model, its results and conclusions.

Chapter 6 summarises the model, the main contributions and the limitations of this research, and the potential for further investigation.

Chapter 2

Background on Virtual Organisations

2.1 Introduction

Not long ago, most of the developments in electronic commerce were associated with a passive web query type interaction between customers and suppliers, who wanted to buy or sell a good. However, in recent times, electronic commerce technology has evolved not only to allow more complex transactions but also to enable the emerging of new entities known as Virtual Organisations (VOs).

VOs underpin two important trends in distributed computing that have acquired strategic importance for industry and government. These trends include business-to-business e-commerce and e-science. Today, businesses advertise globally for their resources, and organisations no longer require their members to be located in the same place to perform their work. VOs enable the formation of strategic alliances and electronic value chains.

The concept of VO has been used to describe the aggregation of autonomous and independent organisations connected through a network and brought together to deliver a product or service in response to a customer need [53]. VOs are not exclusive to e-commerce, and they have been extended to diverse areas such e-science [62].

In this chapter we introduce the central subject of this dissertation. First, the concept of Virtual Organisations is presented and its relation with similar forms of organisation is discussed. Next, we provide two categories of VOs, namely Static VOs and Dynamic VOs. The life-cycle of VOs is explained and the problems associated with the formation of VOs are discussed. We next examine the shortcomings of current work in VOs related to the problems of VO formation, and analyse how such problems provide a context for developing models that consider the nature of partners not only in their selection but also

to foster cooperation among them. This analysis is then used to justify the introduction of reputation in the formation of VOs.

2.2 VOs and Related Concepts

In this work a VO is *a temporary alliance composed of a number of autonomous entities (representing different individuals, departments and organisations) each of which has bounded problem solving capabilities and limited resources at their disposal, that come together to share skills or core competences and resources in order to better respond to customer need or business opportunity, and whose cooperation is supported by computer networks* (adapted from [51]). Figure 2.1 shows the VO. In this model, entities (*service providers*) offer services that complement one another to form a VO.

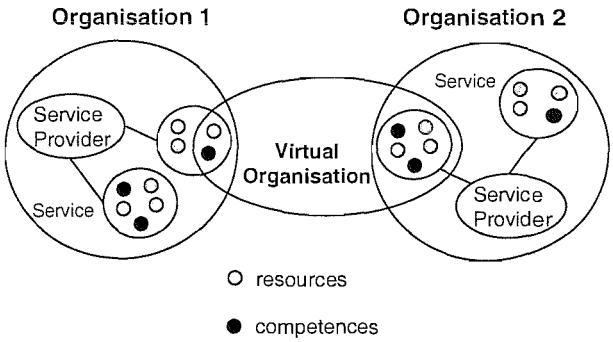

FIGURE 2.1: The Virtual Organisation Model.

These entities share places (the Internet, laboratories, marketplaces) and sometimes compete with one another in a virtual market place. Initially, each entity attempts to attract the attention of potential customers, who are looking for service providers, and eventually tries to sell them its services describing the cost and quality of the service.

Sometimes, however, due to their limitations and after realising of the potential benefits to be obtained from pooling resources and capabilities, one or more of the entities may try to group with competitors or partners. The formation of groups of competitors is called a *coalition*. Examples of coalitions are the groups of customers that come together to get goods at a volume discounts (*buying clubs*) [73].

On the other hand, entities group with partners with complementary expertise in order to offer a new type of service, by sharing skills or core competences and resources. When this potential is recognised, the relevant entities go through a process of trying to form a new VO to exploit the perceived niche.

There is no guarantee of success in the grouping process due to the independent nature of the entities involved. However, if it succeeds, the collection of independent entities have to start acting as a single conceptual unit. This means that these entities need to

cooperate and coordinate to deliver the services of the newly formed organisation. In dynamic environments, it is likely that the VO is no longer viable because the requirements of the customer change or the context varies at any time. Then, entities will be forced to disband or re-arrange themselves into a new organisation that better fits the prevailing circumstances.

A basic characteristic of a VO is that communication between cooperative entities is supported by computer networks. Camarinha-Matos in [6] says that due to the wide variety of different networked organisations, a number of related concepts such as those in Figure 2.2 have been introduced.

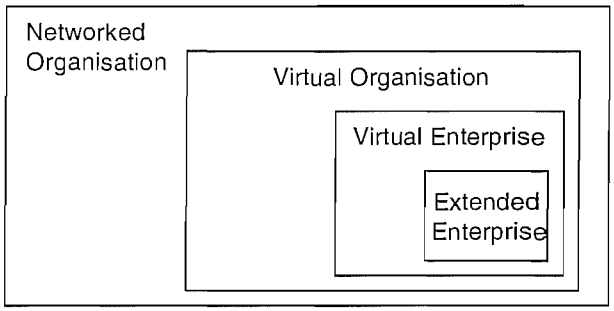

FIGURE 2.2: The Virtual Organisation and related concepts.

1. *Extended Enterprise.* This term is applied to those organisations in which the dominant entity *extends* its boundaries to all or some of its suppliers.

2. *Virtual Enterprise.* Different from an *Extended Enterprise*, a Virtual Enterprise include other types of organisations forming a more democratic structure in which the cooperation is peer to peer.

3. *Virtual Organisation.* This concept is similar to a Virtual Enterprise, but it is not limited to an alliance of enterprises. An example of Virtual Organisations could be a virtual municipality organisation, where all the organisations involved in a municipality, such as city hall, municipal water distribution services, tax council, etc., are associated via a computer network.

4. *Networked Organisation.* This is the most general term referring to any group of organisations inter-linked by a computer network, but without necessarily sharing skills or resources, or having a common goal.

2.3 Characteristics of VOs

VOs pose significant challenges from a managerial and computational point of view. These challenges are entailed in five factors that distinguish the VO from other organisational forms, according to Byrne: excellence, opportunism, technology, trust, and openness [5].

1. *Excellence.* Each partner in the VO has its particular core competence at which it excels. Companies must make rational decisions about the process of partitioning the global goal of the VO in order to make an efficient use of their resources.

2. *Opportunism.* The VO is opportunistic by nature and highly adaptable, and it may be disbanded and/or reformed relatively quickly. *Reactivity* (the ability to react to changes) and *flexibility* (i.e. being able to adapt its structure to satisfy the customer need) are important benefits of a VO [46]. Partnering underpins these elements of the concept *virtual.* Companies must choose the right partners for the right reasons: because they are dependable, can be trusted, and offer the best products or services.

3. *Technology.* The VO strongly relies on communication technology and information networks to identify opportunities, partners and customers, while these technologies at the same time provide the infrastructure enabling the collaboration itself.

4. *Trust.* The VO depends on the collaborative actions of its members. Each partner specialises in doing a part of the process and cedes partial control of its operations to its partners, which obviously has some risks. Companies must build trust with their partners. In other words, trust between the partners is a central premise of the VO, and all partners are reliant on each other, as their success depends on mutual cooperation.

5. *Openness.* Companies have to be more open in dealing with different partners. Every joint venture must offer an utility opportunity for everyone. The VO must serve the interests of all parties.

Although the VO concept has been more commonly applied to define the temporal association of independent organisations that come together to exploit and apparent market opportunity, it could be also used to refer to a group organisations inside another organisation. Of course, this requires further explanation of how a VO is formed inside a traditional organisation. To explain this, consider the following example that involves the refinery business.

The refinery business entails tasks that span several departments and process large amount of data. Among others, these include crude procurement, logistics and scheduling. Decisions in these tasks are normally taken based on business policies. These decisions impact the overall economic performance of the refinery. However, with volatile crude prices, excess refining capacity and fluctuating demands and prices for different products, refiners are turning to new approaches to manage their business profitability.

Decision-making is distributed across various departments in a refinery. Each department solves subproblems, but local improvements do not necessarily assure that the overall process is moving towards the optimum. Often, the department objectives are

conflicting, and thus all decisions do not contribute positively to the overall performance of the refinery. Temporary VOs are created for assisting different departments in taking group decisions. For example, a VO can be set up for planning and scheduling an unexpected demand for petrol, which may involve the participation of various departments; for example the departments of procurement, sales and storage, among others.

This example illustrates the central features of the VO. The success of the refinery depends completely on the match between requirements and offerings of resources (storage availability, refining capacity, etc.), on the one hand, and the abilities and desires of potential clients, on the other. For a refinery to be efficient, it must account for the demand side, or the needs and concerns of potential clients, as well the supply side, or storing and refining services available. Note that both the request of services (e.g., product demand, storage, distillation process, etc.) and the set of available services (store, refining operations, transport of crude, selecting shippers) vary dynamically over time. The VO dynamically carries out the process of assigning services to request to accommodate the dynamic character of the refinery business.

As in the above example, VOs have long been used by several industries. Many of them have been built over time and have lasted several years in the same relationship. Cisco Systems is one of the best-known examples of a long-term VO. Cisco focuses on developing new products and selling these products to customers while leaving everything else to its business partners. However, as mentioned before, within today's high-speed economy, it can be envisioned more dynamics VOs, which must be rapidly created for pursuing specific market opportunities and after satisfying their goals they could be disbanded or reformulated to pursue other market opportunity.

2.4 Categories of VOs

Although there is no strict academic definition regarding VOs, different VO models feature common business and technical characteristics and attributes. Two well-defined categories of VOs can been identified. These are:

- Static Virtual Organisations
- Dynamic Virtual Organisations

Next in this section, these two generic VOs categories are further analysed and certain examples are provided to better clarify the different models, concepts, and benefits and drawbacks that both approaches share.

2.4.1 Static Virtual Organisations

Ouzounis mentions in [54], that in Static Virtual Organisations (SVOs) a set of partners is linked together in a static and fixed way, i.e. their actions are tightly integrated. The relationships among the partners, are pre-defined, tightly coupled, fixed, well integrated, and customised among the partners. The number of partners is fixed and pre-determined and thus, the structure of the VO is static and predetermined.

Typical examples for SVOs are models that have been applied in the automotive manufacturing business. In this case, a big automotive manufacturer has a network of suppliers, distributors, and distributors that are working together in different phases of the production, distribution, and reselling process. The big manufacturer has specific needs and requirements and enforces his requirements in order to increase the degree of automation and decrease the production and distribution costs. The network of suppliers, resellers, and distributors closely cooperate with the central partner or *initiator*.

A more recent approach to automate the process of forming a SVO is to use a virtual marketplace or a directory service where potential VO members register their resources and services. The virtual marketplace provides matchmaking services to VO's initiators that want to locate VO partners. This approach takes advantage of the new, open, Internet economy and significantly improves the formation process of the VOs. However, after the formation of a VO, the relationships among the partners remain static and fixed, while the evolution of the VO in terms of new members, that might provide better services with better terms, is impossible [54].

Marketplaces can be used in a more effective and dynamic way not only in the formation phase of VOs but also in the operation phase. This means that the partners that are involved in the provision of a complex goal are changing continuously and dynamically according to the requirements of the customer and the goal. In this case, for every goal, a new VO is created in a dynamic way, addressing the needs and requirements of the customer and the individual partners. The deployment of marketplaces not only for the establishment of VOs but also during the provision of complex goals can lead to significant improvements.

2.4.2 Dynamic Virtual Organisations

According to Ouzounis [54], in Dynamic Virtual Organisations (DVOs) a set of partners is linked dynamically, on demand, and according to the requirements of the customers, by deploying a virtual marketplace. Partners do not have fixed relationships and thus the VO is not static and might change continuously based on market driven criteria.

The virtual marketplace provides services for the registration of partner process offerings based on some generic, well-known, globally specified process templates. VO initiators

that want to form VO relationships can register offers in the marketplace related to certain complex goals. Whenever a VO initiator wants to use a particular service, searches the marketplace, and locates all the potential partners that can provide the service. As soon as the list of VO candidate partners for one particular process has been found, the selection process starts. The selection process between potential partners is usually performed through negotiation. The negotiation process might be either manual, or automated, while the result of it is usually a short term contract that regulates the relationship among the involved partners [54].

Marketplaces are usually organised around certain globally specified service or product that can be offered by the different service providers. The marketplace is a matchmaking mechanism that brings potential service providers together with potential users of these services. Although marketplaces and matchmaking mechanisms have been used for some time for business-to-consumer electronic commerce purposes they have not been actually deployed for dynamic VO purposes. The main reason is the lack of technologies that enable the easy and flexible definitions of services, mechanisms for automated negotiation, and autonomous interaction among different partners.

2.4.3 Life-cycle Model of VOs

Fisher et al. are pioneers in the use of intelligent agents for the design and operation of VOs. In their work [31], two phases are identified for setting up VOs. In the first phase, a suitable service partitioning is worked out. In other words, the service that impels the creation of the VO is divided to be assigned to its members. In the second phase the partners are selected. This second phase includes four sub-phases that go from the identification of potential partners to the mapping of partial sub-processes and partners. Camarinha-Matos et al. identify four stages in the life-cycle of a VO including: the *creation*, *operation*, *evolution*, and *dissolution* stages [6]. Faisst in [30] coincides with Camarinha-Matos in some phases and establishes five phases in the life-cycle of a VO: *identification*, *formation*, *design*, *operation* and *dissolution* phases.

There are certain similarities in the phases of the VO life-cycle described above, which we summarise in four phases:

- *Identification of Needs* is what will motivate the creation of the VO. By describing the products or services to be delivered by the VO, it is possible to identify the needs which guide the conceptual design of the VO.
- *Formation* considers the rational search and selection of the partners which will compose the VO based on its specific knowledge, skills, resources, costs and availability. This phase concludes with the contract negotiation process.

- *Operation* includes the control and monitoring of partners' processes, the resolution of conflicts, and possible VO reconfiguration due to partial failures.
- *Dissolution* is the last phase, which consists of breaking up the VO, the distribution of the obtained profits and storage of relevant information.

Although the advantages of VOs are well known at the conceptual level [25, 5], setting up a VO is far from trivial and several problems must be faced. Camarinha claims that for a VO to have the benefit of being an agile organisation in terms of fast reaction and adaptability to business opportunities, it requires *flexible* and *generic* infrastructures that support the full life-cycle of VOs and cope with the problems associated with them [7]. In this work we focus on the problems derived from VO formation and we assert that *reputation* is an important factor that affects the structure and outcome of a VO.

2.5 Problems of VO formation

In the literature, researchers [7] have found that some of the problems that a VO faces while it is being formed are:

1. Task division. The problem of task division can be formulated as the optimal decomposition of the VO's global goal into small pieces of work or tasks. Then, tasks can be passed to potential partners who may decide to tackle them or alternately decompose them into subtasks.
2. Task allocation. This problem poses a question related to the criteria for assigning the set of tasks identified in the task division phase to the appropriate partner. This problem involves subsequent subproblems: how are partners with the capabilities required discovered, and which eligibility criteria should be used in order to select a partner? These questions are designated as the partner search and selection problems.
3. Partner search. This problem establishes the need for mechanisms for identifying and locating potential partners. This is because in practice it is not possible for a partner to have a list that identifies the capabilities of partners, particularly when the set of tasks demands combinations of specialists in different areas to share their expertise to accomplish the global goal of the VO. Moreover, the problem is more complicated when several candidate partners are capable of accomplishing a certain task.
4. Partner selection. Since the selection of partners is considered an important activity in the formation of VOs [55, 51, 33], it has received an enormous amount

of attention so a number of pieces of work have populated the VO research topic. Geringer argues that the selection of the *proper* partner influences the mix of capabilities and resources which will be available to the VO and thus the VO's ability to achieve its strategic goal [33]. Hence, identifying which criteria a VO might use in attempting to select a partner is vital.

5. Contract bidding and negotiation. Partners in a VO should negotiate the terms by which they provide their services or resources. We discuss the difficulties of contracting in VOs below in this section.

Since VOs must be capable of rapidly adapting to different necessary processes, which include modifying the structure of the organisation by adding or removing partners in a timely and effective manner, there are other problems in the operation phase that are directly linked to the formation phase and must not be ignored:

1. Performance assessment. Camarinha argues that as part of the selection process and during the operation of the VO, an assessment of the partnership performance of the potential partners, based on their cooperative behaviour, should be made [7]. Here, the problem is to find measures of the performance of the partners in order to capture their real capabilities and detect possible defectors.

2. Monitoring and coordination of task execution according to contracts. Once VOs are formed, they require ongoing attention to operate effectively as a whole. This attention is given by monitoring the relationships not only between the current partners, but also with potential ones [15]. Monitoring is required to capture the extent to which services meet the expectations of the consumers.

2.5.1 Current solutions and shortcomings

Computing scientists have proposed diverse tools, protocols and frameworks for setting up VOs [51]. Some tools and protocols are designed for supporting collaborative interactions and standardising the communications between the VO's partners. By contrast, frameworks have bigger scopes and they aim to provide functionalities related to the creation, operation and dissolution of VOs. Here, we discuss the criteria used by current frameworks that address the problems of the formation phase of VOs.

2.5.1.1 Partner Search

Diverse approaches have been proposed [57, 18, 16] to solve this problem, but there is a characteristic that such approaches have in common, that is, the use of *directory services*. Directory services are repositories of information where potential partners register their

core competences or offered services and also provide *search mechanisms* for partners to be easily discovered by *brokers*. The aim of the brokers is to match the task to the potential partner capable of accomplishing that task. A variant of the latter approach suggests that partners should be selected from a pool of individuals, previously formed, that have the potential and the will to cooperate with each other. These pools are called *clusters* [6] or *clans* [35].

Shortcomings. The approach of using pools reduces dramatically the search effort for locating potential *trusted* partners. However, it is not clear how the pool evolves (for example what happens when a certain required capability is not found in the pool), nor how the relationships created by the partners after cooperating in diverse VOs influences the formation of pools neither how such pools affect the behaviour of the partners.

2.5.1.2 Partner Selection

The most popular criterion of selection is based on the *price* asked by potential partners to accomplish a task. This criterion for allocating tasks is common in market-based approaches for distributed resource allocation. In the context of the formation of VOs, researchers have adopted market-oriented approaches for allocating tasks through the interactions of trading, i.e. buying and selling, among potential partners and the initiator of the VO. The most suitable set of partners who are to be selected are the ones with the *lowest price* bids [51, 16, 31, 56].

Nevertheless, Camarinha asserts that the open market model is not adequate for forming VOs, as the involved organisations need to first *trust* their partners due to the liabilities involved [7].

Moreover, alliances are considered risky because of the unpredictability of the behaviour of the partners and the likely costs to an organisation from *opportunistic* behaviour by a partner if it occurs [36]. That is, partners may free-ride by limiting their contributions to an alliance or simply behaving maliciously.

Shortcomings. It is important to note that prior work regarding partner selection based on market-oriented approaches has neglected the self-interested nature of the potential partners and has assumed that partners are fully competent, cooperative and trustworthy. This naive assumption about the *benevolent* character of the potential partners used in the current frameworks [51, 16, 31, 56] has serious consequences:

- the potential partner may not have the required capabilities for performing the task; or
- it could have the capabilities but not the *intention* to do so; or

- it could have the propose of cheating the others by claiming the benefits from the alliance but without performing its tasks; or
- it could hide its lack of abilities by blaming the others from uncompleted or poorly performed tasks.

CONOISE [56] deserves an important clarification. CONOISE is based on self-interested agents and does not neglect the self-interested nature of potential partners in VOs. In CONOISE, a self-interested rationale is behind the decision of either joining a VO or determining who is the most suitable set of partners. Hence, an agent joints a VO only when it is profitable, and it only accepts other agents into the VO with the lowest price bids for performing a task. However, it is clear that CONOISE assumes implicitly that agents completely trust each other, and questions of deception and fraud are ignored while a VO is formed.

We argue that a complementary criterion incorporating aspects of partners real *competences* and *reliability* should be taken into account in partner selection decisions, since these aspects can influence the cooperation between partners. Therefore the partner selection problem is discussed from this perspective later in Section 2.6.

2.5.1.3 Cooperation

Although researchers recognise the advantages of VOs as a form of organisation to create customised assets and cope with product demand uncertainties, complex tasks under time pressure and frequent exchanges among partners comprising the VO based on cooperative and trust relationships, a small number of frameworks include mechanisms for fostering cooperation and trust. Instead, they employ legal contracts to enforce the partners to accomplish their tasks in the VO. However, contractual obligations are exceeded when situations that have not been previously specified appear, for instance when the demand for a product claims a drastic reduction in the delivery time and organisations are forced to adjust their services to meet the new demands. In such situations trust is important to allow that the partners adjust their tasks to the new situation. Additionally, Jones [40] mentions that formal contracts may exist between some pairs of members, but these do not define the relations among all the partners.

As an example, consider a VO formed by three partners SP1-SP3 working synchronously to deliver a customised and complex service. After a certain time, the client decides to modify the characteristics of the service implying changes in the synchronisation of the tasks. If the VO is based exclusively on the contracts among the partners, then new contracts need to be renegotiated to fulfil these changes even though additional tasks may not be required. Otherwise, if the VO is based also on trust relationships, partners can adapt their tasks to the new requirements without spending time in the negotiation

task. Now, consider that SP1 is the initiator of the VO, and SP2 and SP3 negotiate separate contracts with SP1 for providing tasks t2 and t3, respectively. Task t3 depends on the outcomes of t2. Hence, when the VO is in operation, (1) SP3 cannot enforce SP2 to provide t2's outcomes as required by task t3, (2) it is likely that SP3 blames SP2 for its own failures in task t3, and (3) either SP2 or SP3 may refuse to cooperate with each other.

Flexibility to adapt themselves makes VOs an attractive approach for facing demand uncertainty in dynamic environments, but if it implies renegotiating contracts each time the VO is re-arranged, then the possible advantages vanish because many opportunities can be lost whilst the renegotiations take place. Then, to enhance cooperation on shared goals, the VO requires to rely more heavily on other forms of control apart from contracts such as *social control.*

2.5.1.4 Social Context

Up to now, frameworks have made use of a *dyadic* perspective (i.e. a peer to peer notion that does not consider the action of other partners or the relationships in which they themselves are already embedded) to treat each interaction as a discrete independent event, as occurring in an ahistorical context. Both cooperative and uncooperative actions are rewarded or punished when the VO exists and do not have future implications. An interesting alternative to foster cooperation is the use of *social mechanisms* of control.

VOs are not formed in a vacuum, they are embedded in societies or communities of individuals or organisations. Their members form *social networks* through the relationships created by previous interactions that link them directly with each other or indirectly by third parties. This links different VOs together and spreads information about third parties among those within the *social network*, which allows information, and norms to move across VOs boundaries [40]. Having this in mind, reputation is a social mechanism of control that provides information about the reliability and goodwill of potential partners, and when such information is spread in the social network, it fosters trust because it relays the detection of, and serves to deter, deceptive behaviour.

2.5.2 Discussion

One of the major problems that we observe in current solutions is that they underestimate the nature of the current and future partners of VOs. A common flaw is assuming that the partners selected are fully competent and honest. Nothing could be further from reality. Partners represent organisations or individuals who want to maximise their utilities by joining a VO and, hence they have a strong incentive to misrepresent the value of their contributions and enjoy more benefits of cooperative associations [4].

Further, partners are selected regarding the abilities they claim to have, but it is possible that they do not have such abilities. However, due to lack of information about past interactions it is difficult to detect and control those situations.

Undoubtedly, one of the cornerstones of a VO is cooperation. Partners agree to share their capabilities and resources to achieve a common goal. However, it is possible that a partner refuses to cooperate even if it agreed contractually to do so because it could be more profitable to accept a fine for breaching the contract than decline a better opportunity in another VO. Hence, it cannot be assumed that partners are cooperative and do not drop their obligations.

As can be seen, neglecting the self-interested nature of partners while forming or operating a VO influences their outcomes. By contrast, we argue that by introducing reputation in the process of formation and operation of a VO, partners are endowed with the ability to reason about the behaviour and competences of their partners. In the next sections, partner selection and cooperation are discussed in more detail.

2.6 Partner Selection

Geringer in [33] empirically provides a simple two-fold typology of categories for selecting partners. The typology suggested is based on the distinction between task-related criteria and partner-related criteria. Task-related criteria refer to those variables that concern the viability of achieving the VO's objectives, irrespective of whether it involves multiple partners. In contrast, partner-related criteria refer to those variables that become relevant only if the organisation involves the presence of multiple partners. Examples of these variables are trust between the partners, and the degree of favourable past association between partners and reputation.

Based on Geringer's typology, we found in the literature that both criteria have not deserved the same recognition; researchers have given more emphasis to task-related and utility criteria when they provide algorithms for selecting partners. For example, in [3] an algorithm is provided for selecting the most beneficial partner based on an expected utility when a coalition is formed. Petersen and Divitini [55] propose a model for support the selection of Virtual Enterprise Teams that is based on the roles required for performing a set of activities. Norman et al. [51] present agent-based models for the automated formation and maintenance of VOs based on the list of tasks and their scheduling and on the *price* asked by the potential partners. A remarkable exception is the work of Griffiths and Luck [35] who introduce a mechanism for agents to form coalitions that considers not only the capabilities needed to the achievement of the strategic goal but the believed trustworthiness of the potential partners. However, this model fails when there is lack of information about others trustworthiness or when there is not enough information about their capabilities.

In order to cope with the omissions in the literature regarding partner-related criteria, we propose the use of reputation. *Reputation* plays an important role here by providing information about the past behaviour of potential partners and their capabilities. By comparing the aptitudes of an individual inferred from its reputation with abilities required to perform a task, it is possible to determine if such individual can be trusted to accomplish the task well. Hence, for a VO it is not necessary to experience and obtain undesired outcomes by interacting with a *harmful* potential partner if reputation information about it is considered whilst the partner selection decision is made.

2.7 Cooperation in VO

Despite the technological efforts to support the communication and exchange of information in order to facilitate the *operation* of VOs, what determines the success of a VO is the partner's commitment and cooperation. Das and Teng [23] cite a lack of cooperation and the opportunistic behaviour of partners as causes for their relatively high rate of failure of strategic alliances.

Hence, partners cooperate on a short-term basis to reach a common goal; once the goal is satisfied, the desirable next action is re-arranging the VO structure that better fits the prevailing circumstances and, eventually, exploits the relationships already created. However, it is not possible to guarantee cooperation among partners, mainly if they are competitors, unless common rules (*social norms* and *contracts*) or agreements are adopted.

Social norms prescribe social acceptable conduct (such as benevolence or altruism) and forbid socially unacceptable conduct, and are supported by social incentives and sanctions [21]. Potential sanctions include exclusion from the *society* for short periods or indefinitely, loss of future points of interaction, and loss of reputation. Conte establishes that social norms emerge thanks to reputation [21].

2.7.1 Limitations of Contracts

Most commercial transactions rely on the legal system to assure performance of promises, which are written into explicit or implicit contracts. Contracting is often suggested as a medium for inducing cooperative behaviour among the members of VOs. This is partially advisable, but some reasons may indicate the contrary.

2.7.1.1 Time- and resource-consuming

Legally enforceable contracts are usually very time- and resource-consuming in the sense that the creation of a contract implies several phases to be developed before the law

recognises the promises granted in the contract as a duty. In the case of VOs, contracting may be more difficult to establish because it could mean finding agreements in the provision of a good or service with each partner and for each service provided.

2.7.1.2 Inflexibility

There are some issues which are not feasible to describe in a comprehensive contract. For example, it is not possible to restrict the partners to use the information exchanged for their own benefit. That is, partners agree on sharing enough technical information in order to fulfil the requirements of a service. However, such information can also be used by them to exploit their former partners, for instance by demanding lower prices in future trades whenever a service provider has found a way to reduce costs.

McCutcheon [47] affirms that when organisations are involved in a specific, discrete economic transaction, an explicit contract spells out before trading commences each party's tasks and duties in every conceivable eventuality, as far as human capacity for anticipation allows. However, forming strategic alliances usually means opening up more informal channels of communication and broadening the range of expectations, many of which cannot be effectively captured in a legal contract. As a result, the parties involved must have a means of keeping the risks restrained as they develop more reliable ties among them.

2.7.1.3 Weak enforcement

A contract and the penalties imposed for breaking such a contract cannot deter a participant from dropping out of its obligations in the case that the participant considers it is more rewarding to do so. That is, if a participant decides to breach a contract in spite of the penalties, then the provision of the whole service will require both a new provider and a new contract.

2.7.2 Trust in VOs

Most of the literature [61, 47] indicates that organisations use *trust* to reduce the risks associated with establishing this form of organisation. Trust offers a supplementary mechanism to contracts to foster cooperative behaviour. Through trust, potential partners may assess the perceived risk in interactions. In this sense, partners can estimate the extent to which a potential partner is able to fulfil its task in the VO.

Gulati [36] mentions that an important concern for organisations entering alliances relates to the predictability of their partners' behaviour. A detailed contract is one mechanism for making behaviour predictable, and another is trust.

Thinking about contracts as a control mechanism, we agree with Das' affirmations [24] in the sense that while trust leads to low risk perception without doing anything about the partner, control is a more proactive and interventionist approach and leads to low risk perception too but through affecting the behaviour of the partner.

The resulting reduction in complexity to evaluate the risk makes things much simpler for the partners; complexity is not used here in the strict computational sense but to refer the process involved in the contracting process to achieve an agreement that is time and resource consuming. Hence, trust is much more flexible than contracts. A contract usually specifies background conditions that must obtain for the contract to be valid. An allowable defence against breach of contract is that the contract was not valid at the time of the breach (independently of intentions).

As an example, consider again a VO formed by three partners *SP1-SP3* performing tasks t1-t3 synchronously to deliver a customised and complex service. Task t1 output is the input for t2, and t2 output is the input of t3. Let's analyse two cases. In the first case cooperation is ruled exclusively by a formal contract. In the second case cooperation is ruled by a formal contract and supported by trust relationships formed through of a history of positive interactions between the three partners. Now, suppose that after certain time the client decides to modify the characteristics of the service implying changes in task t2. Although these changes in t2 represent an extra effort for SP2, this fact will be known for SP1 and SP3 just after they have interacted several times. SP2 claims an extra charge for the changes in t2. In the first case, both SP1 and SP3 may distrust SP2's motives for claiming extra charges and decide to dissolve the VO, losing an opportunity of deliver the service to the client. In contrast with the first case, in the second case because of past history of positive interactions, SP1 and SP3 may agree with SP2's claims and continue cooperating in the VO. But if they realised that such claims were false they may punish SP2's behaviour by damaging its reputation and avoiding future interactions with it.

As can be seen, trust is required to cover situations that are not covered by formal agreements. An unwritten agreement accompanies cooperative agreements by which partners, who have mutual trust, understand what else is required or expected from them beyond formal contracts. Trust is important because cooperation between partners cannot work if formal contracts and agreements has to be done on every detail.

Time is an irreplaceable prerequisite for trust-building in partner relationships. Sometimes, due to the highly dynamic life cycle of a VO, partner relationships could be characterised by service-oriented relationships that may entail no past history, nor plans for future association. In these temporary relationships, time is a vital but unusual component in the trust-building process. A lack of a prior history of working together as well as no face-to-face communication could create a sense of distance among the

partners, weakening the identity as a group. In such situations reputation is a valuable source of information for building trust.

2.8 Discussion

Our position is that not only is the partner selection process critical for obtaining high levels of reliable performance in a VO, but it provides the conditions that enable a VO to develop confidence that its members will pursue mutually compatible interests (i.e. cooperate) rather than act maliciously. Reputation provides valuable information to detect free-riders, and it acts as a social control in two ways, (1) as a warning level, to protect others from transgressors, and (2) as a punishment: individuals that suffer from bad reputation are likely to be excluded from VOs, and will not find partners for exchange.

On the other hand, prior successful relationships affect alliance success. First, they allow an organisation to better understand the capabilities or resources it is accessing and the partner's likely behaviour. And second, prior interactions build trust between partners and a mutual understanding of how the organisations work together. A positive reputation formed thanks to positive prior interactions signals that a partner is *trustworthy* and decreases the perceived likelihood that the partner will defect [68].

Reputation also affects the interactions in a community. According to Raub [58], if an individual anticipates that his current behaviour will affect not only the immediate consequences he will face in the actual situation but also the later behaviour of his partners and thus his own future consequences, then he has an incentive for trade-off between short-run effects of present decisions and their long-run effects on reputation. This trade-off can significantly affect the individual's behaviour. Hence, the reputation of partners can be used for instilling the motivation to cooperate in the members of a VO. This motivation also establishes trust among the members of the VO. The reputation of a potential partner also motivates it to act in a cooperative fashion and deter any abusive behaviour. If a partner becomes indifferent to its reputation and continues to act in an opportunistic way, it is rejected from the VO.

In conclusion, the overall goal of addressing the problems of VO formation and the need for enhancing cooperation and fostering trust during the operation of VOs provides a context for reputation. Thus, reputation:

1. is essential in the decision to select a partner, because not only is the task taken into account but also the characteristics of the potential partner;
2. offers a medium for enhancing cooperation by detecting and deterring deceptive behaviour;

3. is an alternative mechanism of control due to making partners aware of the future implications of their current behaviour; and finally,

4. fosters trust among the partners of the VO by revealing their capabilities and expected behaviour.

Finally, reputation is concerned with three major kinds of coordination decisions: those involving the initiation of a partner into the network, the selection of a partner for a task, and the termination of a partner from the VO [15].

2.9 Conclusions

In this review chapter, we focused on the concept of Virtual Organisations and their characteristics and categorization, which is based on the duration and strength of the relationships created between their members. More specifically the different terms and definitions of VOs that have been proposed in this area were presented and analysed. Although a fully agreed concept of VO has not emerged in the academic community, this thesis proposes and adopts one. Then, the two broad categories of VO, namely Static and Dynamic were presented, and their main characteristics were explained. Furthermore, a life-cycle model for the formation and operation of VOs was presented. The life-cycle model specifies the main activities required for each of its phases, from the identification of needs to the dissolution of the VO. We assert that the success of a VO depends on the *proper* selection of the potential partners and in the relationships of cooperation that they develop during the operation of the VO.

The core of this chapter is the identification of the problems related to the process of formation of VOs. Though multiple tools, protocols, frameworks and projects have been proposed for enabling the automated formation and maintenance of VOs, the nature of the partners and how their capabilities and trustworthiness influence on the achievement of VO goals haven been underestimated. In the light of theories and empirical studies from Management and Organisational Sciences, we discussed how the significance of partner-related information not only in selecting partners but in fostering cooperation. We consider that the current work done in VOs from the perspective of Computer Science can be enhanced with the presence of concepts such as trust and reputation.

In general, the problem of modelling VO partners able to reason about the behaviour and competences of current and potential members is far from trivial, and although some research has been done, more is needed [8, 7]. To provide an effective model of reputation that meets the needs for information and allows partners to interact with others in an environment of confidence, we have to define the requirements of VOs based on their characteristics and analyse the current work in reputation systems. Thus, on the basis of these requirements, a model of reputation will be proposed.

Chapter 3

Background on Reputation

3.1 Introduction

The study and modelling of reputation has gained the interest of scientists from different fields such as: sociology, economics, psychology, and computer science. In the realm of computer science, several researchers [19, 38] have found that reputation has an effect at two levels in an electronic community. At the personal level, reputation plays a crucial role in the interactions between two or more entities; that is, an entity builds an internal representation, *direct reputation*, in order to evaluate the behaviour of a given partner after each interaction. At the social level, the process of transmission of personal evaluations about the behaviour of a current or former interaction partner not only affects the opinion of others, known as *social reputation*, about such a partner; but at the same time influences the reputed partner's behaviour.

In this thesis, reputation is an individual property that results from the transmission process of personal evaluations about an individual's capabilities and intentions at a social level, where the members of a community ascribe a reputation so as to label their behaviour. Certainly, this process involves the dissemination of personal beliefs among the community members and the acceptance or refusal of those beliefs.

These evaluations are necessarily formed and updated over time with the help of different sources of information. Computational models of reputation have considered basically two sources:

- the direct interactions and
- the information provided by other members of the society about experiences they had in the past.

In this chapter we explore some of the computational models used to represent the concept of reputation in the realm of computer science and, eventually, we establish

clear reputation model requirements derived from the specific problems faced in VO formation cited in Section 2.5. We start with a critical view of the work that has been carried out on reputation in the context of electronic communities. Next, we provide a general introduction to the requirements for reputation mechanisms in on-line communities based on the desiderata outlined by Zachira and Maes [77], and then a more detailed discussion of the requirements for reputation mechanisms in the context of VOs. We next examine the current reputation models and provide a description of the weaknesses and omissions in relation to VO formation requirements. Finally, we lead to what needs to be done to meet the needs of VO formation tasks.

3.2 Reputation in Electronic Communities

In the first wave of mass Internet expansion, huge numbers of electronic communities such as electronic market places, chat rooms, discussion groups have proliferated. On-line communities have fostered interaction between people or businesses who may have never met. With an increasing range of interaction, risk also increases. For example, in the field of electronic commerce, it is particularly clear that when interactions are made between unknowns in on-line auction sites such as eBay, potential customers, especially buyers, feel uneasy, uncertain and insecure [60].

Prior to the Internet, customers decided on whom they could rely based on the information provided by Better Business Bureaus (BBB), which kept records of references of vendors including complaints, or based on their past personal experience and person-to-person gossip. These references formed the reputation of individuals and corporations.

The first step to lessen the risk involved in anonymous interactions in electronic communities, particularly in electronic auctions, was the emergence of computer systems based on a natural extension of the BBB approach. As their predecessors, these systems, known as *Reputation Systems*, collect, distribute, and aggregate feedback about participants' past behaviour. Before making on-line purchases, customers check an available record in the BBB system that contains the reputational status of the seller. After the purchase, in the case of misbehaviour, dissatisfied customers have the chance to transmit their complaints to the BBB itself. The BBB endeavours to settle the dispute; if it is not possible, future customers are warned in order to avoid interacting with that seller in the future [19].

Trust is a basic ingredient in electronic commerce, and buyers only interact with those customers who they think will not try to harm or deceive them. This approach meets the typical need for adding trust in the *on-line store* but it is not viable in direct interactions between private individuals. To oversee private interactions and encourage potential customers, a different *mechanism* needs to be devised. For example, in [52], Nowak presents the dilemma of cooperation faced by a population of individuals having the

option to help one another or not. The decision of cooperation is based on the internal *image* that individuals have of others. This *image* is represented as a numerical value and it can be accumulated (or decremented) for direct interaction among individuals. Image is a sort of simile of reputation because it is a numeric representation of an individual's opinion about another's attitude.

Electronic auctions have been commonly the first economic arena for the implementation of reputation systems. According to Resnick and Zeckhauser [59], marketplaces bring together large number of users who are essentially unknown to each other. Different from conventional partners, e-customers hardly take advantage of direct contact. From the provider's point of view, this modality of interaction could deeply and negatively affect exchanges. For example, in conventional exchange if customers are not satisfied with their buying, they then have the option of returning the merchandise and claiming back their money. However, they are deterred from taking similar action when the seller lives on the other side of the ocean.

To shelter users from the uncertainty of distance and anonymity, auction service providers resorted to centralised reputation systems. In general, these centralised reputation systems provide, for each individual, a reputation rating that can range from a simple *good* or *bad* flag to a full vector of quality characteristics, such as responsiveness, trustworthiness, speed of payment (as a buyer) and of delivery (as a seller). Carter [10] argues that the main, often unique, source of information used to build the ratings is transaction evaluation, consisting of a couple of ratings, provided by the buyer's evaluation of the seller and vice versa.

Electronic auctions need to enforce trust among strangers and reliability in order to strengthen commitments and responsibility. In long-term relationships, trust is built naturally. When individuals interact with one another over time, the history of past interactions informs them about their abilities. The expectation of rewards or retaliation in future interactions creates an incentive for good behaviour. By contrast, trust is much more difficult to build among strangers in electronic auctions. Strangers lack known past histories or the prospect of future interaction, and they are not subject to a network of informed individuals who would punish bad and reward good behaviour. It could be considered that the stranger's good name is not at stake. Given these circumstances, the temptation to deceive others outweighs the incentive to cooperate, since there are no future implications.

We can think of two levels of reputation. The first level is the personal level, which oversees direct interactions between private individuals, and is called *direct reputation.* The information collected individually from direct interactions is a global or averaged evaluation of a given interaction partner on the part of an individual. Direct reputation serves to identify friends and partners and to avoid enemies based exclusively on personal experiences.

Social reputation, on the other hand, is the second level of reputation and emerges thanks to the propagating of *direct reputation*. In particular, it proceeds from the personal level to the level of social propagation and from this level back to that of personal again. For example, after buying several items over the Internet, a buyer *b1* decides to propagate, within its community, its experiences with the seller *s1*. Buyers *b2, b3* and *b4* receive the opinion of *b1* about *s1* and include such opinion as part of their knowledge. Moreover, social reputation serves as a social control mechanism because targets will be aware not only of what people *think* about them (direct reputation) but also of what they *are* in the eyes of others (social reputation) [19].

3.3 Desiderata for Reputation Mechanisms

Zacharia and Maes in [77] outline the desiderata for reputation mechanisms in on-line communities as a way to influence socially the behaviour of the participants in the community. These desiderata establish the following:

1. *The reputation of any participant should never fall below the reputation assigned to a new participant.* This principle aims to avoid participants adopting a new identity, entering into the community, abusing their membership by cheating, and free riding without suffering the consequences of their acts.

2. *The reputation of a new participant should start with the minimum value allowed by the system and improve during the activities of the participant.* Zachira recognises that this requirement could be relaxed if alternative approaches allow participants to acquire reputation by other means such as membership payments.

3. *Disincentives for performing of fake transactions should be created.* This prevents participants from building their own reputation based on fake interactions with colluded partners. Zachira suggests a clever proposal that consists of keeping the evaluation of the most recent interaction when two participants happen to interact more than once.

4. *The opinion of participants with high reputation should have more relevance than others when the reputation of a participant is updated.* An obvious supposition considers that participants with high reputation will give more truthful ratings than others. At a certain level this could be true, but in most cases an individual may be an excellent service provider but a poor adviser. This privileged situation could be used by highly reputed participants to expel competitors from a competitive environment. For example, a well reputed seller could act as a buyer, and after interacting with his competitors, the seller could give a false opinion about his interaction partner in order to discredit him.

5. *Reputation should be able to quantify the subjective expectations of the participants.* Quantifying expectations of participants is a complex task that requires rich ratings that allow participants to express their subjective opinions.

6. *Reputation should be adaptable to changes in the participants' behaviours.* Reputation mechanisms act as a predictors of the future behaviour of a participant. So reputation values should be closer to the latest behaviour of the participants than their average performance.

3.4 Requirements for Reputation Mechanisms in VOs

The objective of this section is to delineate the requirements for building a reputation system in order to serve as a decision-making variable in the selection of partners, promote cooperation, produce trust and induce *good* behaviour in the members of a VO.

While the previous chapter discussed the need for reputation and its effects on the formation of VOs, it paid very little attention to the analysis of specific mechanisms for forming and communicating reputation in the VO literature. On the contrary, reputation systems have recently been a research topic in the realm of computer science. Here, the development of algorithms and systems has been oriented to manage feedback, from members of electronic trading communities, regarding past transactions with other members of that community [27].

Indeed, on-line markets have been the main target of current reputation systems, and by building *trust*, they encourage the exchange of goods and services among strangers [27]. Although VOs and markets are in essence different forms of coordination (VOs are a type of Network Organisation), both organisational models deal with coordination of independent entities and have to make judgements about the trustworthiness of potential partners [28]. Then, some of the general principles and guidelines of the extant reputation systems complemented by experimental studies in network organisations can serve as a starting point for detailing the requirements of reputation systems in VOs.

3.4.1 General Requirements

Resnick [60] affirms that a reputation system, in the context of electronic marketplaces characterised by buyer-seller interactions, must meet three challenges:

1. provide information that allows buyers to distinguish between trustworthy and non-trustworthy sellers;

2. encourage sellers to be trustworthy, and

3. discourage participation from those who are not trustworthy.

In other words, a reputation mechanism should indicate whether a seller (from now on called service provider) is *good* or *bad* with respect to the characteristics of the service provided, foster trust by ensuring that the behaviour of a service provider toward any buyer (from now on called consumer) becomes publicly known and finally it may, therefore, affect the behaviour of the entire community toward that service provider in the future [27]. To face those challenges, reputation systems require algorithms for eliciting, collecting, aggregating and distributing information from consumers and service providers' feedback about each other's behaviour.

Conte [20] affirms that there are several objectives for building reputation systems, and accordingly, several criteria to evaluate their performance. Conte identifies at least two general objectives: *economic efficiency* and *social efficiency*. The former aims to stimulate the exchanges, and the latter aims to increase the rate of satisfaction of users. Economic efficiency is evaluated by the volume of transactions and the level of prices, and social efficiency is evaluated by the number of contracts honoured and the rate feedback provision. Hence, the algorithms for collecting, aggregating and publishing reputation information depend on the objective of the reputation system.

According to the classification given by Conte, it is clear that the objectives of a reputation system for VOs are more aligned with social efficiency than with economic efficiency. This is because, as seen above, the role of reputation in VOs is to create an environment of *confidence* where partners can develop mutual trust and eventually cooperate to achieve a goal that in the long run will benefit not only the partners in the VO but the whole community.

Now lets discuss how the algorithms for collecting, eliciting, aggregating and distributing reputation should be built to meet the general requirements of a reputation system for VOs.

3.4.1.1 Eliciting feedback

A reputation system for VOs should rely on *private* monitoring of interaction outcomes (i.e. evaluation of direct experiences with a service provider) and third party feedback submission (i.e. interaction evaluations shared by others partners with the same service provider). Third party feedback is needed because of two reasons. First, when the interactions with another partner are scarce it is not possible to assign it a reputation based just on direct experiences [65]. It is in these situations when the information coming from other partners is useful. Second, and more important, the information provided by others indicates both how a partner is evaluated regarding its reliability, competences and trustworthiness, for example, and to what degree a partner is known in the community.

Third party feedback introduces two important new considerations. First, ensuring that sufficient feedback is provided and second, inducing truthful reporting [27, 60]. Several researchers [27, 60, 78] agree that in a market environment, voluntary feedback will be under-provided because of two causes. First, once feedback is spread in a community such as an electronic market, it becomes a public good that can *freely* benefit everyone. Second, in order for providing a feedback, the consumer must assume the risk of transacting with the service provider. Then, a reputation system for VOs must: (1) provide incentives for a partner to rate others and share their evaluations, (2) consider whether feedback is honest so that partners in a VO can differentiate the feedback given by less trustworthy partners from that given by trustworthy partners.

3.4.1.2 Collecting feedback

Reputation systems must specify the type, form and quantity of information they solicit, which may vary from rating transactions as *positive*, *negative* or *neutral*, or by using numbers. VOs require that: (1) the reputation information collected indicates the measures to evaluate the performance of the partners concerning the executed task or the claimed capability for executing a task. Additionally, as reputation is context-dependent [44, 48] (i.e. service providers may have more than one ability) thus (2) partners should be reputed in relation to the contributed service in the VO. Context itself is a multi-dimensioned concept [44, 65], and can include factors such as importance and utility of a service [45], quality, delivery time, price, etc. Finally, because VOs are organisations of *peers* that are equal in their *status* (i.e. there are no hierarchies) and are independent entities, and because there is no central authority that controls each peer behaviour, then (3) each partner should be capable of collecting and interpreting their own information.

3.4.1.3 Aggregating feedback

The process of feedback aggregation depends on both the kind of information to be used for representing reputation, and how often the provision of services changes over time and therefore so does its feedback. There are two approaches for aggregating feedback. The most common feedback aggregation mechanism acts as an impartial rating aggregator: it accumulates individual feedback and publishes an unbiased measure such as the sum or average of that feedback per asset exchanged or per service provider [26]. The second approach for aggregating feedback does not aggregate all the history of feedback. Instead, it disregards very old feedback in order to show a reputation value closer to the current performance of the individuals rather than their overall performance [77]. Dellarocas found that in dynamic environments where the quality of the services varies, simple aggregation mechanisms fail to induce cooperation [26].

Each partner in the VO possesses a true ability with respect to the services provided. At any given moment, these true abilities may not be well recognised by the members of the VO. In order to maximise its effectiveness and fulfil better the client's requirements, a VO needs to fully recognise each partner's true ability. Therefore, VOs require (1) an aggregation mechanism where reputation is to stand for the true ability of a service provider. This requires an adaptive mechanism to adjust the values of reputation to the true abilities rather than a simple accumulation of feedback. Additionally, (2) the aggregation mechanism should be time-sensitive in the sense that the aggregation of feedback depends on the time when the feedback is submitted.

Finally, when aggregating feedback it is important to know how *reliable* is the feedback and the reputation value computed. There are many elements that can be considered as determinants of reliability according to the feedback source. Some [77, 65] mention that the reliability of reputation values computed from direct experiences is affected mainly by the number of interactions and the variability of the feedback. A different approach should be considered when opinions are aggregated. In [75] it is suggested that a key element is the trustworthiness of the partner giving the feedback. This is based in the assumption that untrustworthy partners are more likely to submit false or misleading feedback and trustworthy partners are more likely to be honest in the feedback they provide.

3.4.2 Specific Requirements

Specific requirements of a reputation model for VOs come from two sources. An obvious source is the set of general requirements itself. These must be adapted not only to stimulate the formation of VOs but also to increase the rate of satisfaction of their members.

The second source relates to the particular disadvantage of VOs in creating flexible and reliable alliances, which are not addressed by the existing principles and guidelines for building reputation models in the context of e-commerce.

Therefore, in order to put into effect the general requirements for building a reputation system in the context of VOs, it is essential to give an exact description of what is required by VOs in terms of the following points listed in Section 3.4.1:

1. Provide information to distinguish between trustworthy and untrustworthy partners. Partners should identify who is a potential swindler to avoid interacting with it and who is an honest individual to interact with it instead. Information to make such decisions should be collected individually to provide each partner with a personalised model of the trustworthiness of their interaction partners. Moreover, partners can change their behaviour in relation to their individual goals. Hence,

it should be specified how individuals collect and manage their information to recognise trustworthy and untrustworthy partners.

2. Encourage participants to be trustworthy. A key disadvantage of VOs is implicitly inciting their members to abuse each other because there is no central trusted authority controlling each member's behaviour. Contracts offer one way to control VO partners, but they have problems when the tasks requested are not well defined. Hence, reputation in a VO must serve as a means of social control that enforces cooperative behaviour.

3. Discourage participation from those who are untrustworthy. Partners in VOs should be aware of the consequences of behaving dishonestly. Partners should not only consider immediate outcomes but also long-term outcomes of their acts. This is because an individual can deceive its partners in the VO until it is discovered, and then it can move to another VO without receiving punishment. Thus, a reputation model should include a means for addressing this. Firstly, it is needed to provide capabilities for disseminating reputation information and secondly, it is needed to exert control over agent behaviour through collective sanctions in order to discourage participation of dishonest partners.

In summary, the specific requirements are derived directly from the general requirements, but they provide a more detailed description of how to satisfy such requirements in a VO. Specific requirements are grouped in five categories:

- Non-functional. This category refers to the requirements for managing reputation information and includes challenges such as storage requirements, updating criteria and type of reputation (individual or group).

- Trust Enhancing. This category refers to the sort of information that reputation must provide for enhancing trust. This includes information about the competence, predicted behaviour and risk as properties of a potential partner.

- Sources of Reputation. This describes the minimal sources of information to be considered for computing the reputation of a partner.

- Control Enhancing. Castelfranchi [14] mentions that reputation is considered as a medium of *Social Control* that is intended to enforce some agent's actions to some desired behaviour.

- Social Influence. This category refers to the effects of individual reputation on the participants of a VO.

3.4.2.1 Non-functional Requirements

1. *Distributed reputation management.* As explained before, VOs do not depend on the presence of any centrally trusted authority. Moreover, individuals require to maintain their personalised models of the trustworthiness of others at a capability level so that they will be able to know which capability caused cooperation to fail and why [35]. By contrast, a centralised management of reputation offers a biased aggregation and perspective of reputation because it aggregates feedback making no distinction between the preferences of partners submitting feedback who may not coincide with the interests of the VO. As a consequence, personalised evaluations of risk and trust are hindered.

   Requirement. There is a need for distributed mechanisms that facilitate the partners in a VO to collect, store, manage and disseminate other's reputation in a personalised fashion.

2. *Dynamism.* VOs facilitate integrating multiple autonomous, diversely skilled partners under intense time pressures to create complex products or services [40]. Due to limitations in time and the intense task pressure, VOs require that their members develop fundamental mutual trust fast. The partners do not have time to engage in extensive interactions to learn about each other.

   Requirement. Partners should be able to quickly use a reduced number of interactions to estimate the reputation of a partner and at the same time take partner selection decisions without having a significant impact, in terms of time consumption, on the formation of a VO.

3. *Adaptability.* VOs operate under high levels of demand uncertainties generated by unknown and rapid shifts in consumer preferences [40]. Demand uncertainty creates changes in the structure of the VO, which is forced to adapt itself by reallocating tasks or redefining them. In these circumstances, organisations must be able to learn how to work together to accomplish the VO goal, their respective skills, and each other's goals. Their learning feeds into periodic evaluations of the VO, which, in turn, leads partners to make adjustments to their relationships and identify when changes in the efficiency of the partners is due to the adaptability process or due to abusive behaviour [40]. The rapidity with which partners *discover* this behaviour is crucial.

   Requirement. Reputation must be updated dynamically to adapt the values of reputation towards true quality of the service. This suggests that the updating process should be a *learning* process about other's true abilities, that captures the observed performance through the reputation of the partner. This leads to discarding updating methods that diminish the impact of strategic changes in partner behaviour that intend to milk *high* values of reputation by intentionally deteriorating the provision of the service.

3.4.2.2 Trust Enhancing Requirements

1. *Predictability.* The behaviour of each partner in the VO usually offers clues to the others about its capabilities and *intentions*, so it is possible to make predictions about its future behaviour. The main objective of *predictions* is detect any misconduct of the partner early enough, so that the VO can take necessary steps to protect itself from adverse effects of partner misbehaviour.

   Requirement. The reputation system must provide information to predict the future performance of a partner and eventually the risk involved of interacting with it. That is, based on a partner's previous performance, reputation must provide an indication of its future performance and willingness to accomplish a task. This is a rich source to build trust in VOs because it allows to recognise possible intentions and the level of commitment to the global task, and finally estimate the risk associated with a certain partnership.

3.4.2.3 Required Sources of Reputation

1. *Expectations.* As mentioned before, because of demand uncertainty it is not possible to fix the terms and conditions of tasks because the attributes of a new service cannot be known in advance or may change over time. Hence, the absence of an exact definition of the new product or service that specifies its attributes in detail and how it could be delivered makes it difficult to evaluate the performance of a service provider and at the same time to estimate its reputation.

   Requirement. The reputation system should consider implicit sources of information, such as expectations, in addition to explicit ones founded in contracts, to compute the reputation of partners in VOs.

2. *Group Reputation.* The reputation of a partner may be estimated from two points of view. Firstly, partners evaluate each other according to their interactions and secondly, the end-consumer of the good or service evaluates the VO performance as a whole.

   Requirement. The reputation system should include the evaluation of the end-user in order to compute the reputation of the partners in the VO.

3. *Witnesses and Neighbours.* Apart from direct experiences, the reputation of a partner also depends on both the information given by previous interaction partners (i.e. *witnesses*) and the reputation of its previous partners (i.e. *neighbours*).

   Requirement. The reputation system should provide mechanisms for collecting and aggregating reputation information from third party sources such as witnesses and neighbours.

3.4.2.4 Control Enhancing Requirements

1. *Dissemination.* The potential of reputation as a mechanism of *social* control relies on the capabilities of the reputation system for disseminating information about the behaviour of partners.

   Requirement. A reputation system for VOs requires a mechanism that allows reputation information to be disseminated *efficiently* in the *social network* [21]. In this context, efficiency refers to specifying the type of information that must be disseminated and the use of the links created by the partners through previous relationships.

2. *Collective sanctions.* Reputation exerts its control through *collective sanctions.* Collective sanctions encourage partners in resisting incentives for short-term opportunistic behaviour and range from access resource restrictions to ostracism (exclusion from the society for short periods or indefinitely) [21]. For example, when there is an opportunity of continuing with the same VO structure, the fact of seeking an alternative service provider for replacing the actual one may be considered as a punishment. On the other hand, maintaining the same partners may be considered as a reward.

   Requirement. VOs require mechanisms for enforcing collective sanctions providing punishing to those who do not punish deviants. This should be done by hurting the reputation of those who recommends someone whose performance does not meet the expectations [40].

3.4.2.5 Social Influence Requirements

1. *Dependence.* The individual reputation of a partner may directly affect the reputation of the rest of the members in the VO and vice versa. This is a completely different mechanism of propagation of reputation in the sense that the reputation of a partner not only depends on its actions but the rest of the members' actions. Hence, the way that an entity is seen by others in an electronic community, apart from being affected by its reputation as a service provider, is influenced by contagion of its partners'reputation.

   Requirement. The reputation system for VOs must consider this effect when the reputation of a potential partner is computed to take decisions about its possible integration into the VO.

3.4.3 Dynamics of reputation in VO

From the above requirements, the reputation of a partner should change over time and can be depicted as a curve with the following properties:

1. A continuous curve describing the variations of reputation values at each instant in a finite period of time.

2. The values of reputation should not increase or decrease infinitely, instead they should be limited by a minimum and a maximum value [77]. It has been demonstrated that when reputation values are allowed to increase ad infinitum, as is the case of eBay, buyers and sellers can cheat 20% of the time but still maintain a monotonically increasing reputation value.

3. As discussed in Section 3.4.1.2, the values of reputation should measure the perceived capability of a partner with respect to a particular ability (i.e. the provided service).

4. The values of reputation should not remain static unless the partner maintains the provision of the service in a constant way and under the same circumstances (i.e. without changes in the client preferences, product demand, importance of the task, etc.).

5. The values of reputation should show variations when: the service is not provided with the same qualities (price, quality, delivery time) or when the service is not provided for a period of time. In the latter case, reputation values should show a gradual decrement as the time passes and the service is not provided.

6. As discussed in Section 3.4.2.1, when changes in the behaviour of a service provider are detected, the values of reputation should show those changes accordingly, considering the magnitude of variations of behaviour observed in the past and the partner's tolerance for uncertainty and risk. Partner's previous experiences may serve to magnify or enhance the reputational information; others may serve to suppress the content received [29]. For example, if the service provider has maintained the variation in qualities of its service inside the interval $-\delta, +\delta$ in, lets say, more than the 80% of the interactions, then the change in its reputation should be gradual. Otherwise, the reputation value should immediately reflect those changes. The updating process is essential in building reputation so it is discussed in more detail above.

7. As discussed in Section 3.4.1.3, reputation values must be associated with a measure of *reliability* to indicate the degree of certainty of calculated reputation. As the partners gain experience and the variation in the reputation values decreases, this reliability increases.

The goal of the each partner estimating the reputation of a service provider in a highly dynamic environment such as a VO is learning in order to develop quick adaptability to that environment. That is, if the learner is too slow to adapt, its partner may be exploiting it with impunity. But, if the partner presents totally reactive behaviour,

then it can immediately throw away almost the whole history of feedback. This can be counter-productive if the environment is noisy or the partner misreads the feedback, since once thrown away, the past feedback cannot be brought back. This adaptability is called *disposition* [35, 45].

Disposition determines how reputation is updated after interactions. Optimists increase the reputation of a service provider more than the pessimists after a successful interaction, and conversely, after an unsuccessful interaction pessimists decrease the reputation of a service provider more than optimists.

Partners must decide how much past information to take into account for assessing the reputation of an interaction partner, which should be affected by the following factors: (1) the fluctuation of reputation values observed in the past; (2) the current value of reputation; and finally (3) the extent of the partner's optimistic or pessimistic disposition.

3.5 Weakness and omissions in current Reputation Models

The design and implementation of on-line reputation systems has so far been the research domain of computer scientists [1, 45, 48, 59, 63, 65, 77]. The emphasis of past work in the area has been on developing algorithms and systems for collecting, aggregating and extracting useful information from sets of user ratings that are used as signalling quality and quality control mechanisms in on-line trading communities.

The Internet has contributed to the development of reputation systems due to it fosters interaction between people or businesses who may have never met. With an increasing range of interaction, risk also increases. For example, in the field of electronic commerce, it is particularly clear that when interactions are made between unknowns in on-line auction sites such as *eBay*, potential customers, especially buyers, feel uneasy, uncertain and insecure [60]. In this domain, researchers and practitioners look at the ways in which information about all transactions mediated in an electronic marketplace can be aggregated and processed by the market makers or other trusted third parties in order to help on-line buyers and sellers assess each others trustworthiness. According to Resnick and Zeckhauser [59], marketplaces bring together large numbers of users who are essentially unknown to each other. Particularly in eBay transactions, auctioneers resort to *centralised* reputation systems to shelter users from the uncertainty of distance and anonymity.

In general, these centralised reputation systems provide, for each individual in the system, a reputation rating that can range from a simple *good* or *bad* flag to a full vector of quality characteristics, such as responsiveness, trustworthiness, speed of payment (as a buyer) and of delivery (as a seller). One of the most prominent centralised reputation

systems is SPORAS [77]. This system presents some advantages over eBay, which are briefly discussed later in this section.

In recent years the field of multi-agent systems has been actively researching on-line reputation systems as a technology for building trust and inducing good behaviour in artificial societies of software agents. The distributed nature of multi-agent systems has led to the development of *decentralised* reputation systems [63, 1, 65, 48]. There is a lot of research on decentralised reputation systems. Here we just mention that work that is most relevant in our perspective for VOs.

3.5.1 General Weaknesses

Before starting with our discussion, it is important to note that most of the current reputation systems have been designed for encouraging trust in the relationships between sellers and buyers in the electronic commerce domain. These reputation systems aim to reach an *economic efficiency* [20] by stimulating the exchanges of well defined assets between sellers and buyers but do not take into account the fact that the buyers and sellers are selfish and therefore will not share information unless some benefit can be derived from doing so. In this context, buyers may benefit from others's information but without assuming the risk of transacting with a deceitful seller. On the contrary, as discussed before, VOs rely on reputational information not only for enabling buyers to choose the best sellers in the system but for controlling partners' behaviour and establishing cooperative relationships between partners to achieve a common goal.

Now, although there are other reputation systems, mainly those in the realm of multi-agent systems [49, 64, 65, 69, 76], that have broader aims beyond promoting buying and selling transactions, they do not motivate the use of reputation by some agents to elicit good behaviour from other agents. This is a key requirement for VOs that rely on reputation for imposing *collective sanctions* against bad behaviour on the part of the partners.

Finally, there are some aspects that are exclusive to VOs, and consequently, they are not taken into account such as (1) the impact of the VO's reputation as a whole on each one of the partners (*contagion*), (2) the necessity for building trust rapidly under time pressure situations, and (3) the adverse conditions for establish trust-based relationships where the partners may have short-term incentives for cooperating.

The purpose of this section is to discuss extant reputation systems in order to detect their weaknesses and omissions to cope with the requirements of VOs.

3.5.2 Weaknesses and Omissions in Reputation Systems

Reputation systems can be classified according to the way they collect and aggregate the reputation information: *centralised* and *distributed* reputation systems.

A centralised reputation system is characterised by a central management (i.e. the processes of collection, aggregation and displaying) of reputation information controlled in many cases by a trusted third party or *intermediaries*. Rates, or evaluations of transactions, provided by users indicate the level of satisfaction in a range from *unsatisfactory* to *satisfactory*. Here, reputation is conceived as a global value that starts from a minimum value for new users and it is built up and updated over time to reach a maximum value.

There is a lot of research on distributed reputation systems, particularly for MAS. Here we mention the most relevant work. The relevancy of these reputation systems relies on the use of *social networks* or *trust nets* for providing more personalised reputation values through of the retrieval and aggregation of rates from a selected group of referees. Similar to human societies, this assumes that agents are related to each other through either roles that interconnect them or communication links established between one another. These links create networks, where agents can transmit information about each other. Information takes the form of a trust rate (e.g. good or bad, seller delivers late, or bad quality services) as explained in the previous section. Such a rate could be shared by the different nodes of the social network, thus giving rise to the concept of reputation.

3.5.2.1 SPORAS

Zacharia and Maes in [77] present SPORAS, which is a *centralised* reputation system that establishes reputation for users in an on-line community (for example chat rooms, auctions or newsletters groups), based on the aggregation of *rates* given by users after each transaction.

Reputation in SPORAS aims to predict the future performance of the users. In order to make accurate predictions using a small computational space, a recursive and adaptive algorithm for updating reputation is used. The algorithm is recursive in the sense that reputation value is calculated continuously using the previous value of reputation; and it is considered adaptive because it reinforces or weakens the reputation value depending on the latest rates obtained. Then, this aggregation method allows newer rates to count more than older ones. SPORAS, however, gives newcomers low initial reputation values and therefore reduces their chances of being selected as a possible interaction partner.

SPORAS also provides a measure of the reliability of the reputation values that indicates the accuracy of the estimated reputations based on the reputation deviation due to inactivity, variations on the behaviour or misbehaviour.

Of special interest in the context of the present discussion is the policy that SPORAS takes for updating reputation based on an adaptive algorithm that corrects the reputation of a user based of its latest rates received. However, there are some features that we consider as weaknesses:

1. SPORAS is a centralised reputation system, and it is not viable for VOs where partners need personalised reputation values calculated from assembled rates of those they trust already rather than those they do not know. Moreover, mediators are designed and operated by parties whose interests may sometimes diverge from those of the electronic community. One mediator per VO may be suggested, but this brings more problems because when the organisation is disbanded, partners need the information centrally collected to make their own decisions about reuniting with former partners or when they are asked for references about former partners. Hence, it implies another problem: to locate all the mediators from past VOs who keep reputation information about a former partner.

2. Reputation in SPORAS is seen as a single concept, and there is no notion of the context dependency required in VOs (see general requirements in previous section). Partners may have different values of reputation for different services they are capable of providing, including that they require a reputation as reputation referees.

3. This is a special case of the above point. Although the assumption made in SPORAS to make reputation values dependent on the reputation of the entity who is providing a feedback is correct, it mixes two different dimensions of reputation. While a user can be reputed as completely unable to cheat on deals, nonetheless that same user may be a bad evaluator of other users. That is, being an excellent service provider does not mean being an honest evaluator.

4. SPORAS does not consider the impact of social relationships among the users in a virtual community. For instance, reputation is not used for enforcing good behaviour, as discussed in 3.4.2.

3.5.2.2 REGRET

REGRET is a reputation system developed by Sabater and Sierra [64, 65] that adopts a sociological approach for computing reputation in societies of agents trading well defined products inside an e-commerce environment. Although REGRET provides a very simple method for aggregating rates (or *impressions* that are the result of evaluating direct interactions) based on the weighted sum of the impressions (more relevance is given to the recent ones), its major contribution is the vision of reputation through of three dimensions. These dimensions are called the *individual dimension*, *social dimension* and *ontological dimension*.

REGRET emphasises both individual and social components of social evaluations. That is, whereas the individual dimension is the effect of past experience with a given agent, the social dimension refers to reputation inherited by individuals from the *groups* they belong to. Here, groups of agents are not reunited intentionally in order to solve a problem cooperatively, instead they are formed implicitly around the agent evaluated (or *target agent*) or the agent evaluator (or *source agent*). The first group is formed by strategically gathering the agents around the source agent and that had a trade relationship with the agent in question in the past. These agents are called *witnesses*. The second group is formed by the agents that are in the neighbourhood of the target agent, that is, the agents (or *neighbours*) that are related with the links created through interactions. Finally, the third group is formed by those agents that have a similar *role* within the system.

Thus, there are three measures of social reputation: (1) experience with other members of the group (i.e. neighbours) to which the agent evaluated belongs, (2) how the other members themselves evaluate their fellows (i.e. witnesses), and (3) what the members of the group (i.e. witnesses) think about the other group (i.e. neighbours). Hence, this system contributes an interesting model of group reputation, which has the merit of evaluating agents as entities within a given social structure.

Additionally, the ontological dimension allows agents to combine different types of reputation to obtain new types, but more important, agents can share this ontological representation of reputation to understand each other's impressions (e.g. a print service provider being good might imply good printing quality for one agent, but may imply high printing velocity for another). However there are some problems with this model.

1. VOs require the reputation of a partner to be computed using both explicit (issues detailed in contracts) and implicit information (expectations). However REGRET ignores the importance of the expectations while assessing the reputation of a partner. REGRET is based on contracting, which means that the evaluation of the performance of a partner is made by comparing the issues and its minimal values specified in the contract against the current characteristics of the service provided. However, as mentioned before in Section 3.4 this fixed set of characteristics that could define a service may not be presented a priori before the formation of a VO. It could be argued that contracts provide expectations, and this is in part true, but such expectations are limited by the contract itself.

2. As discussed earlier, VOs require to a certain extent that the reputation of a partner is assessed in a *reactive* form to detect possible opportunistic behaviour. However, REGRET's main idea consists of emphasising the freshness of information. Computations in REGRET give a *fixed* high relevance to recent rates over the older ones according to a time dependent function, and, moreover the rates

are aggregated in a way (as shown above) that can be sensitive to noise since they are simply summed.

3. VOs require that reputation be assessed swiftly in order to detect misbehaviour. REGRET, on its part, requires a minimum number of interactions to make correct evaluations of reputation but it is likely that partners will not interact the minimum number of times to provide a reliable reputation value.

4. REGRET does not handle the problem of lying (strategically) among agents. Rates are obtained in a cooperative manner rather than in a competitive environment.

5. Finally, as mentioned at the beginning of this discussion in Section 3.5.1, REGRET does not motivate the use of reputation to foster good behaviour.

3.5.2.3 Yu and Singh Evidential Model

Yu and Singh [76] present a model of trust where agents assist their users and collaborate to find trustworthy parties to deal with, or untrustworthy parties to avoid, either based on their prior interactions or on testimonies given by trustworthy agents called *witnesses*. In their model, Yu and Singh establish a range of values to rate the quality of a service after an interaction. These rates are named *evidence* and they are used to estimate the level of trustworthiness of an agent.

Yu and Singh consider the problem faced when there is a lack of information. They propose using the theory of Dempster-Shafer for combining distinct beliefs based on different bodies of evidence. The context is the following: an agent may receive good or bad rates (+1 or -1) about another, but when an agent receives no rates (good or bad), how should it classify this case of no information? In Yu's work that the lack of belief (or disbelief) in an hypothesis implies that all beliefs have the same probability of being true, which is referred to as a state of *uncertainty*. For instance, if there is not enough evidence to believe that an agent is dishonest, this does not imply that such an agent is honest; instead, it means that its real behaviour has not been revealed yet. Dempster's rule allows the combination of beliefs obtained from various sources (saying an agent is trustworthy, untrustworthy, or unknown to be trustworthy or not) to be combined so as to support the evidence that a particular agent is trustworthy or not.

Another contribution of Yu and Singh is the idea of *referral networks* to identify reliable witnesses that could provide trusted information about a potential partner. Referrals can be seen as pointers to other sources of information. Agents in a referral network collaborate in two ways with each other. Firstly, an agent can propagate upon request its own experiences or *local beliefs* about a potential partner if there were any prior interactions between them; otherwise it will point to alternative sources of information called *referrals*. A series of referral networks is named a *TrustNet* and it is dynamically

built by incorporating the distinct referrals about the same potential partner. Yu and Singh present a method to represent a TrustNet and then provide an algorithm to combine the testimonies of the witnesses located at the end of each referral network with the local beliefs in order to evaluate the trustworthiness of a potential partner.

Yu and Singh's model presents some weaknesses:

1. Their model does not deal with the possibility that the agents may lie about the rates of another agent. They assume all witnesses are totally trustworthy. However, an agent could find some benefit lying about its rates of a partner if it is able to discredit others such that it appears to be more reliable than them.

2. In their model, evidence is timeless, which brings some problems. For instance, partners may act opportunistically at some time intervals so as to maintain an equilibrium between positive and negative evidence. However, as information is aggregated without considering the time when the interaction was evaluated, it is difficult to detect when such behaviour is going to happen or under what circumstances.

3. Similar to REGRET, their model does not motivate the use of reputation to foster good behaviour.

4. VOs require that the reputation of a partner reflects the impact of the actions of the other members in the organisation. Although, Yu and Singh's model considers the social effects of reputation based on the references retrieved in the TustNet, it does not take into account the effect of VO reputation as a whole over the members of the VO.

3.5.2.4 Schillo's Probabilistic Model

Schillo et al. [69] propose a probabilistic model of trust oriented to decide *with whom* to interact rather than *how* to interact, where agents are evaluated after a number of interactions by two parameters: the probability for honesty/dishonesty and for altruism/egoism behaviour. In order to determine if an agent is honest/dishonest based on direct interactions, Schillo's model proposes a simplistic approximation based on the probability calculated from the number of events an agent behaved honestly/dishonestly and the total number of interactions.

In their reputation system, Schillo et al. propose that agents choose their partners based on testimonial evidence and their own observations. Agents create a model of other agents behaviour by gathering data on their past interactions from two sources. When their own observations are not sufficient, they take into account the testimonies of witness agents. Such witnesses are mainly those with whom they have interacted

before. As the information from witness agents could be false, agents also evaluate the behaviour of the witnesses. Even though Schillo considers that the information gathered from witnesses is truthful, he argues that witnesses might hide information to bias others' decisions by transmitting only dishonest and egotistic behaviour.

By using the evaluation of their own observations, agents can infer how often the witnesses have been lying and have a better approximation of the credibility of the agent that is evaluated.

Schillo, similarly to Yu's model, proposes a graph-like structure called *TrustNet* to store the agent's observations and the testimonies that the witness communicated. Each agent stores this graph where nodes represent agents. The owner of the net is represented as the root node. The edges carry information on the observations that the parent node agent communicate to the owner of the net about the child node agent.

Schillo's model present some weaknesses and omissions:

1. Similarly to Yu's work, in Schillo's model information is timeless, and partners in a VO may act opportunistically, and because information is aggregated without considering the time when the interaction was evaluated it is difficult to detect when such behaviour is going to happen under which circumstances.
2. There is no context concept, and VOs require that reputation indicates the capabilities of a partner in a specific competence or context.
3. The influence of partners reputation on the others is not considered.
4. As in the above models, there is no incentive for good behaviour.

3.6 Conclusions

To sum up the main points highlighted in this review chapter, we start by saying that reputation is the key to understanding cooperative interactions in open environments where a large number of individuals who are essentially unknown to each other, and may not share the same objectives, are brought together. In particular, reputation enables participants of electronic communities to take decisions when there are uncertainties about the trustworthiness of a potential interaction partner and, since the coexistence of self-interested and even competitive partners in general inspires mistrust, reputation must be considered to explain both the behaviour of self-interested participants that interact with other partners of similar characteristics, and the attitude of potential partners towards the reputed participant when its past behaviour is made public.

Reputation systems have been applied with great success in different domains. Electronic commerce, particularly electronic auctions, and MAS have been the main fields

of development and implementation of reputation models. Their aims are very diverse, from encouraging well defined asset exchanges between strangers, to promoting cooperation in an environment of self-interested agents. We recognise that current reputation systems have interesting advantages over those developed in the field of electronic auctions such as eBay feedback system.

At personal interaction level, reputation systems endow individuals with information about the average behaviour of fairly and well known individuals in order to decide if it is worth establishing an interaction relationship with a specific individual. Reputation is generally evaluated from the opinions of an individual's behaviour obtained from the history of direct experiences and testimonies that are solicited from elsewhere in the environment. However, these systems, and the models upon which they are based, are argued to be applied directly for estimating the reputation of future and current partners in VOs, which are highly dynamic and heavily dependent on their members for their success.

A key problem with existing reputation models is precisely the way in which reputation is evaluated, which is by using an incremental method to aggregate the history of evaluation of experiences. The underlying assumption of this approach is that the communities where individuals interact are relatively static: there is little variation in their members, which establish long term relationships in order to build their reputations steadily, based on relatively predictable interactions. In reality, one of the failings of this assumption is that it holds only in relatively closed environments; as scale and dynamism increase as is the case of VOs, the incremental model of reputation simply fails. Although some research has been done to provide a dynamic approach for modelling reputation and meeting the requirements of groups of cooperative individuals, such as VOs, more is needed.

Another important omission is the fact that current reputation systems do not motivate the use of reputation by some individuals to elicit good behaviour from other individuals. This is particularly important, because VOs require that reputation operates as a *social control* mechanism, first identifying dishonest partners in the society and second imposing collective sanctions on them in order to encourage VO partners in resisting incentives for misbehaviour.

By taking this perspective of personal interaction, many of the gaps in current reputation models, and the reputation systems based on such models, can be filled and, more important, a model for selecting partners based on the characteristics of the potential partner can be proposed.

Chapter 4

A Direct Reputation Model

4.1 Introduction

In a VO, partners elect to cooperate for the achievement of a common goal by contributing with a specific task that either is a part of, or helps to reach, a common goal. In order to complete that goal successfully, VO must consider the nature of those it may cooperate with and the likelihood that they will execute their tasks correctly.

Trust offers a mechanism for modelling and reasoning about the capabilities and willingness of potential partners in order to assess perceived risk in interactions. *Reputation* and *personal experiences* with the subject of trust are the main sources of information to determine the *degree of trust* that a trustor can give to a trustee. In the proposed model, reputation fulfils two functions in a VO. Firstly, reputation and experiences are used to identify potential partners for cooperation. Secondly, reputation is used as a medium to *intensify* the partner's motivation towards cooperation because partners are aware of the opinion others have about them and as a result of that opinion they can be accepted or avoided as a partners in a VO.

In this chapter we introduce a model of trust based on reputation information, and the definitions of the terms used. This reputation model meets some of the requirements discussed in Chapter 3; specifically we focus in those described in Section 3.4.2.1 and Section 3.4.2.2. Though important, the rest of the requirements are not relevant in this model and they are not covered in this thesis. We start defining mathematically the concepts of reputation, impressions, trust and risk. Then, we discuss two possible approaches for updating reputation: subjective probability and reinforcement learning. Next we describe the formulae used in our model for updating reputation, and computing trust and risk based on reputation values.

4.2 Definitions

4.2.1 Virtual Organisation

Norman and et.al. [51] describe a Virtual Organisation (VO) as *a temporary alliance composed of a number of autonomous entities or partners (representing different individuals, departments and organisations) each of which has bounded problem-solving capabilities and limited resources at their disposal.* Partners in a VO come together to share skills or core competences and resources in order to better respond to customer needs or business opportunities, and whose cooperation is supported by computer networks.

4.2.2 Reputation

Liu et al. define the reputation of an agent as *a perception regarding its behaviour norms, which is held by other agents, based on experiences and observations of its past actions* [44]. This definition covers the principal characteristics of reputation such as the sources of reputation (experiences and observations of its past actions), the subjectivity of reputation (perception) and the fact that reputation of an agent is created from the point of view of other agents so it is unable to decide its reputation.

However, it ignores the fact that reputation is the effect of propagating subjective evaluations of experiences inside the community where the VO is embedded. The information propagated about an agent not only affects the attitude of the community towards it but influences the behaviour of the agent reputed. In other words, once subjective evaluations are propagated, reputation operates in two ways, (a) as a warning label, to protect others from transgressors or to attract potential partners, and (b) as a social control, punishing those agents that suffer from bad reputation by excluding them from any form of cooperative association. To accentuate this *social influence* of reputation, we define reputation as follows:

Reputation of an agent is *a perception regarding its intention and competences, which is held by other agents through the formation and dissemination of subjective evaluations which are based on experiences and observations of its past actions.*

Evaluations are subjective because the same behaviour can cause different evaluations in different individuals or situations. Here, these evaluations are called *impressions.* From the definition, the observed behaviour of others is collected through: (i) direct experiences, with interaction histories serving as a strong reference for reputation evaluation or (ii) via the testimony of others, known as recommenders. On the basis of the source of reputation, two concepts of reputation may be derived, namely: *Direct Reputation* and *Social Reputation.* Direct Reputation is a key concept in this thesis and it is defined in detail further in this section. Although important, the concept of Social Reputation

lies beyond the scope of our research and it is not defined, and we only make reference to it as another source of reputation different from direct reputation.

An important characteristic of reputation is that it is dynamic and changes over time. Hence an algorithm for updating reputation over time will be specified later in this chapter.

4.2.3 Direct Reputation

Our model of direct reputation is based upon SPORAS, which is discussed in the literature review in Section 3.5.2.1. SPORAS was chosen because some of its design guidelines are consistent with our specific requirement such as *dynamism*, *adaptability* and *predictability*, described in Section 3.4.2. In particular, since SPORAS is an adaptive reputation model that updates the reputation values after each interaction, removing the effects of *obsolete* data in some manner, it is ideal for environments where the behaviour of VO members is changing through the time because the relationships among them are themselves changing as a function of their interests and goals.

Direct reputation (DR) is defined as the weighted average evaluation that an agent (source) makes of another's competence (target), and tells the extent to which the target is *good* or *bad*, with respect to a given behaviour or action. Direct reputation is context-dependent and an agent is reputed according to the service provided. For example, an agent may be well reputed as a printing service provider but bad reputed as a file storing service provider.

VOs provide an environment where agents may offer the same service with different qualities for different reasons such as demand uncertainty or as a result of dishonest behaviour. In this sense, we adopt the ideas of Shapiro [70] expressed in his analysis of the economic effects of reputation in such environments. Shapiro proved that, the most efficient way to estimate a seller's reputation (i.e., the way that induces the seller to produce at the highest quality level) is a time-discounted average of the recent ratings.

We are basing our algorithm for computing direct reputation on an approach which approximates Shapiro's desiderata. Hence, direct reputation is computed as the average of *impressions* received within the most recent time. This is supported by Shapiro's work, where he proved that efficiency is higher: (1) the higher the weight placed on recent quality ratings, and (2) the higher the discount factor for older ratings. In this work, the impressions given in the recent time are contained within the time window:

$$W = [t - \epsilon, t], \tag{4.1}$$

ϵ defines an interval in the time that limits the set of interactions and whose impressions are used to compute a direct reputation value. As above mentioned, impressions are weighted from 0 to 1 to indicate the notion of importance of the impression in relation to the rest for calculating reputation. Taking only the most recent impressions is equivalent to using an average calculation where weights are non-zero for impressions received within the recent time and 0 otherwise. The direct reputation values vary in the range of [0,1] and are used only to represent comparative values in a continuous space from bad reputation (values near zero) to good reputation (values near one). The direct reputation of i in the perspective of j in the context k is represented as:

$$DR_{ij}^{k} \in [0,1].$$

4.2.3.1 Reliability of Direct Reputation

Reliability is the attribute of the measure of reputation, which indicates after evaluating repeatedly through impressions the behaviour of a partner, the degree of variability of the reputation values. Furthermore, reliability indicates if the number of interactions is not enough to make a correct judgement about a service provider. In general, reputation should not change dramatically if it was calculated a minimal number of times to establish the abilities of a service provider. Mathematically, reliability must be estimated as the average of the deviations in the predictions, that is, the errors in the estimation of reputation. Following the ideas of Zacharia and Maes [77], reliability can be computed as follows:

$$ReDr_t \in [0,1],$$

$$ReDr_t^2 = \left(\lambda * ReDr_{t-1}^2 + (imp_t - DR_t)^2\right) / \lambda^t, \tag{4.2}$$

where $ReDr$ is the reliability of the direct reputation value computed at time t, imp is the impression obtained after an interaction and DR is the previous value of reputation. Reliability of direct reputation, $ReDr_t$ must be parametrised by i,j and k. However, these parameters were omitted in the notation in order to make it more readable. Finally, $0 < \lambda < 1$ is a forgetting factor that allows the removal of the effects of *obsolete* values of reliability by ensuring that the initial conditions, when there is a small number of values of reputation, are quickly discharged. Moreover, λ ensures that all data are not weighted equally over time, and the most recent impressions have more weight than the older ones. So, Equation (4.2) is a recursive estimation algorithm that takes into account the possibility of behavioural characteristics of an agent changing through time. In such cases, the inference is that the interactions are themselves changing through the passage of time as a function of the changing socio-economic conditions and preferences.

As can be seen, the values of $ReDr_t$ are constrained to be in a particular range. This is because by comparing such values it is possible to determine when the reputation of an agent is reliable and when it is not. Hence, in order to be consistent with our earlier definitions, $ReDr_t$ is constrained to be in the range [0,1].

Now, bearing in mind that *ReDr* is broadly a measure for indicating the *variability* of the reputation values computed, values of *ReDr* near to 0 mean low variability or high reliability and values of *ReDr* near to 1 mean high variability or low reliability. Hence, a value of *ReDr* near to 1 can mean either that the agent has not been active enough to provide an accurate value of reputation, or that the agent's behaviour has too much variation, or even that the agent's behaviour is too disparate to make an appropriate evaluation.

4.2.4 Impression

We define an impression as an evaluative opinion that is formed by any entity (individual, organisation, etc.) based on a discrete experience with other partner, coupled with the partner's performance. An entity can get an impression by evaluating either its observations or its direct interactions. Computationally, an impression is the value assigned to the service provider by the service consumer during a transaction that indicates how close are the services provided by an agent i to the expectations of the agent j requesting the service from agent i. An impression is related with a dimension d that describes just one of the qualities of the service as required by the agent j. For example, a partner can get different impressions for its efficiency or the quality of its services. The group of dimensions needed for evaluating the whole performance of a service provider is denoted by the set of enabling qualities Q_{ij}, and is context-dependent. For example to evaluate an agent in the context of a printing service, two dimensions may be taken into account among others: printing quality and rapidity. Mathematically, the impression appears as follows,

$$imp_{ij}^{d} \in [0,1],$$
$$Q_{ij} = \{d \in k | k \; is \; a \; context\}. \tag{4.3}$$

where i is the service provider whose interaction with the service consumer j left in it the strong impression *imp* in relation to dimension d, and Q_{ij} is the set of dimensions for evaluating a service provider in the context k. The numbers used for impressions are merely reference values for making comparisons. Each consumer establishes a personal threshold of *acceptable* values for the dimension d evaluated. This personal threshold may be based on:

- the agreed values of a contract, when the interactions consumer-service provider are fixed by contractual terms; or
- the values that constitute a standard, when standards are available to indicate the permissible values of a particular dimension; or
- the values obtained empirically, when the consumer has previous experience consuming a particular service and it is available to estimate optimal values for the dimensions involved.

Once a personal threshold of acceptable values is established, this threshold is compared with the actual values of each dimension after providing a service. Then, lower values than the threshold confer insufficiency or inadequacy for fulfilling the requirements of a service, and greater values than the threshold merely express that an agent is useful for accomplishing a task. Finally, values in the threshold indicate that the service provider is good enough for accomplishing a task.

When evaluating a service it must be indicated how the increment of the value of the quality affects the value of the impression. For instance, when the evaluation refers to the number of pages printed per minute, if the number of pages increases, the impression also increases, while when it refers to the delivery time, the impression value decreases. Considering a numerical example, a service provider i is asked for a service with a quality of 0.4, but it delivers the service with a quality of either 0.5 or 0.3. Because increments in the quality are considered positive, in the first case (quality = 0.5) i will have an impression value of 1.0 in the context of quality of service in the opinion of j. However, in the second case (quality = 0.3) i will have an impression value of 0.75 because the difference between the quality required and the quality received is of the 75%. That is $imp_{ij}^{QoS} = 0.75$.

4.2.4.1 Importance of impressions

The importance of an impression is the significance or *weight* that the impression has in relation to the total of impressions obtained from direct interaction with the same partner in a period of time T. Mathematically, the importance w of the impression obtained at the instant t is a numerical value in the range [0,1] but it has to satisfy the following property:

$$\sum_{t=1}^{T} w_t = 1. \tag{4.4}$$

4.2.5 Trust

Based on the definitions of Jonker [41] and Castelfranchi [11], trust may be defined as: *the attitude of an agent towards an interaction partner about how likely it is to act honestly and cooperatively in situations where the achievement of a common goal depends on partner's actions; that results from reasoning about its capabilities, willingness, reliability, intentions, etc.* The trust attitude is characterised by beliefs and expectations and is relative to the competence of the agent, to the behaviour predictability of the agent and its honesty. For example, an agent might trust the commitment of another agent with respect to a certain goal, or it might trust that an agent is capable of performing a certain task.

4.2.5.1 Dimensions of Trust

Some researchers [12, 24, 47, 66] provide a particularly useful basis for classifying trust dimensions, namely, *competence trust* and *goodwill trust*. Competence trust (CT) is concerned with the expectation that the agent has the resources and capabilities for accomplishing a task competently and reliably, but this does not imply that the referred agent has the intention to do so. Goodwill trust (GT) is about the expectations developed that the agent has the willingness and commits itself to accomplish its task.

4.2.5.2 Beliefs and Expectations

According to Castelfranchi [11], the nature of the beliefs that compose trust are *positive expectations*, not neutral beliefs. In trusting an agent, a principal believes that such an agent has the right qualities, competence and willingness to complete a task. For example, a principal may express his trust in a highly skilled agent a to perform a printing task as: *I firmly believed that agent α is capable of accomplishing the task t* or *It can be expected that agent α is to accomplish the task t successfully.*

4.2.5.3 Degree of Trust

The *degree of trust* is the measure of such beliefs and expectations, and among other things, is the determinant for taking the risk of whether to cooperate, and with whom. In order to establish the degree of trust in a partner, an agent requires a scale of values where it can indicate the different levels of trust that can be put in an agent. These values on the scale are relative to a personal threshold. Similarly to Griffiths and Luck's work [34, 35], we define the degree of trust in an agent i, from agent j's point of view, as a numerical value from the interval between 0 and 1: $DT_{ij} \in [0, 1]$. These values are

only for reference and allow comparisons between distrust, values approaching 0, and full trust, values approaching 1.

The degree of trust is context-dependent and it can be related to a particular capability or willingness of an agent to accomplish a task. For example, to evaluate the trustworthiness of an agent in terms of its competence as a printing service provider, we should assign a trust value for the degree of trust for the speed and quality of printing. These values are refer to a personal threshold that signals when the abilities observed in an agent are *sufficient* to trust in it. For example, an agent may establish the value of 0.5 to indicate the threshold to be considered trusted. Hence, $DT_{ij}^{speed} = 0.8$ and $DT_{ij}^{quality} = 0.3$ indicates that in the opinion of agent i, agent j is highly trustworthy in terms of the speed of printing but the quality that can be expected is poor and it is not trustworthy in that context.

4.2.6 Risks in VOs

An organisation in a VO exposes itself to risks of failure because of many reasons. For example, due to demand uncertainty, or as a result of rapid shift in consumer preferences, the marketed product of a VO may no longer be required. Another reason for failure may be those situations when the buyer of a customised product or service decides not to purchase it, and blames the VO for not meeting his expectations, or when the complexity of the VO's goal and the time pressures impede the achievement of the goal. This kind of risk is *objective* because it is inherent in a given situation [40]. However, in this work we focus on *perceived risks*, particularly on those related to the uncertainty attending the cooperation among partners.

4.2.6.1 Perceived Risk

Perceived Risk (PerR) refers to those ambiguities, as perceived by prospective partners, about the future events that may negatively impact on the performance of the VO, which refers to the degree to which the partners' objectives are achieved in a VO. While the success of traditional organisations depends solely on the organisation's actions in comparison with the general and competitive environment, the success of VOs is contingent additionally on the effective cooperation between the partners. Hence, VOs are not only constrained by the objective business environment, but also by the conduct of each partner.

4.2.6.2 Relational and Performance Risk

In this model we adopt the two types of perceived risk defined by Das [24]: *Relational Risk* (RR) and *Performance Risk* (PR). The former is defined as the likelihood and consequences of not having satisfactory cooperation. This risk arises because of the potential misbehaviour on the part of the partners due to their individual interest not necessarily being congruent with that of their partners. Performance risk is the probability that the VO objectives are not achieved due to lack a of competence of the partners.

4.3 Gathering Experience

A general approach suggests that individuals interact with each other a minimal number of times in order to be able to predict others' future behaviour. At the end, once an individual has learned the behaviour of the rest, it can identify which are the most reliable, or it can reason about the acceptable strategy to deal with them.

In order to form and update reputation, individuals should establish the way that reputation is represented internally (when they are reasoning with reputation) and externally (when they are spreading reputation). We will discuss two approaches to update reputation based on subjective probability and reinforcement learning as a measures of reputation.

4.3.1 Subjective Probability

It could be argued that subjective probability is a useful metric for reputation because it provides a systematic way to update prior estimations or beliefs to posterior beliefs based on the state of information available to whoever evaluates reputation. To a certain limited extent this is true, but subjective probability has several deficiencies to be considered as a proper measure of reputation:

- It strongly depends on the initial probability evaluations and it ignores the question of how people actually evaluate the probability of an uncertain event. Sometimes, assumptions are imposed on individual judgements to permit the construction of probabilities [49]. For example in the event of flipping a coin, it is assumed that the coin is *fair* and only two possibilities are considered with the same odds 50-50. But, in the real world, this simple game could have different results that depend on the person who is flipping the coin for example. Moreover, to assure a desired level of confidence they require a minimum number of observations m. However, human behaviour is unpredictable and it may change in the $(m+1)$th observation. This would imply that either the prior set have to be disregarded and wait until a

new m number of observations, the new observation with the latest increase in the uncertainty must be included in the estimations of reputation. Also, the calculated probability is dominated by prior assumptions. In the above example, suppose two friends A and B flip a coin to decide who pays for the lunch. After B, who is the owner of the coin, has won several attempts, it is reasonable to think that A would start to suspect that B is cheating him. However, if A has the prior belief that B is honest then future observations will not change A's beliefs according to subjective probability.

- The latter shortcoming leads us to conclude that subjective probability does not offer a way to take advantage of the recent observations [74]. That is, there is not a function of feedback and experience that indicates whether a probability is more or less accurate after considering the subsequent events.

- Subjective probability follows the rules of probability, but reputation does not [42]. In particular consider transitivity. For example, if after an interaction, A believes that B has a reputation of 90% of being honest and recommends B to C, then C could conclude that B's reputation in its opinion is:

$$R_{CB} = R_{CA} \cdot R_{AB}. \tag{4.5}$$

  However, it may be that B changes its behaviour and decides to cheat C. In this case, we could think that while estimating probabilities, only the observations are considered rather than the entity doing the estimations.

- Moreover, subjective probability suits well when it is used to infer possible causes from the information observed (effects) in physical phenomena [22]. However, we think that it is hardly accurate to infer human behaviour because it can change without an apparent cause at any time and probability does not make any distinction between the observations as a function of the time they were collected.

4.3.2 Reinforcement Learning

Reinforcement learning is defined as the problem faced by an *agent* that must learn to find optimal *actions* by interacting with its environment. Agents learn what actions must be taken at a given time or under certain circumstances or *states* that maximizes its cumulative *reward*. As a consequence of each action an agent receives a numerical *reward*, but this *reward* depends on the *state* in which the *action* was taken.

For example, consider a robot which can perform two actions, turn on or turn off the lights in order to maintain a room illuminated from 8 am to 10 pm. These actions are taken twice a day, once in the morning and once in the evening. There are three states: a) room illuminated by daylight, b) room illuminated by artificial light or c) room in

darkness. There is one restriction: the room must not be illuminated by daylight and artificial light at the same time.

The robot is rewarded positively for turning the lights on before 10pm when there is no more daylight. However, the robot will get a negative reward if the lights are turned on during the morning, because of the restriction.

Now, each time an agent has to make a decision about what to do based on its current stage of learning, it will take into account which state-actions lead to the highest reward, that is, it will map state-action relations against rewards. This mapping is also referred to in the reinforcement literature [72] as *policy*. Furthermore, each action leads to a new state.

The reward can be immediate or delayed depending on the action taken. For example, an agent may take a series of actions where it receives no reward but at the end arrives at a state with a high reward, called delayed reward. In our example, consider that the robot decides to turn off the lights after 10 pm. Although this action does not have a reward, it leads to a new state: room in darkness, which will change during the morning of the next day. Thereby, in the morning, the room will only be illuminated by the daylight avoiding negatively rewarding the agent.

Reinforcement learning has been applied in situations more complicated than the situation described above such as learning how to control an elevator or how to allocate communication channels efficiently. In the elevator example, the elevator is the *agent*, the *states* of the elevator are defined by the position, direction and speed. Passengers who want to leave the elevator must press a set of buttons to indicate the floor where they leave the elevator. The elevator is considered as an independent entity able to decide which *action* to choose. For instance, if the elevator is going up, it can decide not to stop at the first floor where there is a passenger who wants to get on. The goal may be to minimize the *waiting time* of the passengers and for this purpose the agent has to learn a *policy*. The diagram in Figure 4.1 shows the agent-environment interaction in reinforcement learning.

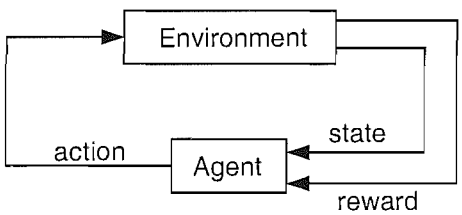

FIGURE 4.1: The interaction agent-environment in reinforcement learning.

4.3.2.1 Exploitation and Exploration

The absence of *instructions* indicating the correct actions in reinforcement learning problems leads agents to understand their environment by explicit *exploring* possible actions through trial-and-error interactions. Once agents are capable of identifying which actions provide the highest *rewards*, they can *exploit* them in order to maximize the total amount of rewards received in the long run, that is, the *expected return.*

Reinforcement learning problems have been faced at two levels of complexity. The simplest one, known as *evaluative feedback*, focuses on learning how to take actions in a single *state* by *evaluating* the immediate rewards received from each action, and at the same time it is concerned with the trade-off between exploitation and exploration.

The more complex is related with most real reinforcement learning problems and involves learning *policies* in environments with multiple *states* and, in some cases, delayed *rewards.* Algorithms to solve this problems are based on estimating *value functions* that can be either a *state value function* or an *action value function.* The former function estimates the *value* of a state s under a policy π as the expected reward received for following π after starting in s. On the other hand, the latter function estimates the value of the pair, action a and state s, under the policy π, as the expected return received for following π after starting in s, taking the action a.

4.3.2.2 Applications

Reinforcement learning has been applied in many areas, including: process control, scheduling, resource allocation, queuing, and adaptive games. For example, reinforcement learning has been traditionally used in repeated games against an unknown opponent such as the Iterated Prisoner's Dilemma [67]. Singh [71] has applied reinforcement learning to the problem of resource allocation in mobile telephone systems.

4.3.2.3 Evaluative Feedback

In this work we are interested in the *evaluative* aspect of reinforcement learning, this for simplicity we will initially study in depth the case of a single state in reinforcement learning, that is, the evaluative feedback approach.

Evaluation of the actions taken through the training information instead of getting *instructions* of the correct actions, is the characteristic that makes reinforcement learning particularly distinct from other types of learning. While an agent is interacting with the environment it receives rewards or feedback for its actions. The evaluation of feedback does not indicate if an action was the best or the worst but how good the action taken was, also known as the *value* of that action. Evaluative feedback is a simplified aspect of

reinforcement learning where the value of actions is estimated without considering how these actions affect either current or future *states*.

Stationary environment. When an agent is estimating the value of an action in a stationary environment, the real value of an action a is denoted as $Q^*(a)$ and can be calculated as the average of the reward received. That is, the value of an action at the time t that has been taken k_a times and produced rewards $r_1, r_2, ..., r_{k_a}$ is estimated as:

$$Q_t(a) = \frac{r_1 + r_2 + ... + r_{k_a}}{k_a}. \tag{4.6}$$

However, there are limits in the resources that make it difficult to maintain in memory the set of rewards when we want to estimate the value of many actions after multiple interactions. An iterative formula is derived in [72] to calculate the average of $k + 1$ rewards is:

$$Q_{k+1} = Q_k + \frac{1}{k+1}(r_{k+1} - Q_k). \tag{4.7}$$

The intuition behind the latter formula is the following:

- Each Q_{k+1} is an estimation or prediction of the real value of an action, denoted as Q^*.
- The estimations are based on updating of previous evaluations.
- For each experience a feedback r_{k+1} is received as evidence of the real value of an action. The difference between this feedback and the previous estimation Q_k is the error in the prediction of Q^*.
- An increment in the previous estimation updates the new estimation by approximating Q_{k+1} to Q^*. This increment is a fraction of the error in the prediction of Q^*.

The latter formula is an update rule that follows the general form:

$$NewEstimate \leftarrow OldEstimate + StepSize(Target - OldEstimate). \tag{4.8}$$

The expression $Target - OldEstimate$ can be seen as the *error* in the estimate. The target is assumed to indicate a desirable direction in which to move. The *StepSize* parameter is also known as *learning rate* that is denoted by $\alpha_k(a)$. In the iterative formula $\alpha_k(a) = \frac{1}{k_a}$, it is correct to assign functions that decrease along the time in a stationary environment to indicate that the learning process is reduced as the estimations get closer to the real value and the number of interaction increases.

Nonstationary Environment. Nonstationary environments, in which the received rewards for performing an action change continually, are more frequent in reinforcement learning problems. A natural approach is to give more importance to the recent rewards than the oldest ones. Now, contrary to stationary environments, where the learning rate is assigned firstly a value which decays progressively as the *real* value of an action is learned, in nonstationary environments Sutton [72] suggests to use a constant learning rate α, in the interval $0 < \alpha \leq 1$. Constant α is a good choice here because continuous learning is needed. But that constant has to be affected by the age of the reward in order to weight recent rewards more heavily than the past ones. So, according to Sutton Q_k can be calculated as a weighted average of the past rewards and the initial estimate Q_0:

$$Q_k = (1-\alpha)^k Q_0 + \sum_{i=1}^{k} \alpha(1-\alpha)^{k-i} r_i. \tag{4.9}$$

Here, it is important to note that the weight $\alpha(1-\alpha)^{k-i}$ given to the reward r_i depends on how long, *(k-i)*, such a reward was received. The weight given to each reward r_i decreases as time passes and the number of rewards increases. Sutton also suggests that when it is required that the step-size parameter varies from step to step, $\alpha_k(a)$ must satisfy the following conditions given by stochastic approximation theory to assure convergence to the true action values:

$$\sum_{k=1}^{\infty} \alpha_k(a) = \infty \tag{4.10}$$

and

$$\sum_{k=1}^{\infty} \alpha_k^2(a) < \infty. \tag{4.11}$$

Stochastic approximation theory explains the behaviour of a physical phenomenon or process for which it is not possible to write a *deterministic* model that allows exact calculation of the future behaviour of the phenomenon or process. Hence, stochastic approximation theory proposes models that can be used to calculate the *probability* of a future value lying between two specified limits. That is, it considers *approximations* that give the lower and upper bounds on the expected behaviour of the phenomenon or process.

Therefore, in order to make better approximations, condition (4.10) guarantees that the steps are large enough so that any initial conditions or random fluctuations will eventually be overcome. The second condition guarantees that eventually the steps become small enough to assure convergence of the action values.

Combining RL and Reputation. The initial idea is to use the RL method to estimate the reputation of a partner. We argue that there are certain similarities between reinforcement learning and reputation estimation. Both are concerned with the problem of prediction, that is, of using past experience with an uncertain system to predict its future behaviour. In both cases feedback is received after an interaction and it is used to estimate the values of actions in order to make action selection decisions. We consider the ratings given by the partners after each interaction as the rewards.

Although we only consider evaluative feedback in the reputation problem, we think that RL could be introduced to learn trustworthy thresholds to decide if it is convenient or not to interact with a particular individual. In this sense, states can be seen as mental states of trustworthiness and by performing an interaction (for example, buying a good) the agent is transferred to a new mental state (with an increase or decrease in the level of trust in that partner), which is reinforced by the rating given to that interaction.

4.3.3 Conclusion

It is clear that a measure different from probability is required for reputation. In this sense, we suggest to use a measure of reputation based on individual prior experiences that can be updated using an adaptive approach such as reinforcement learning.

4.4 Computational Model

Now that reinforcement learning has been chosen as a suitable mechanism for updating reputation, we define the computations required for evaluating the level of satisfaction of the actions taken by an agent and updating its reputation based on such evaluation.

4.4.1 Updating Direct Reputation

Each agent updates its reputation value of a service provider every time it receives impressions from either direct (immediate or observed interactions) or indirect experiences. Our first proposal to update the reputation values (after receiving t rated experiences or impressions) consists in the use of the following reinforcement learning based action update rules:

$$DR_t = DR_{t-1} + \alpha \cdot [imp_t - DR_{t-1}]. \tag{4.12}$$

Reputation, in Eq.(4.12) can be interpreted as the aggregation of the previous value of reputation plus a factor that strengthens or weaken that value. The factor mentioned

indicates how close is the recent impression to the past reputation. This factor is an indicator of how well the previous reputation predicts the latest given impression.

Note that although we omit the index *k, i and j* to make the expression more readable, DR_t makes reference to the reputation of an agent *target* in the opinion of an agent *source* for the context k. The update rule in Eq.(4.12) is a linear function. A linear function is required in an open environment where the number of interactions is reduced and reputation cannot be updated in the long term through a non-linear function because an agent is able to cheat in many occasions before the reputation is updated. Instead, reputation must be updated immediately after any interaction. Newcomers and agents with low values of reputation have an incentive to the improve their performance and get better values of reputation due to learning rate that also can be seen as a *forgetting factor*. That is, if α is near 0 then all the previous history will be forgotten, otherwise, if α is near 1 then the previous history will be preserved.

The learning rate α. The factor α is also known as a learning rate. We will discuss its meaning in Eq.(4.12) in order to fulfil the principles defined above. First, α is an indicator of how long past experiences will last in the memory of the system. For example, while low values of α mean that early experiences will have more influence in the system than recent ratings, high values of α indicate that early experiences will soon be forgotten. For our aims, we will consider α as a function $\alpha(DR_{t-1}, imp_t)$ with the following properties that are based on the ideas of Carbo et al [9]:

- The function $\alpha(DR_{t-1}, imp_t)$ will determine how fast the reputation value changes after an experience and how this affects the memory of the system. This depends on the accuracy of the predictions suggested by the *impressions* received; that is, how much similarity exists between the expectation formed by the previous reputation values and the last rating. We consider the initial value for the function $\alpha(DR_{t0}, imp_{t0}))$ to be 0.5. That is, as the agent starts to learn, it will be careful with the first impressions until it learns how to estimate better its predictions.

- Similarity will be estimated through a similarity function $\beta(DR_{t-1}, imp_t) \in (0, 1)$:

$$\beta(DR_{t-1}, imp_t) = 1 - e^{-10 \cdot ABS(E-imp)}, \tag{4.13}$$

  where E is the estimated rating based on the past reputation and rating:

$$E = \frac{DR_{t-1} + imp_{t-1}}{2}. \tag{4.14}$$

- Finally, the function $\alpha(DR_{t-1}, imp)$ is updated as follows:

$$\alpha(DR_t, imp) = \frac{\alpha(DR_{t-1}, imp) + \beta(DR_{t-1}, imp)}{2}. \tag{4.15}$$

Analysing the function α we can observe the following behaviour:

- By comparing the predicted reputation values and the latest rating in the α function we could get some conclusions. First, when they are quite different then α gets values near to 1.0, which means that the information previously stored is not useful and it will be updated. On the other hand, if ratings are similar to the reputation values then α gets values slightly greater than 0.0 and the reputation values are held fixed.
- The β function contains an exponential function that intensifies the differences between the estimated reputation and the ratings.
- Once the learning rate α is adjusted properly using the latter rules (a slowly decreasing is observed), the reputation updating formula converges to the real reputation value.

The factor of 10 in the $\beta()$ function requires a more detailed explanation: it is the acceleration factor of the similarity function $\beta()$, which slows down the changes for very similar estimated values of reputation and impression. The bigger the value of the acceleration factor, the steeper the similarity function $\beta()$. The value of the acceleration factor is chosen to be 10 so that the $\beta()$ function remains above 0.9 for dissimilarities whose values are above 0.25. The behaviour of the $\beta()$ function with different values of the acceleration factor is shown in Figure 4.2, which plots $\beta()$ for 6 values of the acceleration factor, ranging from 5 to 100.

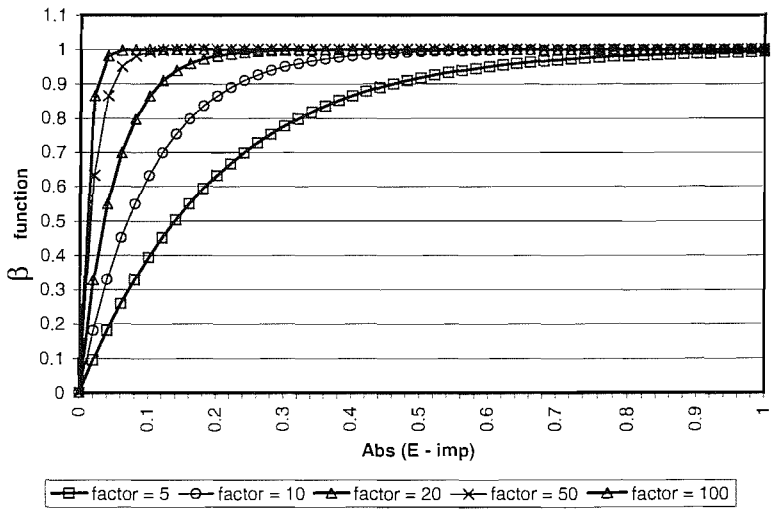

FIGURE 4.2: The behaviour of the beta function with 6 different values of the acceleration factor, ranging from 5 to 100.

4.4.2 Computation of the Degree of Trust

The degree of trust can be inferred through the evaluation of the experiences gained in direct interactions and the opinion of others who interacted with the subject of trust in the past. Direct evaluations are also known as *Direct Reputation, DR* and the trust a partner has in potential partners is a function of the service provider's reputation. The following relationship between reputation and trust is expected: An increase in the partner's reputation should also increase the trust from other agents for that partner. On the other side, others' opinions about the behaviour of a potential is important when there is a lack of interactions with the potential partner, and these opinions are also known as *Social Reputation, SR*. In Figure 4.3 it is illustrated how direct experiences and the opinion of *recommenders* contribute to compute the degree of trust that can be put in a partner. Hence, the degree of trust can be computed as follows:

$$DT = wr * ReDr * DR + (1 - wr) * ReSr * SR, \tag{4.16}$$

where wr is a parameter that reflects the agent's degree of reliance on its own experience represented by DR. Depending on the reliability of DR and SR, $ReDr$ and $ReSR$ respectively, the agent assigns different weights to experiences and recommendations by using wr. Because agents tend to rely on their own experience more than on others agents's recommendation, usually $wr > 0.5$.

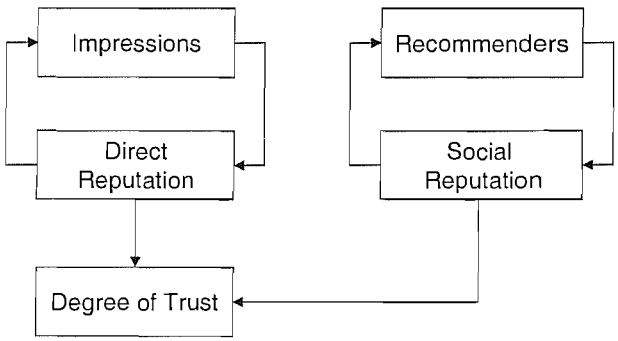

FIGURE 4.3: Model Trust-Reputation

4.4.3 Trust and Risk

Griffiths describes risk in terms of trust and establishes an inverse relationship between them. This relationship is based on the argument that an agent with a high trust value is more trusted than an agent with a low trust value, and represents less risk. Similarly, in this work we adopt such a relationship between the degree of trust, DT and risk, R, as follows.

$$R = \frac{1}{DT}. \tag{4.17}$$

Goodwill trust is about partners's good intentions and integrity. It is about whether an organisation has a reputation for dealing fairly and caring about its partner's welfare in alliances. With such a reputation, partners feel more assured that the potential partner will cooperate in *good faith*, rather than behave dishonestly. Then, it can be suggested that goodwill trust reduces a partner's perceived relational risk in VOs. Hence a relation between Relational Risk (RR) and Goodwill Trust (GT) can be established in our model.

On the other hand, competence is based on the various resources and capabilities of a organisation. These resources and capabilities provide the basis for competence or expertise that is needed in VOs. Moreover, partners that have been successful in previous alliances tend to build a reputation for competence. A partner's competence suggests that its tasks can be accomplished successfully, which implies a low performance risk. In this sense, because competence trust gives an organisation a sense of confidence that the partner is capable of accomplishing given tasks in the VO, Performance Risk is related to Competence Trust. Hence a relation between Performance Risk (PR) and Competence Trust (CT) can be established in our model.

Then the above relation is re-written as:

$$RR = \frac{1}{DT_{GT}} \tag{4.18}$$

and

$$PR = \frac{1}{DT_{CT}}. \tag{4.19}$$

Although mathematically, risk (RR,PR) may have values in the interval $[1, \infty]$ we limit its values to $[1, 100]$ to avoid the complexity of working with values of risk greater than 100 whose meaning is the same (high risk).

4.5 Conclusions

In this chapter we proposed a computational model that relates the concepts of trust and reputation in order to provide information for the formation of VOs, particularly on the selection of partners. First, we provided mathematical definitions of the basic concepts used in our model: Virtual Organisation, reputation, impression, trust and risk. Direct Reputation is updated using a reinforcement learning algorithm in order to provide an adaptive evaluation of reputation. Values of reputation are reinforced or weakened according to the similarity of the prediction of future behaviour in relation to the latest behaviour observed.

Our model also provides a differentiation between the risk due to the capabilities of a partner and the risk associated with the willingness of the partner to commit itself to the VO. These measures of risk provide an indication for VO members to decide if it is worth engaging or continuing in a partnership with an individual based on its reputation.

In deciding whether they should enter into a VO with other organisations, agents must identify the desired outcomes from the VO and the likelihood of achieving them. Figure 4.4 shows a decision-making process model, in which organisations consider suitable combinations of trust and control that satisfy a given threshold of risk, RT. Organisations assess the risk threshold that they can tolerate, their degree of trust in the potential partner, and ensure that the level of control they can exert over the VO will meet or exceed the need for control implied by the risk profile.

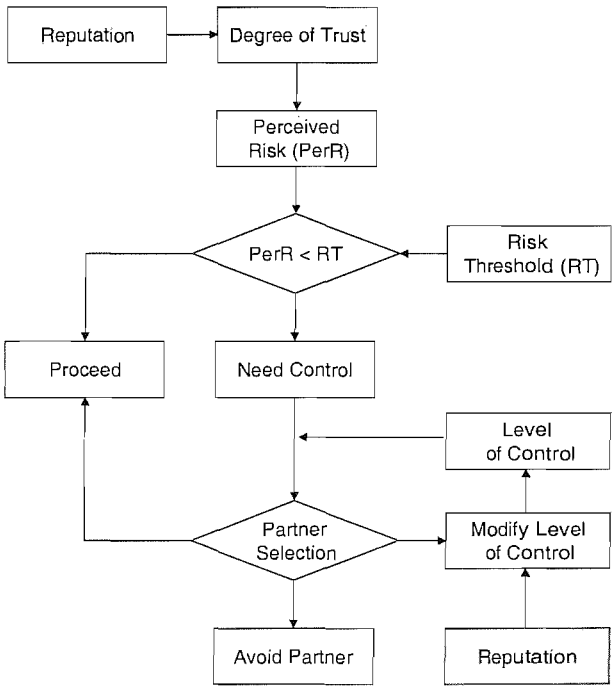

FIGURE 4.4: Model for evaluating potential partners

Reputation helps to estimate the degree of trust an agent deserves. The perceived risk, PerR, is computed as the sum of the performance risk, PR, and the relational risk, RR. Then, PerR is compared against RT. If PR is less than RT, then the potential partner is selected. Otherwise, adjustments in the control mechanism are required. It is possible to use a contract and adjust the issues under negotiation based on the reputation of the potential partner in order to modify the level of control. That is done by introducing those issues where the potential partner is weak into the negotiation process. These actions increase the level of control over the potential partner so that the VO can take the decision of selecting the partner, avoid or decline the participation of the partner, or modifying successively the level of control. Alternatively, changes in the degree of trust may occur as the partners explore different options for structuring the VO.

Finally, if several agents exceed RT, two actions may be followed:

- choosing the agent with the lower risk perceived value and increase monitoring activities during the operation phase; or
- dividing the task in question to reduce its complexity and try a new organisation structure.

Summarising, our claim is that the model developed in this chapter contributes to a better understanding of the way in which individuals should not only estimate the reputation of a partner in a dynamic environment but determine the degree of trust that it deserves and the risk perceived in the partnership.

Chapter 5

Experimental Work

5.1 Introduction

In this chapter we provide experimental evidence to demonstrate the validity of our solution and present our findings. We performed two sets of experiments to evaluate DIRECT (our algorithm for computing reputation based on direct interactions) and show its feasibility and effectiveness. The first set of experiments evaluates DIRECT in terms of its accuracy. The second set of experiments demonstrates the effectiveness of DIRECT for detecting changes in the behaviour of an interacting partner. This experimental evaluation is the first part of the experiments required and it is focused more precisely in two requirements: Dynamism, Adaptability and Predictability. Up to now, we are not covering the total of the requirements because our reputation model is limited, and the rest of the requirements are related with the aspects of control and social influence linked with the transmission of direct reputation among the partners of a VO.

We start by explaining the model used and follow with a detailed description of the experiments performed, including their hypotheses, input and output data, and demonstration criteria. Next we show and analyse the results obtained. Graphs for each of the experiments are presented to compare our results with two other models. Also, tests for significance are performed to show statistically the reliability of our findings. Finally, we discuss our results and make a few concluding remarks.

5.2 Experiments

We consider two agents interacting successively in a finite interval of time. One of them is the service provider (SP), and the other is the consumer agent (CA) of the service provided by SP. In each interaction CA collects information related to the capabilities

of SP. For simplicity, in this experiment, CA considers only the quality of the services (QoS) offered by SP. A third agent is the initiator of the VO (IA) who monitors the operation of the VO and takes the coordination decisions involved in the initiation of a partner in the VO, task allocation, and the termination of a partner from the VO. The initiator's model lies beyond the scope of our experimentation and is not included here. Now, based on this information, CA computes the reputation of SP, which is passed to the initiator of the VO for it to use in the process of forming a VO. The information about reputation of SP provides another basis for negotiation according our model, and thus is important to the process of VO formation.

For comparison purpose, we use two of the models of reputation previously analysed called: SPORAS and REGRET. These models were chosen because, as discussed in Section 3.4.3, reputation systems for VOs should consider the time when the interactions take place in order to update their reputation values, and both SPORAS and REGRET meet this requirement. In order to compare similar values of reputation, SPORAS and REGRET were modified to produce reputation values normalised in the range [0,1]. Additionally, in the case of REGRET, just the *individual dimension* of reputation is considered because it is the only one associated with direct interactions.

In order to validate our model we ran a series of experiments to evaluate the accuracy of the calculated reputation values and the ability to react to abusive behaviour.

5.2.1 Reputation Evaluation Accuracy

The objective of this experiment is to evaluate the accuracy of the reputation model with the basic parameters. The value of reputation must provide a measure of the true capabilities of a service provider for providing a service.

Hypothesis. *A reactive assessment of reputation provides a good perception of the true capabilities of a partner competence after a limited number of interactions.*

In order to verify our hypothesis, we will consider the quality of a service offered by a service provider. We will provide the same input data for the three models (DIRECT, REGRET and SPORAS) as follows:

- Input: We generate 10 series of data representing the quality perceived of a service (QP) during 60 interactions. This data varies randomly in the interval $-z$, $+z$ from a mean value q of the actual quality QoS. To illustrate how input data vary, in Figure 5.1 one of the 10 series of data is shown. As can be seen, we want to model the fact that although SP delivers its services with the same quality, that is $q = 0.5$, the consumer CA may perceive such a quality in distinct way. CA's perceptions may vary around the actual quality q with a standard deviation $\sigma = 0.03$, and according to a Normal Distribution.

- Output: We generate a graph indicating on the x axis the time t when the reputation is estimated, and on the y axis the reputation calculated.
- Demonstration: We determine which reputation curve shows values of reputation that predicts more precisely the quality QoS of the service provided using a minimal number of interactions.

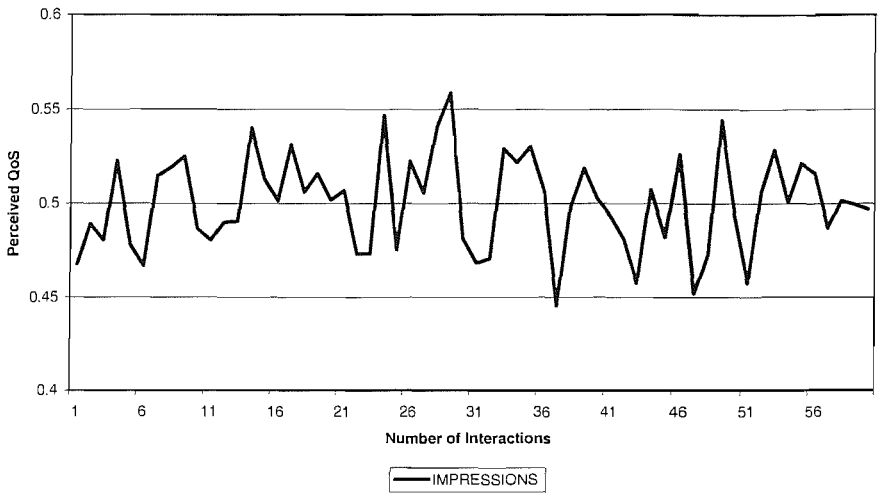

FIGURE 5.1: Input Data for Accuracy Experiment.

5.2.2 Abuse of prior performance.

This experiment is similar to the experiment described in [77] where a SP who joins a VO behaves reliable until it reaches a high reputation value and then starts committing fraud. Thus, in this experiment we aim to show quantitatively which model of reputation offers a mechanism for dealing with deceit. Deceit in a VO is found when a partner deteriorates the quality of its services once it has reached certain level of reputation to exploit it. We measure the rapidity with which agents *learn* oscillating behaviour of their partners in terms of the minimum number of interactions to adjust the reputation of a partner towards true quality of a service.

Hypothesis. *Deceptive behaviour due to changes in the quality of a service is detected more optimally by a balanced reactive mechanism. Optimisation in the detection of fraud and deceptive behaviour consists in finding out signals of such behaviour using a minimum number of interactions but discarding false information (noise).*

In order to verify our hypothesis, we consider a service provider agent who decreases the quality of its services after being considered a reputable supplier, in order to exploit the others by maintaining a minimal level of reputation. We also compare the results obtained with: SPORAS and REGRET.

We provide the same input data for the three models as follows:

- Input: We generate two sets of 10 data series representing the quality of a service during 120 interactions, both data vary using a normal distribution in the interval $-z$, $+z$ from a mean value *q1* and *q2* of quality. In the first set of data the agent provides highest quality during the first 25% of the interactions, 30, after that it decreases the quality in 6.25%, 12.5%, 18.75%, 25%, 31.25%, 37.5%, 43.75% and 50% for the rest of the interactions. To illustrate how input data vary, in Figures 5.2 to 5.5, one of the 10 series of data for each of the decrements in the quality (12.5%, 25%, 37.5% and 50%) is shown. As can be seen, we want to model the fact that after delivering its services with the same quality, $q1 = 0.8$ during the first 30 interactions, and this may be perceived by CA in distinct way, SP reduces the quality of its services to milk the reputation already built. As in the previous experiment, CA's perceptions may vary around the actual qualities *q1* and *q2* with a standard deviation $\sigma = 0.03$, and according to a Normal Ðistribution. Hence, for example in Figure 5.3, the values of QoS perceived by CA vary first in the interval $[0.77, 0.83]$ and afterwards in the interval of $[0.57, 0.63]$. Similarly, in Figures 5.4 and 5.5 the values of QoS vary in the interval $[0.77, 0.83]$ during the first quarter of the interactions but in the following 90 interactions the values of QoS diminish in the interval of $[0.47, 0.53]$ and $[0.37, 0.43]$ respectively.

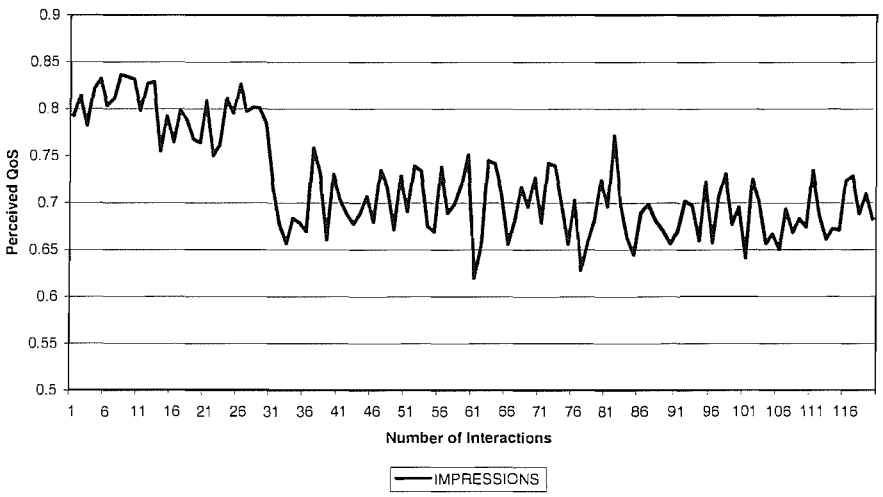

FIGURE 5.2: QoS is decreased in 12.5%.

- Output: We will generate two set of graph. The first set indicates on the x axis the interactions when the reputation was estimated, and in the y axis the reputation calculated. The second set is the conclusion of the first set and indicates in the x axis the minimal number of interactions to detect the change of quality *q2*, and in the y axis the percentage of variation in the initial reputation *q1*.

- Demonstration: We compare which curve contains the values of reputation closer to the mean *QoS*, that is, which model is able to detect deception more optimally.

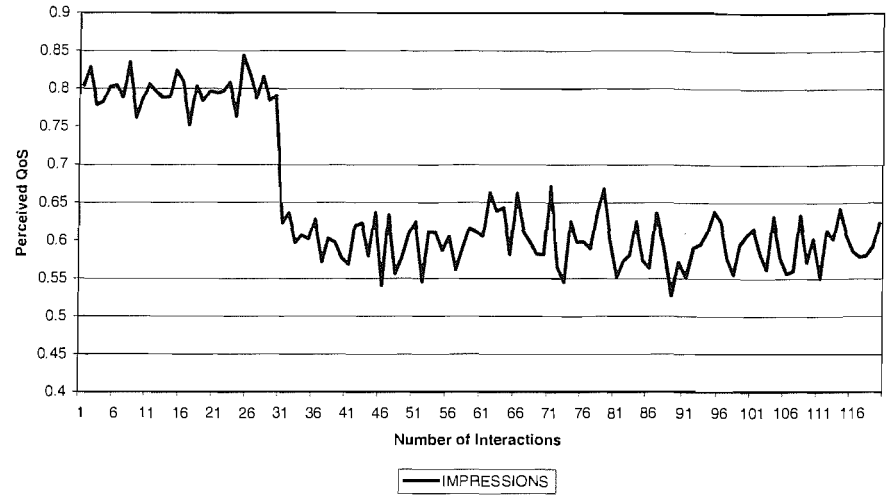

FIGURE 5.3: QoS is decreased in 25.0%.

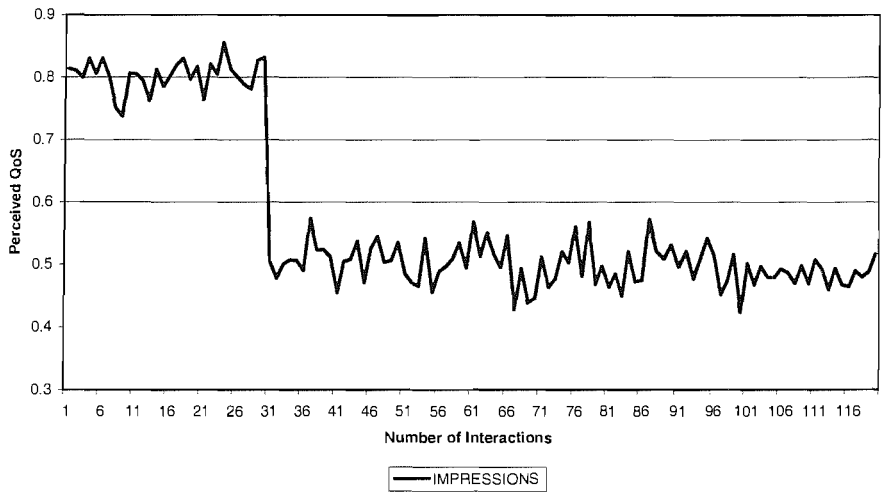

FIGURE 5.4: QoS is decreased in 37.5%.

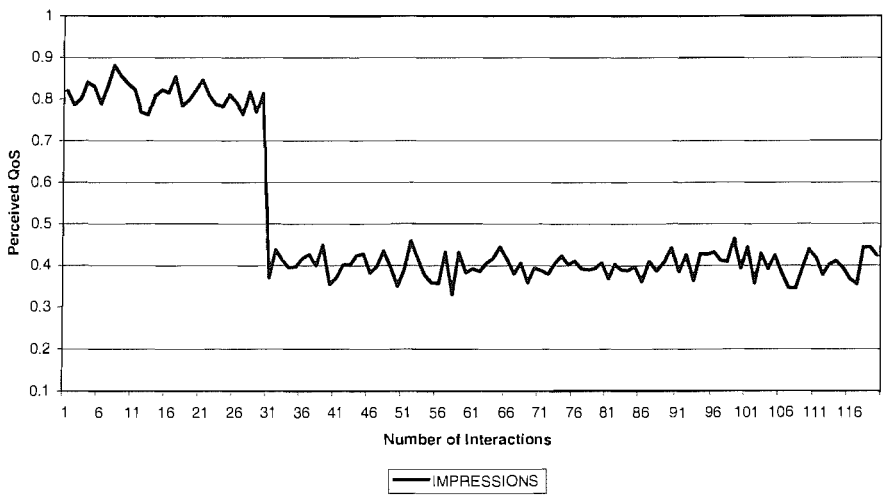

FIGURE 5.5: QoS is decreased in 50.0%.

5.3 Results

Results are divided in two sections to demonstrate the suitability of our model for fitting the requirements already described. In both sections we include Tests for Statistical Significance and the type of statistic used is the T-Test. The aim of T-Tests in this work is to know how reliable are the differences between the models of reputation compared. More in detail, we want to estimate the probability of that what we think is a difference between the models is not just due to random chance but it can be found in any other simulation of the same kind.

5.3.1 Results of Reputation Evaluation Accuracy

In Figures 5.6 and 5.7 the reputation values computed with the three algorithms are shown. Figure 5.7 is a detailed section of Figure 5.6. As can be seen, our proposal obtains similar results as REGRET in a similar number of interactions. DIRECT and REGRET establish the reputation value faster than SPORAS. Although REGRET and our proposal DIRECT use a different aggregation algorithms for computing reputation, both obtain accurate results when a SP maintains the provision of its service without change.

	DIRECT	REGRET
Number of Observations	10	10
Mean	2.5	1.5
Standard Deviation	1.2692955	0.9718253
Variance	1.6111111	0.9444444

$t = 1.8766297, df = 18, p > .05$

TABLE 5.1: DIRECT and REGRET interactions in accuracy experiment.

In our simulations, we compute the number of interactions before the reputation curves generated by each of the algorithms converge towards the actual QoS. The convergence is considered when the calculated values of reputation are in the interval $[q - \sigma, q + \sigma]$.

Tests for statistical significance were calculated from the number of interactions required by each of the models to reach an acceptable value of direct reputation. In this experiment, the range is $[0.47, 0.53]$. In Figure 5.8, the mean and the standard deviation of the number of interactions to converge are depicted after 10 simulations.

In order to test whether differences between the number of interactions to converge using DIRECT and REGRET and between DIRECT and SPORAS, a series of T-Tests were carried out and their results are shown in Table 5.1 and Table 5.2 respectively.

Table 5.1 shows no significant difference ($p > 0.05$) between the number of interactions required by DIRECT and REGRET to achieve an acceptable value of reputation that

	DIRECT	SPORAS
Number of Observations	10	10
Mean	2.5	12.6
Standard Deviation	1.2692955	1.264911064
Variance	1.6111111	1.60

$t = 16.908884, df = 18, p < .05$

TABLE 5.2: DIRECT and SPORAS interactions in accuracy experiment.

indicates the quality of the service provided. On the other hand, Table 5.2 shows a significant difference ($p < 0.05$) in the number of interactions, indicating that SPORAS scores were generally higher than DIRECT.

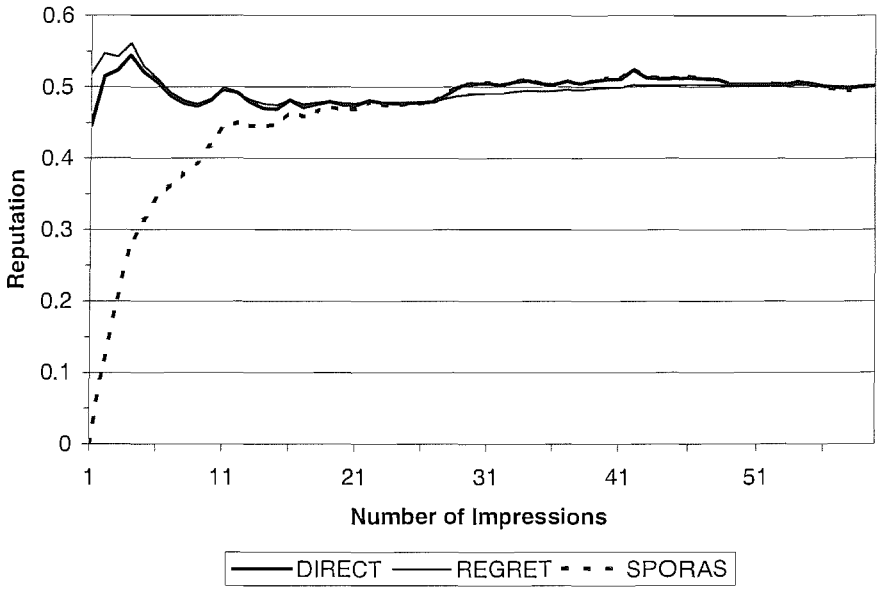

FIGURE 5.6: Building Reputation.

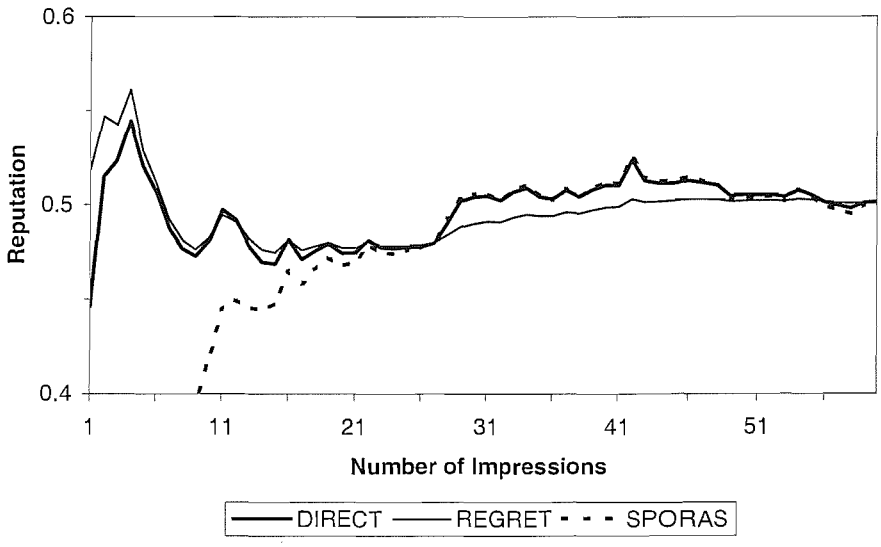

FIGURE 5.7: Detail of building reputation process.

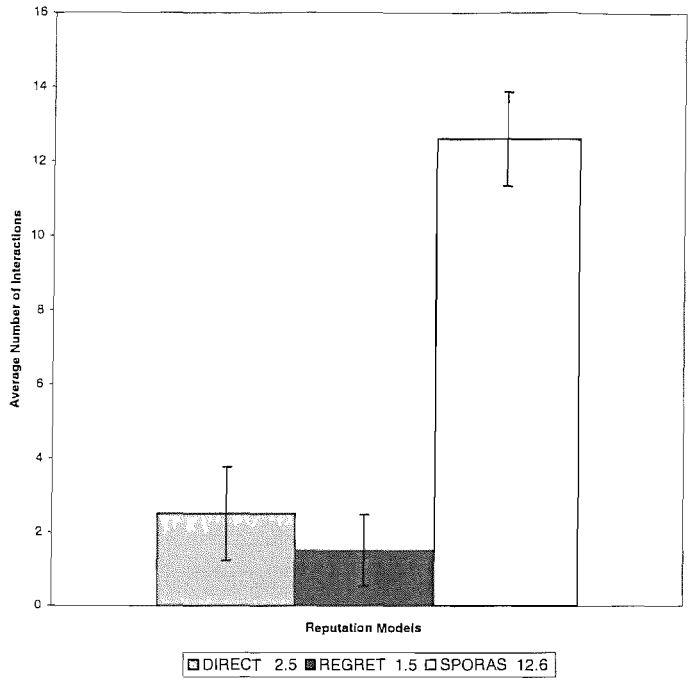

FIGURE 5.8: Number of interactions before reputation converges towards QoS.

5.3.2 Results of Abuse of prior performance

In Figures 5.9 to 5.12 we can see that REGRET requires in general more interactions to adapt its reputation values to the change of behaviour of service providers. In contrast, DIRECT and SPORAS show a more adaptive behaviour and require fewer interactions. On the other hand, REGRET updates reputation values very slowly, hence opportunistic service providers might take advantage of this by getting high values of reputation up to be considered as well reputed and then start to cheat.

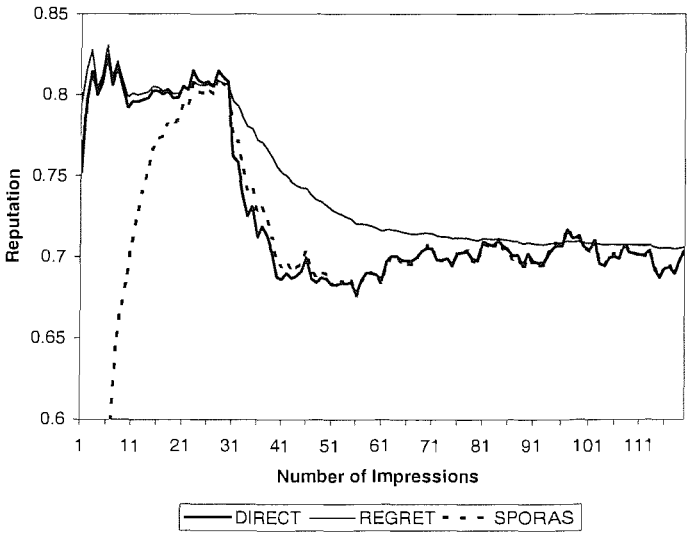

FIGURE 5.9: Reputation is decreased in 12.5%.

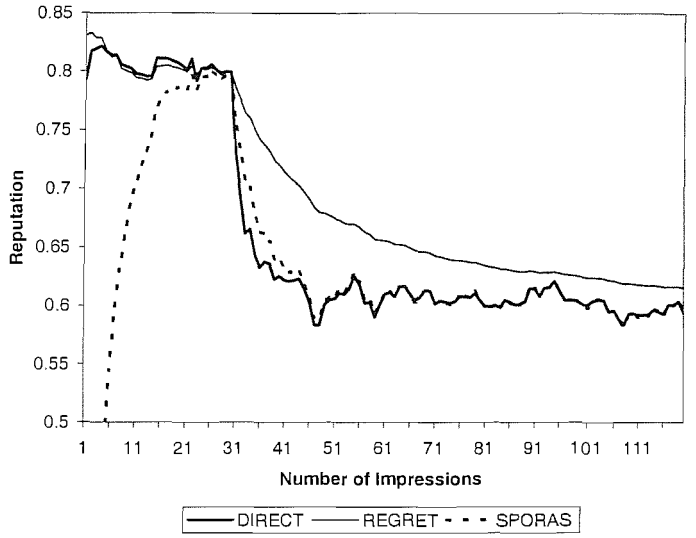

FIGURE 5.10: Reputation is decreased in 25.0%.

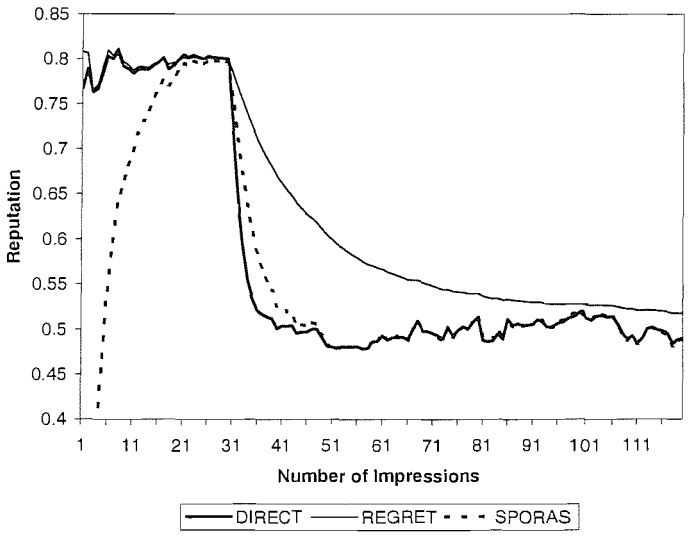

FIGURE 5.11: Reputation is decreased in 37.5%.

In Figure 5.14 the minimum number of interactions is shown to adjust the decrement in the reputation value of 6.25%, 12.5%, 18.75%, 25%, 31.25%, 37.5%, 43.75% and 50%. As can be seen SPORAS and DIRECT require less than 20 interactions to adjust their reputation values to the change in quality of the service, regardless of the percentage in which the quality is reduced. In REGRET, due to the accumulation of experiences, the effect of past experiences on the computation of reputation provokes accelerated increment in number of interactions required to adjust its reputation value.

The key difference between SPORAS and DIRECT is the ability to distinguish changes in behaviour based on the accuracy of the predictions. That is, DIRECT adjusts its values of reputation faster than SPORAS when the expectations created by the reputation values are closer to latest impressions. This can be seen in the slope of both curves.

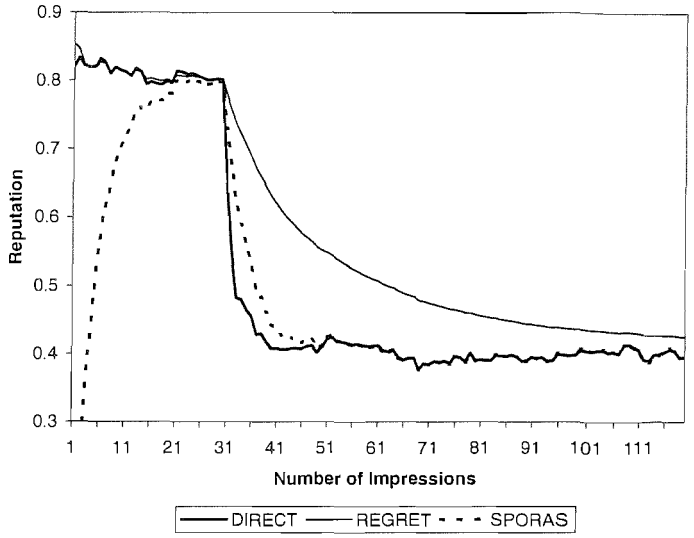

FIGURE 5.12: Reputation is decreased in 50.0%.

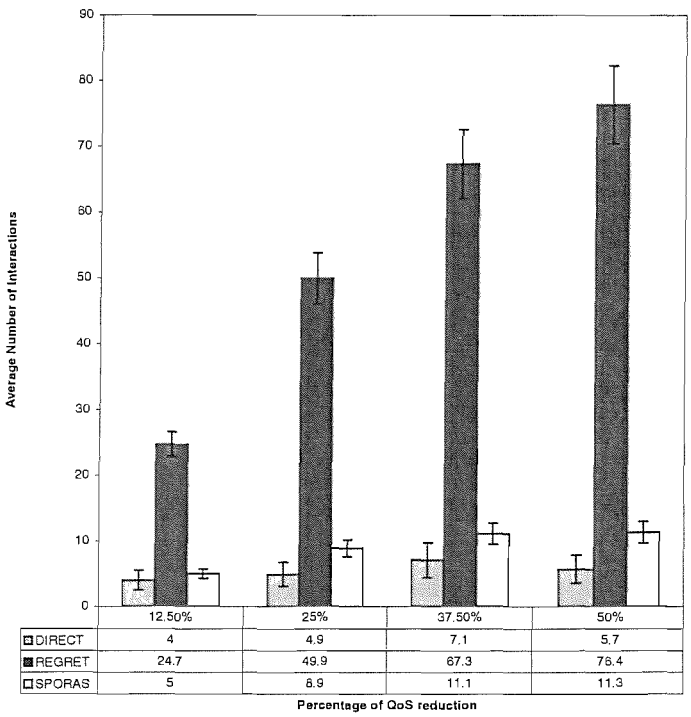

FIGURE 5.13: Detail of building reputation process.

While the rapidity to detect the change in the quality of a service is alike for both curves, DIRECT presents values of reputation closer to the true quality of service.

In this experiment, like the former one, we compute the number of interactions before the reputation curves generated by each of the algorithms converge towards the actual QoS, but contrary to the previous experiment, in this case reputation values should converge to the decreased value of reputation, $q2$. The convergence is considered when the calculated values of reputation are in the interval $[q2 - \sigma, q2 + \sigma]$.

	DIRECT (N=10)	REGRET (N=10)	value of t
Mean with 12.5%	4.0	24.7	14.723707*
Mean with 25.0%	4.9	49.9	25.938100*
Mean with 37.5%	7.1	67.3	41.925742*
Mean with 50.0%	5.7	76.4	27.626077*

*t significant at $p < .05$

TABLE 5.3: Interactions required to detect abuse of prior performance in DIRECT and REGRET.

Tests for statistical significance were calculated from the number of interactions required by each of the models to detect abuses of prior performance in situations where the provider decreases the quality of its service in 12.5%, 25.0%, 37.5% and 50.0%. In Figure 5.13, the mean and the standard deviation of the number of interactions to converge are depicted after 10 simulations for each of four simulations.

In order to test whether differences between the number of interactions to converge using DIRECT and REGRET and between DIRECT and SPORAS, a series of T-Tests were carried out and their results are shown in Table 5.3 and Table 5.4 respectively.

	DIRECT (N=10)	SPORAS (N=10)	value of t
Mean with 12.5%	4.0	5.0	1.500000*
Mean with 25.0%	4.9	8.9	4.819315**
Mean with 37.5%	7.1	11.1	4.372075**
Mean with 50.0%	5.7	11.3	6.606463**

*t not significant at $p > .05$ **t significant at $p < .05$

TABLE 5.4: Interactions required to detect abuse of prior performance in DIRECT and SPORAS.

Table 5.3 shows a significant difference ($p < 0.05$) of the number of interactions, indicating that REGRET scores were generally higher than DIRECT. Additionally, an increment sustained in REGRET scores indicates that REGRET maintains the inertia of the first 30 interactions longer as the difference between *q1* and *q2* increases. On the other hand, Table 5.4 shows no significant difference ($p > 0.05$) between the number of interactions required by DIRECT and SPORAS to detect abuse of prior performance when the difference the difference between *q1* and *q2* is not greater than 12.5%. But for differences greater or equal to 25% Table 5.4 shows significant difference ($p < 0.05$) of the number of interactions, indicating that SPORAS scores were slightly higher than DIRECT.

5.3.3 Discussion

Without DIRECT, and from the point of view of the performance of the two models (REGRET and SPORAS) to estimate values of reputation nearest to the actual quality

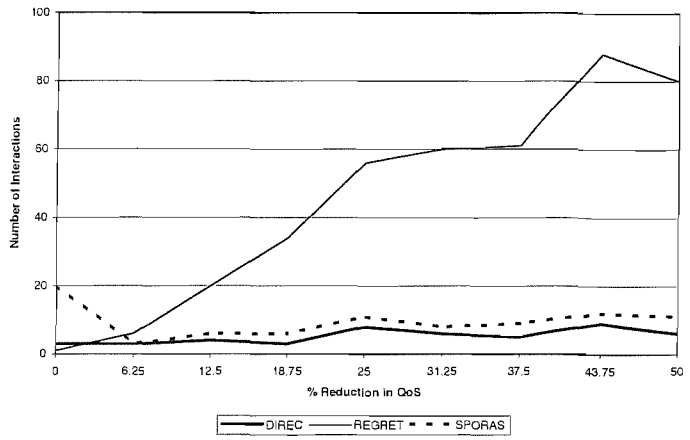

FIGURE 5.14: Minimum number of interaction to adjust reputation values.

of the service provided, there are no ways to conclude that one model of reputation is better than other. This is because, as was shown in the experiments, this depend on the degree of variability in the behaviour of reputed individuals and, of course, on the time passed since the first interaction. If direct reputation is calculated for the first time, an excellent option is REGRET. However, REGRET shows a very poor performance in situations of abusive behaviour. On the other hand, SPORAS is able to rapidly detect changes in the provision of a service, obviously because it is based on a reactive algorithm.

On the contrary, DIRECT offers an acceptable way to address both issues.

- First, in the accuracy experiment, in spite of the noise in the impressions, the values of reputation converge to the quality of the service provided in a reduced number of interactions. Hence, if the service provider constantly shows a good behaviour since the first interaction then the service consumer will not avoid interacting with it because the service provider reputation accurately reflects its behaviour.

- Second, as can be seen, on average DIRECT employs less than 8 interactions to detect abusive behaviour. These results coincide with the work of Banerjee [2]. Banerjee theoretically derived an optimal learning rate that maximizes reactivity while minimising noise-sensitivity. In his work, Banerjee explores a non-stationary environments where a learning agent is capable to react to prevent exploitation by dishonest agents. Moreover, Banerjee emphasis the need to balance reactive behaviour with noise-sensitivity. In a VO environment, high reactivity may have some disadvantages. For example, *impressions* might be imprecise and introduce noise in the reputation values crating two negative effects: volatility and instability in the reputation values. Both effects do not allow any conclusion about the behaviour of a partner so reputation values might be useless. Baneerjee derived a expression for computing the optimal *window* that maximises reactivity while maximising the size of the past *rewards*.

It should be emphasised that the power of direct reputation in VOs is its ability to make accurate predictions about the capabilities of an interaction partner using few interactions. In order to make better estimations from service consumers' experience, direct reputation models must adapt themselves to detect and prevent exploitation by dishonest partners, and at the same time reflect the degree of satisfaction of service consumers. In this sense, if we analyse the evidence from the curves plotting direct reputation values when a service provider intends to abuse its good reputation by reducing the quality of its services, we can conclude the following:

- REGRET does not provide a good estimation because it suggests that the quality of the service provided is lessened gradually but this is incorrect: quality is actually reduced suddenly.
- Although SPORAS detects when the quality of the service changes, it demands that newcomers interact many times before making good estimations.
- Finally, DIRECT not only requires few interactions to make good predictions, but also rapidly recognises when a change in the quality has occurred. Thus, DIRECT describes precisely when the quality of the service changes, but this is not done abruptly. DIRECT employs on average 5.5 interactions to adjust its values of reputation and avoids imprecise *impressions* introducing noise in the reputation values.

This explanation serves to justify the improvement of our model over SPORAS and REGRET to compute direct reputation in VOs.

5.4 Conclusions

In this chapter the accuracy of the DIRECT's values of reputation and DIRECT's ability to detect abusive behaviour using a reduced number of interactions have been verified experimentally. Accuracy of values of reputation refers to the similarity between the behaviour observed of an interaction partner and the behaviour expected based in its reputation. Detection of abusive behaviour refers to the rapidity to detect changes in the behaviour of a service provider that affects negatively the provision of a service, but they are difficult to detect because in the past interactions the behaviour of those providers was appropriated.

Additionally, we have shown some differences between the three models analysed in terms of the number of interactions required to evaluate or detect changes in the behaviour of a partner. Abuse was evaluated in different levels, i.e. the quality of the service provided was reduced in different percentages to evaluate the number of interactions

required by each model to adjust their values of reputation. In this sense, our model obtains acceptable results compared with REGRET and SPORAS.

The experiments defined here can be used by VO's partners to decide when it is worthy to initiate a partnership and with whom, and when a partnership must be terminated. Designers of VOs can also use this information to build VOs where its members are more autonomous and also capable of distinguishing deceitful partners and avoiding them.

Chapter 6

Conclusion and Future Work

6.1 Introduction

Several, researchers have presented models for the automated formation and maintenance of VOs based exclusively on *task-related* and utility criteria (i.e. the list of tasks and their scheduling, and the *price* asked by potential partners), but they have neglected the nature of the participants who are assumed to be fully competent, cooperative and trustworthy. In this thesis we have discussed the problems of such assumptions and proposed a partner selection model based on an evaluation of the risk perceived in a prospective partners in terms of the *degree of trust* inspired. The *degree of trust* determines the decision of whether or not to form or continue an alliance with a current or future partner, and it is a function of the reputation that such a partner establishes, based on its behaviour.

On the other hand, a key issue in the construction of reputation models for VOs is to cope with the high dynamism and flexibility of a VO's structure. The underlying assumption of current reputation models is that the communities where individuals interact are relatively static, and individuals can build their reputations steadily based on relatively predictable interactions. This assumption holds only in relatively closed environments. The work developed in this thesis addresses this issue, providing an adaptive reputation mechanism that is able to compute acceptable values of reputation based on a relatively small set of interactions. In this chapter, the main contributions of the thesis are summarised, and the limitations of the work, and topics for future research, are discussed.

6.2 Current Work

We have provided a critical overview of the state of the art in the field of VOs and reputation. We argue that subjective aspects of partners such as their *competences* and *trustworthiness* should be taken into account in partner selection decisions, since these aspects ultimately influence the cooperation between partners. Moreover, we assert that reputation plays an important role in VOs when members decide who to interact with and when to interact, by providing information about the past behaviour of potential partners, their abilities and reliability. In particular, we have asserted the importance of reputation not only in the formation process of VO but in the operation process.

Additionally, we discussed the requirements for building reputation systems that pursue three basic objectives in the formation and operation of VOs: (1) they provide useful information about potential partners for selecting the most *appropriate*, and eventually enable the formation of VOs; (2) they foster trust among the partners of the VO by revealing each partner's capabilities and predicting its future behaviour; and (3) they offer a means for enhancing cooperation by detecting and deterring deceptive behaviour through imposing *collective sanctions* on defectors.

Identifying these three basic objectives is particularly important, because the various issues concerning reputation model design have either been dealt with separately or ignored in the literature, each forming a different piece of work, without any consideration of how they all fit together. To rectify this position, we presented a list of requirements grouped in five categories to cope with each of the basic objectives mentioned. Although the third objective is out of the scope of this thesis, we proposed a reputation model based on the evaluation of direct experiences to achieve the first two objectives.

Finally, we have provided experimental evidence to demonstrate the validity of the model developed and the fulfilment of the requirements mentioned above, including a comparative analysis of the model proposed in this thesis and two other models. In particular, two aspects were analysed regarding the accuracy of the values calculated and the ability to detect abuses. As a means for supporting our findings, tests for statistical significance were calculated.

6.3 Contributions

In what concerns the introduction of reputation in VOs, and the development of a direct reputation model to enable the members of a VO to take decisions about who cooperate with and when, the work of this thesis provides three main contributions in the field of Virtual Organisations.

- The thesis provides assistance to designers of VOs to identify situations where reputation is required to cope with the uncertainties in the behaviour of current and future partners, and includes a list of requirements classified according to the needs of VOs.

- It provides a direct reputation model based on an adaptive algorithm for computing the reputation of a partner by enforcing or weakening an opinion about a partner's capabilities, trustworthiness or honesty through the evaluation of direct interactions in a decisive time window. This adaptive approach is matter of greatest importance to VOs because the environment where they are formed is highly dynamic and uncertain and we need to identify any possible misbehaviour.

- It provides a comparative evaluation of the proposed model, in relation to two other models, undertaken in order to demonstrate that our model better meets the requirements. In comparison with other models of reputation, we showed quantitatively that our model provides acceptable estimations of reputation in short term relationships such as those in VOs.

These contributions facilitate the development of VOs with characteristics that have not generally been considered previously. In particular,

- since individuals decide on the basis of the profits they can make by joining a VO, they can estimate, based on the reputation of its partners, if it is worth forming a partnership.

- since individuals can abandon a VO without facing future consequences, direct reputation may be spread in the community to warn others about its past behaviour.

6.4 Limitations

The core of our work consists in the functionality of the provided direct reputation model. However, not all the situations or cases could be covered due to time limitations and the complexity of some aspects. In particular, our model cannot be applied in the following cases.

- Erratic Behaviour. We have explained that the reputation of an individual is the result of finding similar behaviour in recent time contained within an optimal time window by comparing the predictions of behaviour and the current behaviour. Consequently, when an individual behaves erratically it is impossible to make predictions and obtain a reliable value of reputation. Our model does not allow for this situation.

- Noise. In our work, individuals compute the reputation of an interaction partner based on the impressions they obtain from each interaction, and these impressions may be obtained by comparing expected behaviour defined in a contract or induced by experience. It is in the latter case where our model is weak, because it is possible that in the absence of contracts or standards, the impressions induced by experience may be biased by noise and lead to incorrect values of reputation.

- Learning by trial-and-error. Although our model can be extended to consider the opinions of others and learn from their experiences to compute the reputation of a potential partner, our model portrays one individual interacting with another, but the bigger picture (who interacts with whom, who learns from whom, etc.) has not been considered.

6.5 Future Work

Although this thesis has answered how reputation is relevant to recognise cooperative partners through direct interactions, it opens up more research opportunities and questions that are unanswered. Moreover, there are other issues that were not faced in this thesis, due to the bounds imposed on the research, and still need to be addressed.

- *Social Reputation.* An important topic of research concerns the construction of a Social Reputation model. Thanks to Social Reputation, the number of partners one has reputation information about exceeds the number of agents with whom one interacts. Therefore, even with low probability of repeated interaction, the number of reputed partners one finally interacts with is higher, and the probability of being cheated is lower.

- *Social Networks.* Social behaviours such as *cooperation*, *trust*, *team working* or *collective sanctioning* that characterise VOs cannot be fully understood without considering the *social network* structures over which a VO's operation occurs. Hence, a research topic is exploring social network approaches for VO formation that considers new partner search and selection mechanisms, and cooperative behaviour incentives.

- *Social Mechanisms of Control.* In this work, it has been assumed that cooperation may be fostered and the behaviour of a partner in a VO may be controlled directly by using a formal contract or indirectly by evicting it from the VO and avoiding interacting with it in the future. However, this approach has many flaws because deceptive partners may join diverse VOs without suffering the consequences of their past. An alternative to foster cooperation is the use of *social mechanisms* of control. Because VOs are embedded in societies or communities of individuals

or organisations, their members may spread reputation information in the society about the misbehaviour of third parties.

- *Negotiation.* The temptation of easy profits, cheating or free-riding in VOs can be ameliorated by endowing partners with efficient negotiation capabilities that constrain the issues with high risk of failure or abuse. Our reputation model can be used together with a negotiation model to improve the success of the entities in their negotiations. The purpose of the negotiation in the VO is to reach an agreement about the provision of a service. The necessity of using reputation in the negotiation process arises when the provision of the service is not as expected. If the provision is worse than expected, probably the VO will not reach its main objective. On the other hand, if the provision is better than expected, then it means that the entity may have lost time in negotiation. For instance, the entity could have accepted a poor offer knowing that in the end it would be enough. Also, arriving at an agreement sooner maximises the opportunities of the VO.

6.6 Concluding Remarks

Our work contributes to research in the design of Virtual Organisations in providing answers for some of the questions that have concerned researchers on Virtual Organisations [6, 8]. Some have argued that contracts are enough for assuring cooperation among the members of a Virtual Organisation. In our work, we have discussed the implications of such assumptions and have shown how partners in a Virtual Organisation can determine the risk of interacting with other members based on their reputation.

We have given the means to consider the possible situations where the reputation of the individuals affects the formation and operation of VOs. We have also provided a model to estimate the reputation of an individual interacting inside a highly dynamic and flexible structure such as Virtual Organisations. Our thesis contributes to a better understanding of reputation and the role it plays in Virtual Organisations. These contributions are a step to towards the computational representation of Virtual Organisations, whose members do not necessarily share the same objectives and interests, that might not trust each other, but that should work together and help each other to achieve a common goal. Such forms of organisations are likely to pervade not only business domains but other environments where collaboration is essential such as e-science projects. Thanks to the contributions of this thesis we are nearer to achieve that goal.

Bibliography

[1] A. Abdul-Rahman and S. Hailes. Supporting trust in virtual communities. In *Proceedings of the 33th Hawaii International Conference on System Sciences*, pages 188–188, Hawaii, USA, 2000.

[2] B. Banerjee and J. Peng. Countering Deception in Multiagent Reinforcement Learning. In *Proceedings of the IEEE International Conference on E-Commerce (CEC03)*, pages 162–174, 200.

[3] B. Banerjee and S. Sen. Selecting partners. In *Proceedings of the Agents 2000 Conference*, pages 261–262, Barcelona, Spain, 2000.

[4] S. Braynov and T. Sandholm. Trust revelation in multiagent interaction. In *Proceedings of CHI'02 Workshop on Philosophy and design of Socially Adept Technologies*, pages 57–60, Minneapolis, USA, 2002.

[5] J.A. Byrne, R. Brant, and O. Port. The virtual corporation. *Business Week*, 8:36–40, February 1993.

[6] L.M. Camarinha-Matos and H. Afsarmanesh. Virtual enterprise modeling and support infrastructures: Applying multi-agent system approaches. In Michael Luck, Vladimír Marík, Olga Stepánková, and Robert Trappl, editors, *Multi-Agent Systems and Applications, 9th ECCAI Advanced Course ACAI 2001 and Agent Link's 3rd European Agent Systems Summer School, EASSS 2001, Prague, Czech Republic, July 2-13, 2001, Selected Tutorial Papers*, volume 2086 of *Lecture Notes in Computer Science*, pages 335–364. Springer, 2001.

[7] L.M. Camarinha-Matos and H. Afsarmanesh. Elements of a base VE infrastructure. *Computers in Industry*, 51:139–163, 2003.

[8] Luis M. Camarinha-Matos. Infrastructures for virtual organizations - where we are. In *Proceedings of ETFA'03- 9th Int. Conf. On Emerging Technologies and Factory Automation*, Lisboa, Portugal, 2003.

[9] J. Carbo, J. Molina, and J. Davila. Trust management through fuzzy reputation. *International Journal of Cooperative Information Systems*, 12(1):135–155, 2003.

[10] J. Carter, E. Bitting, and A. Ghorbani. Reputation formalization for an information-sharing multi-agent system. *Computational Intelligence*, 18(4):515–534, 2002.

[11] C. Castelfranchi and R. Falcone. Principles of trust for mas: Congnitive anatomy, social importance, and quantification. In Y. Demazeau, editor, *Proceedings of ICMAS'98*, pages 72–79, 1998.

[12] C. Castelfranchi and R. Falcone. Trust and Control: A Dialectic Link. *Applied Artificial Intelligence*, 14:799–823, 2000.

[13] C. Castelfranchi and Y.H. Tan. The Role of Trust and Deception in Virtual Societies. In *Proceedings for the 34th Hawaii International Conference on System Sciences*, 2001.

[14] Cristiano Castelfranchi. Engineering social order. In *Engineering Societies in the Agent World, First International Workshop, ESAW 2000*, volume 1972 of *Lecture Notes in Computer Science*, pages 1–18. Springer, 2000.

[15] C. Ching, C. W. Holsapple, and A. B. Whinston. Toward it support for coordination in network organizations. *Information and Management*, 30:179–199, 1996.

[16] J. Collins and M. Gini. A testbed for multi-agent contracting for supply-chain formation. In *Proceedings of the Workshop On Agent-Based Approaches to B2B at the 5th International Conference on Autonomous Agents*, 2001.

[17] J. Collins, W. Ketter, M. Gini, and B. Mobasher. A Multi-Agent Negotiation Testbed for Contracting Tasks with Temporal and Precedence Constraints. *International Journal of Electronic Commerce*, 7(1):35–57, 2002.

[18] J. Collins, M. Tsvetovat, B. Mobasher, and M. Gini. Magnet: A multi-agent contracting system for plan execution. In *Proceedings of Workshop on Artificial Intelligence and Manufacturing: State of the Art and State of Practice*, pages 63–68, 1998.

[19] R. Conte and M. Paolucci. *Reputation in Artificial Societies. Social Beliefs for Social Order.* Kluwer Academic Publishers, 2003.

[20] R. Conte and M. Paolucci. Social Cognitive Factors of Unfair Ratings in Reputation Reporting Systems. In *Proceedings of the IEEE/WIC International Conference on Web Intelligence (WI'03)*, pages 57–60, 2003.

[21] Rosaria Conte. Emergent (info)institutions. *Journal of Cognitive Systems Research*, 2:97–110, 2001.

[22] G. D'Agostini. Bayesian reasoning in high energy physics. Principles and applications. CERN Yellow Report 99-03, 1999.

[23] T. K. Das and BS Teng. Between trust and control: Developing confidence in partner cooperation in alliances. *Academy of Management Review*, 23(3):491–512, 1998.

[24] T.K. Das and B.S. Teng. Trust, Control, and Risk in Strategic Alliances: An Integrated Framework. *Organizations Studies*, 22(2):251–283, 2001.

[25] W.H. Davidow and M.S. Malone. *The Virtual Corporation: Structuring and revitalizing the corporation for the 21st century.* Harper Business, New York, 1992.

[26] Chrysanthos Dellarocas. Goodwill Hunting: An Economically Efficient Online Feedback Mechanism for Environment with Variable Product Quality. In *Lecture Notes in Artificial Intelligence LNAI 2531*, pages 238–252, 2002.

[27] Chrysanthos Dellarocas. The Digitization of Word of Mouth: Promise and Challenges of Online Feedback Mechanisms. *Management Science*, 49(10):1407–1424, 2003.

[28] V. Dignum, H. Weigand, and L. Xu. Agent Societies: Towards Frameworks-Based Design. In *AOSE 2001*, volume 2222 of *Lecture Notes in Computer Science*, pages 33–49. Springer, 2002.

[29] M.j. Dollinger, P. A. Golden, and T. Saxton. The Effect of Reputation on Decision to Joint Venture. *Strategic Management Journal*, 18(2):127–140, 1997.

[30] Wolfgang Faisst. Information technology as an enabler of virtual enterprises: A life-cycle-oriented description. In *Proceedings of the European Conference on Virtual Enterprises and Network Solutions*, Paderborn, Germany, 1997.

[31] K. Fisher, J. Muller, I. Heiming, and A-W. Scheer. Intelligent agents in virtual enterprises. In *Proceedings of the First International Conference on the Practical Application of Intelligent Agents and Multi-Agent Technology*, London, UK, 1996.

[32] I. Foster, C. Kesselman, and S. Tuecke. The anatomy of the grid: Enabling scalable virtual organizations. *The International Journal of High Performance Computing Applications*, 15(3):200–222, 2001.

[33] J. Michael Geringer. Strategic determinants of partner selection criteria in international joint ventures. *Journal of International Business Studies*, 10(6):41–63, 1991.

[34] N. Griffiths and M. Luck. Cooperative Plan Selection Through Trust. In *Multi-Agent System Engineering. Ninth Workshop on Modelling Autonomous Agents*, pages 162–174. Lectures Notes in Articial Intelligence 1647, 1999.

[35] N. Griffiths and M. Luck. Coalition formation through motivation and trust. In *Proceedings of the Second International Joint Conference on Autonomous Agents and Multi-Agent Systems*, pages 17–24, Melbourne, Australia, 2003.

[36] Ranjay Gulati. Alliances and networks. *Strategic Management Journal*, 19:293–317, 1998.

[37] C.P. Holland and A.G. Lockett. Business Trust and the Formation of Virtual Organizations. In *Proceedings of the 31st Annual Hawaii International Conference on System Sciences*, pages 602–610, 1998.

[38] B. A. Huberman and F. Wu. The dynamics of reputations. Technical report, HP Labs, 2002. http://www.hpl.hp.com/shl/papers/reputations/reputation.pdf.

[39] N. R. Jennings and M. Wooldridge. Intelligent agents: theory and practice. *Knowledge Engineering Review*, 10(2):115–152, 1995.

[40] C. Jones, W.S. Hesterly, and S.P. Borgatti. A General Theory of Network Governance: Exchange Conditions and Social Mechanisms. *Academy of Management Review*, 22:911–945, 1997.

[41] C. Jonker and J. Treur. Formal analysis of models for the dynamics of trust based on experiences. In *Proceedings for the Second Workshop on Deception, Fraud and Trust in agent Societies*, pages 81–94, Seatle, USA, 1999.

[42] A. Jøsang. The right type of trust for distributed systems. In *Proceedings of the 1996 workshop on New security paradigms,ACM*, pages 119–131, 1996.

[43] N. Julka, R. Srinivasan, and I. Karimi. Agent-based supply chain management-1: framework. *Computers and Chemical Engineering*, 26:1755–1769, 2002.

[44] J. Liu. and V. Issarny. Enhanced Reputation Mechanism for Mobile Ad Hoc Networks. In *Trust Management: Proceedings for the Second International Conference, iTrust 2004. Lecture Notes in Computer Science 2995*, pages 48–62, Oxford, UK, 2004.

[45] S. P. Marsh. *Formalising trust as a computational concept.* PhD thesis, University of Stirling, Stirling, Scotland, UK, 1994.

[46] M. T. Martinez, P. Fouletier, K. H. Park, and J. Favrel. Virtual enterprise organisation, evolution and control. *International Journal of Production Economics*, 74:225–238, 2001.

[47] D. McCutcheon and F.I. Stuart. Issues in the choice of supplier alliance partners. *Journal of Operations Management*, 18:279–301, 2000.

[48] L. Mui, A. Halberstadt, and M. Mohtashemi. Notions of Reputation in Multi–Agents Systems: A Review. In *Proceedings of the First International joint conference AAMAS*, pages 280–287, Bologna, Italy, 2002.

[49] L. Mui, M. Mohtasemi, and A. Halberstadt. A computational model of trust and reputation. In *Proceedings of the 35th Hawaii International Conference on System Sciences*, pages 188–188, Hawaii, USA, 2002.

[50] N. Nayak, T. Chao, J. Li, J. Mihaeli, R. Das, A. Derebail, and J. Soo Hoo. Role of technology in enabling dynamic virtual enterprises. In *International Workshop on Open Enterprise Solutions: Systems, Experiences, and Organizations, OES-SEO2001*, Rome, 2001.

[51] T. Norman, A. Preece, S. Chalmers, N. R. Jennings, M. Luck, V. Dang, T. Nguyen, V. Deora, J. Shao, W. Gray, and N. Fiddian. Agent-based formation of virtual organisations. International Journal of Knowledge Based Systems, 2003.

[52] M. A. Nowak and K. Sigmund. Evolution of indirect reciprocity by image scoring. *Nature*, (393):573–577, 1998.

[53] E. Oliveira and A. Rocha. Agents advanced features for negotiation in electronic commerce and virtual organisations formation processes. In Dignum F. and Sierra C., editors, *Agent Mediated Electronic Commerce, the European AgentLink Perspective*, volume 1991 of *Lectures Notes in Artificial Intelligence*, pages 77–96. Springer, 2000.

[54] E. Ouzounis. *An Agent-Based Platform for the Management of Dynamic Virtual Enterprises.* PhD thesis, Technical University of Berlin, School of Electrical Engineering and Computer Sciences, 2001.

[55] S. A. Petersen and M. Divitini. Using agents to support the selection of virtual enterprise teams. In *Proceedings of the Fourth International Bi-Conference Workshop on Agent Oriented Information Systems (AOIS2002)*, Bologna, Italy, 2002.

[56] S. A. Petersen and M. Matskin. Agent interaction protocols for the selection of partners for virtual enterprises. In *Proceedings of CEEMAS 2003*, volume 2691 of *Lectures Notes in Artificial Intelligence*, pages 606–615. Springer, 2003.

[57] Alun Preece. A mediator-based infrastructure for virtual organisations. In *Proceedings of Workshop On Agent-Based Approaches to B2B At The 5th International Conference on Autonomous Agents*, Montreal,Canada, 2001.

[58] W. Raub and J. Weesie. Reputation and efficiency in social interactions; an example of network effects. *American Journal of Sociology*, 96(3):626–654, 1990.

[59] P. Resnick and R. Zeckhauser. Trust among strangers in internet transactions: Empirical analysis of ebay's reputation system. Working Paper for the NBER Workshop on Empirical Studies of Electronic Commerce, 2001.

[60] P. Resnick, R. Zeckhauser, E. Friedman, and K. Kuwabara. Reputation systems. *Communications of the ACM*, 43(12):45–48, 2000.

[61] P.S. Ring and A.H. Van de Ven. Structuring Cooperative Relationships between Organizations. *Strategic Management Journal*, 13:483–498, 1992.

[62] D. De Roure, N. R. Jennings, and N. R. Shadbolt. Research agenda for the semantic grid: A future e-science infrastructure. Technical Report of the National e-Science Centre, UKeS-2002-02, 2001.

[63] J. Carbo Rubiera, J. Molina Lopez, and J. Davila Muro. A fuzzy model of reputation in multi-agent systems. In *Proceedings for the Fourth Workshop on Deception, Fraud and Trust in agent Societies*, pages 25–26, Montreal, Canada, 2001.

[64] J. Sabater and C. Sierra. Regret: A reputation model for gregarious societies. In *Proceedings for the Fourth Workshop on Deception, Fraud and Trust in agent Societies*, pages 61–69, Montreal, Canada, 2001.

[65] J. Sabater and C. Sierra. Reputation and social network analysis in multi-agent systems. In *Proceedings of the First International joint conference AAMAS*, pages 475–482, Bologna, Italy, 2002.

[66] Mari Sako. *Price, quality and trust: Inter-firm relations in Britain and Japan.* Cambridge University Press, Cambridge, Great Britain, 1992.

[67] T. Sandholm and R. Crites. Multiagent reinforcement learning in the iterated prisoner's dilemma. *BioSystems*, 37(1996):147–166, 1996.

[68] Todd Saxton. The effects of partner and relationship characteristics on alliance outcomes. *Academic of Management Journal*, 40(2):443–461, 1997.

[69] M. Schillo, M. Rovatsos, and P. Funk. Detecting deceitful agents in artificial societies. *Applied Artificial Intelligence*, 14(8):825–848, 2000.

[70] Carl Shapiro. Consumer information, product quality, and seller reputation. *The Bell Journal of Economics*, 13:20–35, 1982.

[71] S. P. Singh and D. Bertsekas. Reinforcement learning for dynamic channel allocation in cellular telephone systems. In M. Mozer, M. I. Jordan, and T. Petsche, editors, *Proceedings of the Ninth Conference in Neural Information Processing Systems*, pages 974-980. MIT Press, 1997.

[72] R. Sutton and A. Barto. *Reinforcement Learning: An Introduction.* The MIT Press, 1998.

[73] M. Tsvetovat and K. Sycara. Customer coalitions in the electronic marketplace. In *Proceedings of the Agents 2000 Conference*, pages 263–264, Barcelona, Spain, 2000.

[74] A. Wilson. Cognitive factors affecting subjective probability assessment. ISDS Discussion Paper Series 94-02, Duke University, 1994.

[75] L. Xiong and L. Liu. A Reputation-Based Trust Model for Peer-to-Peer eCommerce Communities. In *Proceedings of the IEEE International Conference on E-Commerce (CEC03)*, pages 162–174. Lectures Notes in Articial Intelligence 1647, 2003.

[76] B. Yu and M.P. Singh. An evidential model of distributed reputation management. In *Proceedings of the First International joint conference AAMAS*, pages 294–301, Bologna, Italy, 2002.

[77] G. Zacharia and P. Maes. Trust management through reputation mechanisms. *Applied Artificial Intelligence*, 14(8):881–907, 2000.

[78] G. Zachira, A. Moukas, and P. Maes. Collaborative reputation mechanisms in electronic marketplaces. In *Proceedings of the 32th Hawaii International Conference on System Sciences*, pages 8026–8026, Hawaii, USA, 1999.