A classification scheme for negotiation in electronic commerce*

Alessio R. Lomuscio† Michael Wooldridge‡ Nicholas R. Jennings#

† Department of Computing
Imperial College of Science, Technology and Medicine
London SW7 2BZ, United Kingdom
A.Lomuscio@doc.ic.ac.uk

‡ Department of Computer Science
University of Liverpool
Liverpool L69 7ZF, United Kingdom
M.J.Wooldridge@csc.liv.ac.uk

Department of Electronics and Computer Science
University of Southampton
Southampton SO17 1BJ, United Kingdom
nrj@ecs.soton.ac.uk

Abstract

In the last few years we have witnessed a surge of business-to-consumer and business-to-business commerce operated on the Internet. However, most current electronic commerce systems are little more than electronic catalogues that allow a user to purchase a product under predetermined and inflexible terms and conditions. We believe that in the next few years we will see a new generation of electronic commerce systems emerge, based on *automated negotiation*. In this paper, we identify the main parameters on which any automated negotiation depends. To show the applicability of our classification framework, we use it to categorise a representative sample of some of the most prominent negotiation models that exist in the literature.

*This research was supported by the EPSRC under grant GR/M07076. A previous version of this work appeared in [LWJ00]. This paper has benefited from the comments of several anonymous referees.

1 Introduction

All that hard work has paid off and your paper has been accepted at that conference. You go through the reviews with a smile but get quickly worried when reading the instructions for the authors:

> The venue of the conference, Marina del Sol, is a popular tourist resort and is therefore expected to be busy during the conference. We advise you to book your flight as early as possible.

It will take hours of effort to find a reasonable fare for your flight! But will it? Why not try "Negotia", that latest shopping agent that everyone was talking about last night? Indeed you launch the program, you quickly go through the menus to enter all your personal details and specify your preferred airline, your budget, how flexible you are around the dates and specify that your partner might join you there for a few days after the conference.

After half an hour you receive a report from Negotia with the standard fares together with flight times, stop-overs and availability. A good fare seems to be around 600Euros for travelling immediately before and after the conference. Availability is already very limited on the economy fares.

Included in the report there are also some special arrangements that Negotia has been able to negotiate on your behalf with the airlines computer systems. These include:

- A-Airlines. Out-bound 13 July; in-bound 18 July; cost 500Euros.

  This is a discount of 15% on their standard price that can be obtained by flying two days before the conference. This discount has been offered to you as a regular flyer with A-Airlines because of an unexpected lack of demand for that date. Tickets must be paid in full within 15 days.

- B-Wings. Out-bound 15 July; in-bound 25 July; cost 420Euros.

  A discount of 25% on their standard price. This offer has been negotiated by Negotia by offering to buy two seats (one for your partner) rather than only one. It has been negotiated that only the return flight must be taken together. B-Wings, a major carrier, is currently trying to improve its relatively small load factor in view of the annual general meeting and they are therefore offering special conditions for multiple purchases. The tickets must be paid in full within the next 48 hours.

- C-Air. Out-bound 15 July; in-bound 18 July; cost 350Euros.

  An incredible offer from a charter airline. The ticket must be booked immediately and it is a no-change/no-refund fare. This discount has been obtained by Negotia because you specified that you are able to pay immediately for future travel and the airline is trying to finance a major restructuring operation by pre-selling their flights far in advance. After contacting a third-party watchdog, Negotia is also warning you that the airline has a history of delayed and cancelled flights.

You decide that the second option is the best, and instruct Negotia to put two seats on hold — you will confirm them tomorrow. You save the search Negotia has done for you, quit the program and start thinking about something else.

1.1 Agent-mediated electronic commerce

The term "electronic commerce" (e-commerce) generally denotes an advanced step of modern commerce in which the figures of buyer and seller are replaced by electronic entities [BB97]. It is widely believed that e-commerce will reduce costs for enterprises and provide customers with better bargaining tools. For example, automation can exploit short-terms contracts in combinatorially complex settings [San99] and allow extensive comparisons to be performed between a wide variety of goods.

Given the nature of the tasks involved in e-commerce, it has been suggested that agent technology [WJ95] will play an important role in its development (see [GMM98] for a survey). In such settings, agents will play the roles of buyer, seller, mediator [RNSP97], facilitator, and information provider. In so doing, agents will automate part or all the business tasks involved in e-commerce.

Most current e-commerce applications (servers and clients) may be classified as *first-generation* systems. That is, servers are connected to the Internet and they let the user browse through catalogues containing well-defined commodities (e.g., flights, books, compact discs, computer components, etc.) and make purchases typically by means of a credit card transaction. In this case the system's client can be any basic Internet browser and therefore the users' actions, being reduced to select/accept choices, are somewhat limited. However, several more advanced forms of first generation clients, such as Pricerunner [PCR], Pricewatch [PCW], and countless of others have also been introduced. The main feature of these clients, frequently called *shopping assistants*, is the ability to perform *merchant brokering* on behalf of the user. Merchant brokering is the stage in the consumer buying behaviour where the buyer has chosen the product and is looking for the best deal to purchase it[1].

Other shopping assistants (e.g., PersonaLogic [Per] and Firefly [FF]) offer their services earlier in the transaction by helping the buyer in the *product brokering* stage. This is the stage in the consumer buying behaviour where a decision is made about what product the buyer needs. Yet more advanced shopping agents, notably Tete-a-Tete [TAT], attempt to follow the buyer in more than one phase of the transaction by performing both product brokering, merchant brokering, and even some primitive forms of negotiation.

However, a key problem with all of these first-generation systems is that they are focused on just one aspect of the transaction: price. Although price

[1]The six main stages of consumer buying behaviour can be broadly thought of as: need identification, product brokering, merchant brokering, negotiation, payment and delivery, and service and evaluation. See [Nic66, HS69, EB82, Bet79] for more details on consumer buying behaviour.

is clearly important, it is very often the case that neither customers nor enterprises have price as their only concern. Enterprises usually aim to acquire or maintain regular customers, whilst these customers are, in turn, often interested in non-price aspects of the purchase (such as warranty, delivery time, and track record of the vendor or manufacturer). To cope with this need, more advanced forms of shopping assistant are required. Such software will have to deal with (interact over) the complete range of issues that are relevant to their user.

At this time, several more complex shopping assistants are emerging (such as priceline [PCL] and hotwire [HTW]) that allow the user to choose parameters other than just price such as travel dates, quality of the hotel, and number of max connections for a flight. Nevertheless it is still fair to say though that the importance given to one parameter - price - is paramount. Indeed, with these assistants the user bids for a price given the dates required, and no form of tradeoff among the parameters is present.

When faced with the need to reach agreement on a variety of issues, humans make use of *negotiation*. Similarly, we believe that automated negotiation will become the dominant mode of operation for shopping assistant agents. This *automation* of negotiation can significantly reduce negotiation time (making large volumes of transactions possible in small amounts of time) and can also remove some of the reticence of humans to engage in negotiation (e.g., because of embarrassment or personality). For these reasons, the formalisation of negotiation has received a great deal of attention from the multi-agent systems community throughout the past two decades [DS83, KL89, KM90, RZ94, Kra01, JFL+01]. Through such endeavours, it is becoming possible to develop practical automated negotiators. This will, in turn, facilitate the development of fully-fledged shopping assistants that are able to perform negotiation on behalf of users.

Negotiation strategies and their corresponding properties depend heavily on the specific characteristics of the scenarios under consideration. For example, economists and game theorists distinguish between scenarios in which the value of a good is common to all the agents (for example a typical shrink-wrapped good such as a CD), and those in which it is a private characteristic, which differs from agent to agent (such as a painting, which different agents may value differently) [OR94]. Whether a good has a common value, or whether its value differs from agent to agent, the valuations themselves may be common knowledge or private. These are but two illustrations of parameters upon which any negotiation mechanism crucially depends. However, as we shall see, many other parameters also play a key role in shaping the negotiation. Against this background, the principle aim of this paper is to define (for the first time) the *negotiation space* for electronic commerce. That is, we aim to identify the possible parameters that can be used to classify any negotiation mechanism for electronic commerce. Such a classification is an important step for the development of more sophisticated shopping assistants because it *defines* and *delimits* the design space for agent interactions. Thus, the classification in this paper should be seen from an analytical point of view. It does not offer any immediate practical application, but, instead, it aims at providing a conceptual framework within which protocols and strategies for negotiation can be classified and reasoned about. We argue that the components of the negotiation space that we identify constitute a complex space of possible games; indeed, we have so far been unable to find two distinct negotiation scenarios that would be described in the same manner in the classification. We illustrate our taxonomy by classifying a number of the most prominent models of automated negotiation that have been proposed in the general literature on the subject.

The remainder of this paper is structured as follows. In section 2 we discuss and classify the main variables upon which negotiation in electronic commerce depends. In section 3 we classify a number of references from the literature according to our taxonomy. We conclude in section 4 with a discussion on the future direction of automated negotiation in e-commerce.

2 The Negotiation Space

In its broadest sense, automated negotiation involves the design of high-level protocols for agent interaction. The theme of negotiation is present in many different fields and, as a result, several definitions have been proposed in the literature [RZ94, Lee96]. It is not within the scope of this paper to discuss the merits of these definitions, and so we simply adopt the following definition:

> Negotiation is the process by which a group of agents communicate with one another to try to reach agreement on some matter of common interest.

Two basic components are important when designing an automatic negotiation system: the *negotiation protocol* and the *negotiation strategies*. The former specifies the "rules of encounter" between the negotiation participants. That is, the protocol defines the circumstances under which the interaction between the agents takes place: what deals can be made and what sequences of offers are allowed. In general, agents must reach agreement on the negotiation protocol to use before negotiation proper begins. Reaching agreement on a negotiation strategy may itself be done by *meta-level* negotiation, but we do not consider this process here.

An agent's negotiation *strategy* is the specification of the sequence of actions (usually offers or responses) the agent plans to make during the negotiation. There will usually be very many strategies that are compatible with a particular protocol, each of which may produce a very different outcome. For example, an agent could concede at the first round or bargain very hard throughout the negotiation until a private timeout is reached. It follows that the negotiation strategy an agent employs is critical with respect to the outcome of the negotiation. It should also be clear that strategies which perform well with certain protocols will not necessarily do so with others: the choice of strategy to use is thus a function not just of the specifics of the negotiation scenario, but also of the protocol in use.

A *negotiation mechanism* consists of a negotiation protocol together with the negotiation strategies for the agents involved. There are some properties that are generally considered desirable for a negotiation mechanism (cf. [San99, pp12–14] and [RZ94, pp20–22]).

Computational efficiency: Ideally we seek a negotiation mechanism that is computationally efficient. Although users constantly have at their disposal faster machines, there will be little scope for negotiation algorithms that are, say, EXPTIME-complete[2].

Communication efficiency: All things being equal, we would rather have a mechanism that handles communication among the agents in an efficient way. Broadcasting to all the agents in the system, for example, may not be ideal in this respect.

Individual rationality: A mechanism should be individually rational for all the agents involved. In other words, it should be in an agent's independent interest to participate in negotiation — otherwise, a rational agent will not. If considerations of group utility need to be taken into account, they can be made a component of each agent's private utility.

Distribution of computation: Mechanisms that distribute the computation over the agents involved are preferable to ones in which one server is performing all the computation for the whole system. This is preferred for many reasons, including the desire to avoid the disruptive effects of a single point of failure, and performance bottlenecks.

Pareto efficiency: An outcome is Pareto efficient if there is no other outcome that improves the lot of one agent without making another agent worse off. All other things being equal, Pareto efficient solutions are preferred over those that are not.

Even with these broad guidelines in mind, many different mechanisms can be designed. These mechanisms very much depend on the specific characteristics of the negotiation space, i.e., the set of all possible negotiation games. In the next section we investigate this space in more detail in order to try and separate out the key orthogonal components.

2.1 Parameters of the negotiation space

As already noted, the design of an appropriate negotiation scenario crucially depends on a number of parameters, which can be seen to generate a space of possible negotiation scenarios. A preliminary attempt to classify this space has already been proposed by Wurman *et al.* [WWW98]. In particular, they

[2] One should be careful not to confuse the computational complexity associated with *designing* a negotiation strategy compared to the complexity of *executing* a strategy. The former can usefully be carried out offline beforehand, thereby reducing the computational costs of the agents. The latter obviously needs to be carried out at run-time and must therefore be manageable.

identify variables upon which auction servers can be designed. Some of the parameters introduced there are also present in our classification; but since we are not only concerned with auctions, our classification is more general.

We divide the parameters on which the negotiation can take place into the following broad characteristics:

- cardinality of the negotiation;
- agent characteristics;
- environment and goods characteristics;
- event parameters;
- information parameters; and
- allocation parameters.

Under this characterisation, an agent's negotiation strategy is simply one of its characteristics. We structure the characteristics of the protocol by analysing in more detail some of the key parameters such as events, information parameters, and allocation. The protocol as a whole results from the definition of all these.

Before discussing the parameters of our classification in some detail we would like to point out that this is just one possible way of understanding and classifying literature and the mechanisms there discussed. There surely are other criteria upon which similar classifications can be obtained. Still we find the one that we describe below rather natural, and quite adequate for the task.

2.1.1 Cardinality of the negotiation

We can distinguish between the cardinalities of the negotiation domain itself, and of the interactions that take place, as follows:

- negotiation domain: single-issue or multiple-issue; and
- interactions: one-to-one, many-to-one, many-to-many.

The domain of negotiation can be thought of simply as the set of tuples over which the agents negotiate. These tuples represent the issues on which an agreement is to be made. Elements of these tuples correspond to issues such as price, quality, warranties, delivery time, and so on (see for example [FSJ98]). In the case of single-issue negotiation (for example where the only issue is price) the tuples are singletons. In the case of multiple-issue negotiation, the different issues might be related by some publicly agreed utility function. This makes the process of bid formulation easier because each agent is then able to compute the utility function of each agent involved in the negotiation [VJ00].

Interactions between agents can be classified in terms of the number of agents participating in the negotiation. *One-to-one* negotiation (where one agent is negotiating with exactly one other agent) is important for both theoretical and practical reasons. Theoretically, it is important because of the technical difficulties that this apparently simple setting provides; for example most games that result from one-to-one negotiation can be proven to have multiple equilibria [OR94], and a naive application of game-theoretical tools is therefore not possible. Practically, it is important because of the emerging role of one-to-one relations associated with business-to-business e-commerce scenarios.

Many-to-one negotiation (where many agents negotiate with just one agent) is the standard setting of auctions, which have been popular on the Internet for some time now [QXL, EBA]. In this setting, one agent plays the role of the seller, while many play the role of the buyers. Other settings are possible such as having many sellers and one buyer, as is the case for price comparison engines.

Finally, *many-to-many* negotiation (where many agents negotiate with many other agents) constitutes the most complex scenario. The continuous double-auction is the most complex of these scenarios that is actually in use [FR93].

Notice that whether a given negotiation is one-to-one, one-to-many, or many-to-many is not determined simply by the number of agents. For example, it is conceivable to have a large set of agents, but arrange them such that any negotiation takes place on a one-to-one basis.

2.1.2 Agent characteristics

Agents are the computational entities that participate in the negotiation process. Each agent is assumed to be capable of rating its preferences, so that it can evaluate and choose between different deals. We can further characterise agents depending on their:

- role;
- rationality;
- knowledge;
- commitment;
- social behaviour; and
- bidding strategy.

Agents can interact in the negotiation as buyers, sellers, or both (as in a continuous double-action). This is their type: the role they play in the negotiation. For most negotiation domains, buyers and sellers are obviously the leading figures, but in other scenarios, (auctions for example), intermediaries can have an important role.

Rationality can be modelled to be *perfect* or *bounded* [OR94]. The assumption of perfect rationality generally amounts to the agents being able to perform arbitrarily large computations in constant time. In all practical scenarios agents do not have the ability to perform such calculations, and they are forced to bid or withdraw on the basis of finite computations. Thus, negotiation models that assume perfect rationality (e.g., many of those coming from game theory) have to use approximations in practice, whereas models that explicitly assume bounded rationality are more realistic in this sense.

Agents have knowledge about the goods they bid for, and possibly some knowledge about how other agents value the same and other goods. Depending on how such knowledge is distributed, agents may choose different bidding strategies. Whether or not agents hold private information is crucially important to design the agent's bidding strategy. For example, internal deadlines, and the valuation of the opponent's utility functions can be important parameters of an agent strategy.

Various levels of commitment can be present. For example, after having made an offer, agents might be obliged to stop bidding for similar goods until an acceptance or counter offer is received. Alternatively, agents can have the mechanism and process for reneging upon contracts built into the original negotiation (e.g., [San99, FSJ98]).

Agents can act as individually self-interested entities, as altruistic units of a society, or they can strike a balance somewhere in between. Furthermore, whatever their attitude, they can play as distinct entities or effectively pursue team-formation (or coalitions) in order to get better deals when convenient [VJ00].

An agent's bidding strategy is ultimately the component that decides about placing or accepting offers, making counter offers or withdrawing from negotiation. Although conceptually independent from the other parameters, it is reasonable to imagine the bidding strategy to be somehow related to the commitment, the knowledge, the rationality, and the social behaviour of an agent [FSJ98]. A standard assumption is that agents are individually rational, in that their bidding strategy would be against accepting deals that result in them being worse off than if they do not enter into any deals.

2.1.3 Environments

The negotiation environment can either be static or dynamic. A static environment is one whose variables (e.g., prices of important commodities) are constant over time; a dynamic environment is one in which these change over time. Clearly, there are various degrees of how static or dynamic environments can be, and this is therefore a simplification. Even so, the dynamism of the environment might affect the design of the utility function of the agents in a subtle way. The utility functions of the agents reflect their preferences; so, while in a static environment, one can imagine an agent that does not learn during the process and maintains a fixed utility function, this behaviour would be less likely to produce a positive payoff in a very dynamic environment.

The characteristics of the goods also crucially define the negotiation protocol:

- private/public value of the goods; and

- nature of the goods.

Agents can value the good differently depending on whether it is intended for private use (e.g., a cake) or whether its value depends on how the other agents value it (e.g., bonds). Frequently, both private and public valuations play a part. For example, when buying a car one has to consider both one's own preferences and how the car will preserve its value over time in case one should be interested in selling it.

The object of the negotiation can either be a discrete or continuous set (of goods). For example, a negotiation over a deadline (delivery dates, completion of contracts, etc.) is a negotiation over an uncountable set because time is intrinsically continuous. Usually, a simplification is made with respect to the granularity of the domain, and countable or even finite sets are used. For example, with a deadline one could use, days, hours, or minutes, depending on how fine the model needs to be. This simplification considerably reduces the size of the negotiation space. The simplest case arises when the set is a singleton.

2.1.4 Event parameters

The negotiation protocol is mainly influenced by the ways in which the offers and other events that take place during the negotiation are regulated. Indeed, what follows forms an important part of the specification of the protocol of the negotiation. We can distinguish between:

- bid validity;
- bid visibility;
- clearing schedule and timeouts; and
- quotes schedule.

The first item specifies an important part of the protocol: the criteria for validity of the bids. To be valid, bids often have to be offered at an appropriate time and must satisfy some constraints on their value. For example, in an English auction bids can be made when the auctioneer is calling for bids and must be progressively higher in value. Similarly, procedures for placing bids might be present during a negotiation.

The visibility of the bids is only relevant in the case of many-to-one or many-to-many negotiations. At the extremes, bids can be private messages passed between buyer and seller, or broadcast to all agents. Alternatively, we can have configurations in which only subsets of agents see some selected messages (especially useful if coalition formation is permissible).

A "clear" is the event producing a (temporary) allocation between buyer and seller. Clears can be scheduled at random times or following some other events (such as the first offer matching the buyer's request). For example, during the bidding phase of an English auction each round terminates with a temporary allocation of the good being auctioned to the prospective buyer that meets the auctioneer's call.

Timeouts determine the closing of the negotiation, therefore they transform clears into "final clears", i.e., a final agreement between buyer and seller about the transaction. Clears and timeouts also depend on the allocation parameters (see below).

If third-party quotes are generated during the negotiation process, their number and frequency are also part of the events and they need to be regulated. Intuitively, an excessive number of requests for quotes can significantly slow down the negotiation mechanism.

2.1.5 Information parameters

Both before and during negotiation, information other than bids may pass between the negotiation participants. These messages are either information that can help buyers and sellers reach agreements, or information that can help limit the noise produced by agents trying to buy and sell goods. Such messages can be beneficial in order to save computational time of the agents.

Among the many possible useful messages, we can distinguish between:

- price quotes;
- transaction history; and
- arguments.

Quotes generated by sellers in response to requests by prospective buyers for an indicative price from a seller before starting a negotiation can be useful to all parties, as they can reduce negotiation time. A similar mechanism can be in place in case of sellers asking for possible bids from possible buyers.

The history of similar transactions can also be requested, or unilaterally provided, by a seller agent in order to give credibility to the information it is offering to potential buyer agents. It can be guaranteed by trusted third-parties. Together with third-party quotes, transaction histories can form the basis for argumentation-based negotiation. Such protocols [SJNP98, KSE98] aim to bridge the gap between how negotiation is performed in human and artificial societies. In the former, expert human negotiators often focus on the reasons why an offer is not acceptable and try and persuade their counterpart of the characteristics that an agreement will have to include. In the latter, artificial negotiators traditionally are only able to propose offers to the counterpart without being explicitly able to motivate an agreement. This is often seen as a severe limitation and can limit the flexibility of the negotiation. Specifically, accompanying arguments of the following form provide the negotiatiors with important contextual information that can assist their decision making: "I cannot deliver the required good any quicker because I have to order it from the warehouse", "this is my final offer, take it or leave it", or "I'm not paying this price for the good because last time you sold it to me for a lower price".

2.1.6 Allocation parameters

The allocation parameters only apply in many-to-one and many-to-many scenarios. They govern the winner of an auction when more than one agent has shown an interest in the good.

Allocation is studied in auction theory [Vic61]. The Mth and $(M+1)$-th price allocation policies cover most scenarios, where M is the number of received bids. Note that if the negotiation is multi-issue, in order for this mechanism to work there has to be a commonly agreed function that "weights" the different issues producing a value of utility. Without this, it would not be possible to order the offers as required by the allocation policy.

3 Applying the classification scheme

Having specified the dimensions of automated negotiation research, this section seeks to demonstrate the efficacy of the classification framework. To this end, a number of prominent negotiation models are placed within the framework. The exemplar models were chosen to provide a representative sample of work on automated negotiation. In doing this, our aim is *not* to provide an exhaustive classification of all the extant models, but rather it is to demonstrate models that have either typical or extreme values on the key dimensions of our classification scheme. Thus, the exemplar models show how the classification scheme can both group seemingly disparate models (because they are similar on one or more dimensions) and differentiate seemingly similar models (because they differ on one or more dimensions).

3.1 One-to-one Negotiation in Cooperative Domains [ZR89]

In their influential paper [ZR89] Rosenchein and Zlotkin introduce and analyse formally *task oriented domains* (TODs) and encounters within TODs. TODs incorporate the set of all possible tasks for the agents and a monotonic cost function from subsets of the tasks. An encounter within a TOD is a distribution of tasks for all the agents and a pure deal within an encounter is a redistribution of tasks.

Rosenchein and Zlotkin start from the assumption of considering agents as self-interested utility maximisers (i.e., cost minimisers) and study the properties of the *monotonic concession protocol* (MCP). The MCP operates on a space of possible deals which are both individually rational (no agents are worse off in one of the possible deals than in their initial condition) and Pareto optimal. Negotiation proceeds in rounds, whereby on every round each agent either concedes (i.e., suggests a deal which gives a better utility for the other agent and a less favourable utility to themselves) or offers the same deal offered in the previous round. Negotiation terminates either when both agents are happy with the deal they have reached or when both agents have failed to concede. So, although both agents would benefit from any deal resulting from this protocol, an agreement is not guaranteed.

Zlotkin *et al.* extensively analyse the properties of a negotiation strategy (called the *Zeuthen strategy* from [Zeu30]) with respect to the MCP. The strategy is based on the degree of loss that agents suffer from not reaching a compromise and it can be described as follows. For any step t in the negotiation and for each agent i, define a degree of willingness to risk a conflict, $risk_i^t$ as

$$risk_i^t = \begin{cases} 1 & \text{if } utility_i(\delta_i^t) = 0 \\ \frac{utility_i(\delta_i^t) - utility_i(\delta_j^t)}{utility_i(\delta_i^t)} & \text{otherwise,} \end{cases}$$

where δ_i^t denotes the deal proposed by agent i at time t, and $utility_i(\delta)$ denotes the utility of deal δ to agent i (the index j refers to the other agent).

The Zeuthen strategy works as follows:

- At the first round each agent offers the best deal (according to their own utility) in the negotiation set.
- At round t, agent i calculates $risk_i^t$ and $risk_j^t$. If $risk_i^t \leq risk_j^t$, agent i then proposes the minimum deal that changes the risk balance between the two agents.

The strategy, as it is presented above, is not in a Nash equilibrium; this means that given that one agent is using this strategy, it is beneficial for the other agent to use another different strategy. The reason for this is that if at any step the risks are equal, the agent knowing that the other is using the Zeuthen strategy could benefit from it by relying on the other agent to make the concession.

An extension of this strategy, the extended Zeuthen strategy, is proposed in [ZR89]. The extended Zeuthen strategy is still guaranteed to terminate but it is also in a Nash equilibrium and is guaranteed to maximise the product of individual utilities.

We now analyse [ZR89] in view of the characterisation offered in Section 2.

Cardinality. The negotiation domain is the space of possible tasks for the agents. The set of possible tasks is assumed to be a finite set and the agents negotiate on subsets of this set. So, the negotiation domain is based on a single issue.

The negotiation algorithms are designed for two agents, i.e., one-to-one interaction.

Agent characteristics. The agents are homogeneous in role. So we do not have one buyer and one seller but two agents trying to reach an agreement on a redistribution of tasks in a domain.

The bidding strategy employed is the *extended Zeuthen strategy*.

Agents are assumed to be perfectly rational; in particular they are assumed to be able to compute the negotiation domain before they can enter the negotiation. They are also assumed to be able to compute the utility of any subsets of tasks in constant time.

Agents have perfect knowledge. In particular they know the other agent's utility function and are therefore able to compute the utility of the other agent for any offer.

Agents have total commitment in the negotiation process and when they agree to a deal they will not pull back or restart the negotiation.

Agents are self-interested utility-maximisers. If any global benefit is achieved, this is only a consequence of the specific characteristics of the negotiation domain and protocol. Since there are only two agents in the group, team formation is not possible.

Goods. The value of the tasks in the domain is public. This is because both the utility functions of the agents and the domain of tasks are public.

The negotiation is carried out on a set of tasks, therefore their nature is discrete.

Events parameters. The validity of the offers is regulated by the monotonic concession protocol.

Given that the negotiation is run on a one-to-one basis, bids have to be visible.

Clearing schedules are not present.

Third-party quotes are not applicable.

Information parameters. Price quotes are not present as an agent is able to compute the utility for the other agent of any arrangement in the negotiation space.

Every negotiation is supposed to be a one off. No histories of past negotiations are considered although agents are assumed to have perfect recall about the negotiation currently under way.

Allocation Parameters. Given the one-to-one setting, allocation parameters are not applicable here.

Results proven for the above setting have been extended in the book [RZ94] where scenarios with incomplete information are analysed[3]. In [RZ94] the circumstances under which the domain of negotiation affects the negotiation strategy are also studied. Incomplete information is also studied there and some strategies for maximising private utilities discussed.

3.2 Many-to-many Negotiation in Cooperative Domains [WBK96]

Wooldridge *et al* propose a negotiation-based solution to a common industrial problem: *factory production sequencing* [WBK96]. Briefly, this problem involves a manufacturing plant deciding upon the sequence in which to process products, so as to both satisfy current orders, and minimise overall factory costs.

[3] In [RZ94] *state oriented domain* and *wealth oriented domains* are also introduced. We do not review these here.

They propose that production sequencing may be treated as a *multi-agent negotiation* problem, in which production cells within a factory negotiate over product sequences in order to distribute costs fairly, and hence minimise total factory costs.

This algorithm proposed by Wooldridge *et al.* is a generalised, n-agent extension of the Monotonic Concession Protocol [RZ94] with the Zeuthen strategy [Zeu30]. Negotiation proceeds in rounds, and on the first round, every agent takes an active part by proposing some deal. If a deal has been proposed that makes every agent happy, then negotiation ends successfully. Otherwise, negotiation proceeds to another round, in which some subset of the currently active agents must *concede*. For such conceding agents, there are three possibilities:

- the agent is able to propose a deal that represents a "true" concession, in which case it does so;
- the agent is unable to make a "true" concession, but is nevertheless able to make another proposal, in which case it does so; in this case, the agent will in some sense be "backtracking";
- the agent has exhausted the set of all deals it could propose, in which case it withdraws, and plays no further part in negotiation.

Agents that do not concede in some round put forward the same deal on the next round. In this way, negotiation proceeds with agents conceding and possibly withdrawing, until finally, they find a deal upon which they agree.

In order to select which deal to propose on any given round, each agent initially puts forward its most preferred deal. Subsequently, agents either concede, or else put forward the same deal they proposed in the previous round. In order to decide *who* should concede, agents use the "risk" concept from the Zeuthen strategy — the idea is that the agent with the most to risk from not conceding should be the one to concede. Formally, the risk $risk_i^t$ to agent i at time t is computed as

$$risk_i^t = \begin{cases} 1 & \text{if } utility_i(\delta_i^t) = 0 \\ \frac{utility_i(\delta_i^t) - \min\{utility_i(\delta_j^t) \mid j \in Ag\}}{utility_i(\delta_i^t)} & \text{otherwise.} \end{cases}$$

where δ_i^t denotes the deal proposed by agent i at time t, and $utility_i(\delta)$ denotes the utility of deal δ to agent i.

Whenever an agent puts forward a *new* deal, it uses the following technique:

- it first identifies deals that it has *not* previously proposed;
- using these deals, it then computes the set of deals that improve the lot of at least one agent, without making any other agent worse off;
- using these deals, it computes the set of deal that improve global (in this case, factory) utility;

- using these deals, it finds the most preferred deals that is sufficient to change the balance of risk within the group, so that some other agent will subsequently concede.

Wooldridge *et al.* prove that the algorithm terminates, and show how it can be applied to the production sequencing problem. However, as they note, the techniques proposed are ultimately *heuristics*, as they only provide *rules of thumb* for efficient negotiation. The difficulty in finding more precise negotiation techniques stems from the fact that negotiation occurs between more than two agents.

[WBK96] is an extension to the case of n agents of the work reported in the previous subsection. The only different parameters are therefore the cardinality of the interaction and their bidding strategy.

3.3 Multi-agent negotiation under time constraints [KWZ95]

Kraus et al. analysed a number of problems related to *resource allocation* and *task distribution* among self-motivated, rational and autonomous agents. Their approach differs from the ones presented above in the way it emphasises the importance of the passage of time during the negotiation. This is motivated by two reasons. First, in symmetrical scenarios of resource allocation, the passage of time is a cost for all the parties involved in the negotiation, and, as such, begs the question of researching mechanisms that quickly converge to an agreement. Second, if the scenario is not symmetrical[4], then the negotiation time may be a cost for some agents and a benefit for others. The authors make use of the formal machinery provided by game theory [Bin92, OR94] but differ from [ZR89] in making no assumptions about the agents giving any relevance to the negotiation history when computing the next offer. Still, one important assumption that, in line with [ZR89], Kraus et al. do make is that they use an alternating offer protocol.

The authors are concerned with designing protocols that enjoy properties such as symmetry, efficiency, simplicity, stability (see Section 2), and crucially they tackle the issue of instantaneous outcomes: given that the negotiation cost is taken into consideration, the negotiation should end as quickly as possible, possibly in the first round. Kraus et al. provide an analysis not just for one, but for a family of negotiation scenarios, which share the following characteristics (other assumptions are made and will be discussed later on):

- Even though there may be more than two agents in the model, at every step no more than two agents are active in the negotiation.
- Agents are rational utility maximisers.

Under these assumptions, Kraus et al. focus on the following scenarios:

1. Two agents with complete information negotiating over a shared resource, with no alternative deal.

[4] As it is the case for two agents negotiating over the use of a resource at present entirely controlled only by one of them.

2. Two agents with complete information negotiating over a common goal, with alternative deals available to each of them.

3. Two agents with complete information negotiating over a shared resource, that one has already exclusive access to, with alternative deals available to each of them.

4. Two agents with incomplete information negotiating over a shared resource, that one has already exclusive access to, with alternative deals available to each of them.

5. Several agents with complete information negotiating over a common goal, with alternative deals available to them.

In view of the taxonomy presented in Section 2, the following considerations can be made.

- Cardinality. The negotiation is one-to-one over a single issue.
- Agents. Agents are self-interested, fully rational, fully committed, that can act as buyers or sellers. The bidding strategy and the level of knowledge at disposal varies depending on the scenario.
- Goods. The value of the goods is public.
- Events, information, and allocation parameters are not applicable in this context.

We now turn to analyse in more detail the scenarios above discussed by the authors.

1. Two agents with perfect information negotiating over a shared resource, with no alternative deal.

   In this case, complete information is assumed for the agents, and the (public) utility functions of the agents have the following main conditions[5]. First, disagreement is the worst outcome to every agent; second, both the resource, and the time are valuable to every agent. The nature of the goods is assumed to be a continuum. Under these assumptions, Kraus et al. prove that there exist sub-game perfect equilibrium strategies that guarantee termination of the negotiation process within the first round.

2. Two agents with perfect information negotiating over a common goal, with alternative deals available to each of them.

   This case differs from the one above in two respects: any agent is assumed to be empowered to pull out of the negotiation (each agent has

[5]For other more technical conditions imposed on the agents' utility functions in this and in the following scenarios we refer the reader to the paper.

available some back-up deal), and the goods are assumed to be discrete in nature. As before, a number of assumptions, some of which are reported below, need to be made to guarantee convergence of the mechanism. First and foremost, disagreement is the worst outcome for both agents; second, time is valuable, and actions (the goods to be redistributed) are costly. Third, opting out of a deal is costly over time. By adding three more convergence conditions, details of which can be found in the paper, the authors are able to prove that if the agents use perfect equilibrium strategies an agreement is reached within the first round of negotiation.

3. Two agents with perfect information negotiating over a shared resource, that one has already exclusive access to, with alternative deals available to each of them.

   This is the standard asymmetrical scenario in which one agent is gaining over time, while the other is losing. Apart from this, the setting is equivalent to the one presented above. The condition that the disagreement is the worst possible outcome is still present, and conditions regarding one agent gaining over time while the other is losing are introduced in terms of negotiation costs. Still, the authors show that the threat of the weaker agent of pulling out is enough for the negotiation to end within two rounds at maximum.

4. Two agents with incomplete information negotiating over a shared resource, that one has already exclusive access to, with alternative deals available to each of them.

   This is essentially the point discussed above in which the requirement of complete information is relaxed. Incomplete information games are complex to analyse and require a shift from the notion of sub-game perfect equilibrium to the one of sequential equilibrium (see [KW82] for details). It is beyond the scope of this review to discuss the technical aspects involved in the belief updating process, but we do point to the following three assumptions made by the authors. First, sequential rationality is assumed, i.e., no considerations about future negotiations are made and each agent attempts to maximise its payoff by considering the strategies of the opponent. Second, Bayes's rule is used to update the beliefs and an agent's beliefs should always be consistent with its initial beliefs. The authors prove that, under these conditions, agents using sequential equilibrium strategies force the negotiation to terminate within the first two rounds, and they discuss the extent to which a high probability can be attached to the event of the negotiation ending with an agreement. We refer the interested reader to the paper for more details.

5. Several agents with complete information negotiating over a common goal, with alternative deals available to them.

   Compared to the scenarios above, this involves a shift from one-to-one to many-to-many negotiation, and it concerns a negotiation about work-distribution where agents have complete information. It is assumed that one agent opting out is sufficient for the negotiations to stop (i.e., they have veto power). The mechanism is still an alternating-offer protocol and it is regulated as follows. An offer is put forward by one agent; if it's agreed, it is then implemented, and the negotiation stops; if at least one agent withdraws, then the negotiation stops without agreement; if neither of these applies, another offer is generated by another agent, and another voting round starts. The authors show that if the agents use perfect-equilibrium strategies, then the negotiation stops with an agreement in the first round. To reach this result four main assumptions are used in the utility functions of the agents: actions (the goods on which negotiation is carried out) are costly, time is valuable, agreements are preferred over withdrawals, and opting out becomes more expensive as negotiation progresses.

3.4 One-to-one Negotiation in Mixed Domains [FSJ98]

Faratin et al. developed a generic suite of models that follow a heuristic approach to negotiation in both cooperative and competitive settings [FSJ98, SFJ97]. A heuristic-based approach is chosen to combat the computational short-comings of the more traditional game theoretic models. Generally speaking, heuristic methods acknowledge that there is a cost associated with computation and decision making and so they seek to search the negotiation space in a non-exhaustive fashion. This has the effect that heuristic methods aim to produce *good*, rather than *optimal* solutions. To date, the negotiation model has been deployed in the domains of business process management [JFN+00] and telecommunications network management [FJBS00].

The central concern of this line of work is to heuristically model the agent's decision making during the course of the negotiation. The chosen negotiation protocol is a repeated, sequential model where offers are iteratively exchanged. Under this protocol, the agents are fully committed to their utterances and utterances are private. The space of possible agreements is quantitatively represented by contracts having different values for each issue. Each agent then rates these points in the space of possible outcomes according to some preference structure, captured by a utility function. Proposals and counter-proposals are then offers over single points in this space of possible outcomes, and the search terminates either when the time to reach an agreement has been exceeded or when a mutually acceptable solution, an intersection in the joint space of possible outcomes of both agents, has been reached.

An agent architecture that models the decisions involved in the search for mutually acceptable solutions has also been developed [FSJB99]. Whereas the protocol normatively describes the orderings of actions, the decision making mechanisms describe the possible set of agent strategies in using the protocol. These strategies are captured by a negotiation architecture that is composed of responsive and deliberative decision mechanisms. Decision making with the former mechanism is based on a linear combination of simple functions called tactics, which manipulate the utility of contracts [FSJ98]. The latter mechanisms are subdivided into trade-off [FSJ00] and issue manipulation mechanisms [FJBS00]. The former generates offers that manipulate the value, rather than the overall utility, of the offer. The rationale for the trade-off mechanism is to make proposals that are more attractive to the opponent. This is achieved by providing contracts that are "closer" to the opponent's last offer. The issue manipulation mechanism aims to increase the likelihood of an agreement by adding and removing issues into the negotiation set. The issue manipulation mechanism dynamically alters the structure of the negotiation object, helping to escape local minima in the negotiation dynamics. It does this either by increasing the set of possible outcomes (adding), when negotiation is in deadlock, or, alternatively, removing "noisy" issues that are obstructing the negotiation progress. When taken together, these three mechanisms represent a continuum of possible decision making capabilities: ranging from behaviours that exhibit greater awareness of environmental resources and less to solution quality, to behaviours that attempt to acquire a given solution quality independently of the resource consumption. Moreover, depending on how the tactics and strategies are set up, they can be used to represent cooperative negotiation situations in which the agents seek to find win-win solutions and competitive situations in which agents merely seek to maximise their individual utility.

In terms of the negotiation taxonomy of the previous section, this model is targeted at one-to-one negotiations (although one-to-many can be modelled as a series of interacting one-to-one negotiations) and it deals with contracts that have multiple issues. In terms of their characteristics, the model can be used to describe the negotiation behaviour of both buyers and sellers, the agents have bounded rationality (this is one of the key drivers of the model), the agents are assumed to have very limited information about their negotiation opponent (they do not know their deadlines, reservation values or their utility function), the agents make binding commitments to one another (for which they have to pay a penalty if they renege), and the model allows for both self-interested and altruistic negotiation behaviour. The environment in which the negotiation takes place is assumed to be dynamic and so agents may update their utility functions during the course of the negotiation. The good about which the agents negotiate is assumed to have a private value and the range of issues covered by the negotiation contract are both qualitative and quantitative.

3.5 Argumentation-Based Negotiation [PSJ98]

Parsons et al. developed a negotiation model that allows agents to justify their negotiation positions by sending meta-information (arguments) along with their basic proposals and counter-proposals [PJ96, PSJ98]. In such cases, the role of the argument may be twofold:

- *It allows agents to justify their negotiation stance.*

  An agent might have a compelling reason for adopting a particular negotiation stance. For example, a company may not be legally entitled to sell a particular type of product to a particular type of consumer or a particular item may be out of stock and the next delivery might not be until the following month. In such cases, the ability to provide the justification for its attitude towards a particular issue can allow the opponent to more fully appreciate an agent's constraints and behaviour.

- *It can be the basis upon which agents persuade one another to change their negotiation stance.*

  Agents sometimes need to actively change their opponents' agreement space, or its rating over that space, in order for a deal to be possible. In such cases, agents seek to construct arguments that they believe will make their opponent look more favourably upon their proposal. Thus, arguments seek to identify opportunities for such change (e.g., a car salesman throws in a stereo with a car to increase the value of the good), create new opportunities for change (e.g., a car salesman adds a new dimension to the rating function by highlighting the car's novel security features) or modify existing assessment criteria (e.g., a car salesman gets the buyer to change its evaluation function by convincing him that security is more important than high speed).

In both cases, negotiators are providing arguments to support their stance (hence *argumentation-based negotiation*). Thus, in addition to generating proposals and counter-proposals, the negotiator is seeking to make the proposal more attractive (acceptable) by providing additional meta-level information in the form of arguments for its position. The nature and types of the arguments can vary enormously. However, common categories include: threats (failure to accept this proposal means something negative will happen to you), rewards (acceptance of this proposal means something positive will happen to you), and appeals (you should prefer this option over that alternative for some reason) [SJNP98]. Whatever its precise form, the role of the supporting argument is either to modify the recipient's region of acceptability or its rating function over this region. In so doing, arguments have the potential[6] to increase the likelihood and/or the speed of agreements being reached[7]. In the former case, by persuading agents to accept deals that they may previously have rejected. In the latter case, by convincing agents to accept their opponent's position on a given issue (and to cease negotiating over it).

Concentrating specifically on the model of Parsons et al., the way in which argumentation fits into the general negotiation process was defined in [SJNP98] where a simple negotiation protocol for trading proposals was augmented with a series of illocutions which allow for the passing of arguments. It is possible to think of the passing of an argument using one of these moves as marking a transition from the negotiation protocol to a separate argumentation

[6]Poorly designed argumentation systems also have the potential to increase the length of the negotiation as the various merits of arguments and counter-arguments are debated. However, poor design of the other aspects of the negotiation technology can have similarly adverse effects and so long negotiations are not something specific to the argumentation-based negotiation.

[7]For example, if agents prefer arguments that are more likely to lead to an argument (which requires some metric on the agreement space) it is possible to prove that argumentation leads to quicker agreement [Toh97].

protocol which defines the rules of the game for carrying out an argument dialogue (possible protocols for such dialogues are suggested in [AMP00]). When the argument dialogue terminates, the agents make the reverse transition and pick up the negotiation dialogue once again. The exact argumentation mechanism employed is logic-based [PJ96] and builds on work in argumentation as an approach to handling defeasible reasoning. This makes it possible for agents to handling contradictory statements (which frequently occur during arguments) without collapsing into triviality, and allow conflicting arguments to be resolved. Using argumentation in real agents (as opposed to simple collections of logical statements) means handling the complexities of the agents' mental attitudes, and the integration of the argumentation mechanisms into a complex agent architecture. These issues were discussed in [PSJ98], where it was shown how to augment a standard model of argumentation to work for agents which reason using beliefs, desires and intentions. This model was implemented using a multi-context system to represent the agent's various mental attitudes [SSPJ99].

In terms of the negotiation taxonomy of the previous section, the model deals with multiple negotiation issues in one-to-one settings. The model is symmetric in that it can be used for both the buyers and the sellers and it assumes the agents are perfectly rational. The agents have limited information about their negotiation opponent (they do not know their deadlines, reservation values or rating functions). The model can be used for cooperative negotiations (that seek to develop a joint plan of action) or competitive ones (in which agents seek to maximise their individual utility). The negotiation goods are assumed to be private value and can cover a range of qualitative and quantitative issues. There are two main types of event parameter: those that relate to the negotiation protocol itself (and this is similar to the model of Faratin et al.) and the argumentation protocol (which is the protocol that determines how the arguments can be exchanged). The model allows for a rich and flexible set of information parameters to be exchanged between the agents during the course of the negotiation; indeed this is precisely what the arguments are.

3.6 Persuasive Negotiation [Syc91]

Sycara developed a framework for computer supported goal conflict resolution through negotiation and argumentation. The work was presented in a series of papers, including [Syc87, Syc88, Syc89, Syc91]. The conceptual model was implemented in the PERSUADER system, an integrated computer program that can act as mediator in disputes relating to the labour market.

The framework distinguishes itself from works such as [ZR89] described above as it assumes non-cooperative agents with incomplete information. Indeed, each agent is modelled by its own goals, utility function, and beliefs about each other's utility functions and beliefs. The framework aims at tackling negotiation scenarios in which agents have goals that conflict with each other. In this context, Sycara argues that, since the problem is equivalent to one of a planner for resolution of goal conflicts, it shares with it some common characteristics. In particular, the planner should aim at being *iterative*:

it should handle feedback from the parties involved, and use these to generate improved proposals; it should be able to evaluate which counterproposals might lead towards convergence, and it should be able to recognise changes in the environment. If the agents' utilities are not rigid, then the planner should also be able to generate *persuasive arguments* to put to the agents in order to increase the chances to reach an agreement.

The PERSUADER system was comprised of three agents, two negotiation parties and one mediator, which is engaged in parallel negotiation with both parties. The offers being exchanged are plans, composed of actions that the proposing agents will follow, provided that the other agent agrees. Essentially, the PERSUADER system executes the following loop. The mediator generates a first proposal which is put forward to the parties. These can agree to the proposal or disagree, and motivate their disagreement. If agreement has not been reached the mediator will generate a counter-proposal, which can either be a revised proposal to put forward, or an argument put to the disagreement party aimed at changing its position. In the following we briefly review these three steps. In order to perform all of these tasks PERSUADER integrates techniques from Case-Based reasoning [Ham86] and decision theory. More details can be found in [Syc88, Syc87, Syc89, Syc91].

The initial offer is generated using Case-Based reasoning on a promising instance of similar negotiation cases stored in memory. We do not review here how this process is performed. If previous cases are either not available, or not applicable to the instance under consideration, then PERSUADER uses *preference analysis* [Syc87] (based on multi-attribute utility theory [KR76]) to compute a first proposal.

If the proposal is not accepted by one of the parties, an attempt to persuade the refuting party is made by PERSUADER. This is done by analysing and answering the comments received on the rejected proposal. The persuasive argumentation module works as follows. The agent's utility function is assumed to be a linear combination of several variables; PERSUADER produces arguments either to modify the importance an agent gives to some issue, or to change the value of the issue itself. In order to achieve this, PERSUADER maintains and constantly updates a model of the agents' goals and beliefs. Essentially, the system will first try and introduce new goals in the persuadee's utility function; if this fails, it will try to substitute some of its goals; if this in turn fails, it will try and motivate why some of the goals should be abandoned. Should all these attempts fail, PERSUADER tries to motivate a change in the reservation price in the agent's decision function; to this end it will also use breakoff threats. We refer the interested reader to [Syc91] for more details. If the phase of persuasive argumentation terminates without success, then the system, by using techniques related to the generation of the first offer, produces an amended offer to put before the parties, and the loop can restart.

If we consider the taxonomy of Section 2, the framework can be described as follows. PERSUADER is flexible with respect to the cardinality of the negotiation and it allows multi-issue negotiations. The system is structured with two agents negotiating among each other in a one-to-one setting with the help of a mediator that facilitates agreement in the way described above. Hence, there are three agents in the model. Two are the conceptually equivalent negotiating parties, one is the mediator; the rationality of all of them is bounded. The knowledge of the two agents and of the mediator is limited, and each agent holds private information. The model focuses on the negotiation part of the relation between agents and as such it does not make assumptions about the commitment of the agents involved. But since no agents in PERSUADER reason in terms of agents pulling out of agreed deals, this is equivalent to full commitment. Every agent is self-interested, and the bidding strategy it employs to accept or reject proposals is not known in advance to the persuader. (Note that some assumptions on the linearity of the utility functions are made; we refer the reader to the papers for more details.)

The environment is dynamic, and indeed the persuader periodically checks for modifications in the outside world for revising its own model of it. The tasks on which the negotiation takes place have private value, and the nature of them can be measured both by discrete and continuous quantities.

The event parameters are all handled by the mediator. There are no restrictions to the validity of the bids, although some offers may not be available to all the parties involved because they are handled exclusively by PERSUADER in trying to reach an agreement with one party before putting the offer to the other. There are no a-priori timeouts but this can be introduced as a valid negotiation objective.

PERSUADER is very flexible in terms of what information parameters it can include; indeed the persuasive module can handle messages regarding price quotes, past history, etc.

4 Conclusions

The last few years have witnessed a rapid expansion of business carried out online. While current e-commerce systems offer advantages to both consumers and merchants, it is often the case that they amount to little more than electronic catalogues on which credit card payments can be arranged online. However a new, more flexible, class of e-commerce systems are starting to emerge. In such systems, automated software agents participate in trading activities on behalf of their user. Here we focused specifically on automated negotiation since we consider this to be the key type of interaction in such systems. Automated negotiation, like its human analogue, is a very broad and encompassing discipline. For this reason, it is important to understand the dimensions and range of options that are available. As a preliminary step in this direction, we have developed a classification scheme that is specifically targeted at automated negotiation for e-commerce settings. This classification scheme was illustrated on a diverse range of prominent negotiation models and the results are summarised in Tables 1 and 2. The purpose of this scheme is to provide, for the first time, a comprehensive and systematic basis for objectively comparing and contrasting different negotiation models. Such a classification scheme is vital for developers of second generation e-commerce systems since it provides a means of differentiating competing alternatives for the choice of which model of automated negotiation to exploit.

Having specified the negotiation design space, the next step is to identify the tools and techniques that are appropriate for specific regions of this space. Thus, we believe that no one technique or method of approach will come to dominate in this field [JPSF00]. Rather, there will be a range of alternatives that may be selected according to domain circumstances. The act of mapping the negotiation space in this manner will produce a design repository for interactions in agent-mediated electronic commerce. This repository will then constitute a valuable resource that will enable design expertise and know-how to be shared between developers.

References

[AMP00] L. Amgoud, N. Maudet, and S. Parsons. Modelling dialogues using argumentation. In *Proceedings of the 4th Int Conference on Multi-Agent Systems*, Boston, USA, 2000.

[BB97] J. P. Bailey and Y. Bakos. An exploratory study of the emerging role of electronic intermediaries. *International Journal of Electronic Commerce*, 1(3):7–20, 1997.

[Bet79] J. Bettman. *An information processing theory to consumer choice*. Addison-Wessley, 1979.

[Bin92] K. Binmore. *Fun and games*. D. C. Heath and Company, Lexington, Massachusetts, 1992.

[DS83] R. Davies and R. G. Smith. Negotiation as a metaphor for distributed problem solving. *Artificial Intelligence*, 20(1):63–109, 1983.

[EB82] J. Engel and R. Blackwell. *Consumer behavior*. PCBS college publishing, 4th edition, 1982.

[EBA] `http://www.ebay.com/`.

[FF] `http://www.firefly.com/`.

[FJBS00] P. Faratin, N. R. Jennings, P. Buckle, and C. Sierra. Autonomated negotiation for provisioning virtual private networks using FIPA-complaint agents. In *Proc. 5th Int. Conf. on Practical Application of Intelligent Agents and Multi-Agent Systems (PAAM-2000)*, pages 185–202, Manchester, UK, 2000.

[FR93] D. Friedman and J. Rust, editors. *The double auction market: Institutions, theories and evidence*. Addison-Wesley, Reading, MA, 1993.

[FSJ98] P. Faratin, C. Sierra, and N. R. Jennings. Negotiation decision functions for autonomous agents. *Int. Journal of Robotics and Autonomous Systems*, 24(3-4):159–182, 1998.

[FSJ00] P. Faratin, C. Sierra, and N. R. Jennings. Using similarity criteria to make negotiation trade-offs. In *Proc. AAAI 4th Int. Conf. on Multi-Agent Systems (ICMAS2000)*, Boston, USA, 2000.

[FSJB99] P. Faratin, C. Sierra, N. R. Jennings, and P. Buckle. Designing responsive and deliberative automated negotiators. In *Proc. AAAI Workshop on Negotiation: Settling Conflicts and Identifying Opportunities*, pages 12–18, Orlando, FL, 1999.

[GMM98] R. Guttman, A. Moukas, and P. Maes. Agent-mediated electronic commerce: A survey. *Knowledge Engineering Review*, 13(2):147–159, 1998.

[Ham86] K. J. Hammond. *Case-based Planning: An Integrated Theory of Planning, Learning, and Memory*. PhD thesis, Yale University, 1986. YALEU/CSD/RR #488.

[HS69] J. Howard and J. Sheth. *The theory of buyer behavior*. Wiley, 1969.

[HTW] `http://www.hotwire.com/`.

[JFL+01] N. R. Jennings, P. Faratin, A. R. Lomuscio, S. Parsons, C. Sierra, and M. Wooldridge. Automated negotiation: prospects, methods and challenges. *Group Decision and Negotiation*, 10(2):199–215, 2001.

[JFN+00] N. R. Jennings, P. Faratin, T. J. Norman, P. O'Brien, and B. Odgers. Autonomous agents for business process management. *International Journal of Applied Artificial Intelligence*, 2(14):145–189, 2000.

[JPSF00] N. R. Jennings, S. Parsons, C. Sierra, and P. Faratin. Automated negotiation. In *Proceedings of the 5th International Conference on Practical Application of Intelligent Agents and Multi-Agent Systems (PAAM-2000)*, Manchester, UK, 2000.

[KL89] K. Kuwabara and V. R. Lesser. Extended protocol for multi-stage negotiation. In *Proceedings of the Ninth Workshop on Distributed Artificial Intelligence*, Rosario, Washington, September 1989.

[KM90] T. Kreifelts and F. von Martial. A negotiation framework for autonomous agents. In *Proceedings of the Second European Workshop on Modeling Autonomous Agents and Multi Agent Worlds*, Paris, France, August 1990.

[KR76] R. L. Keeney and H. Raiffa. *Decisions with multiple objectives*. John Wiley and sons, New York, 1976.

[Kra01] S. Kraus. *Strategic Negotiation in Multiagent Environments*. The MIT Press, 2001.

[KSE98] S. Kraus, K. Sycara, and A. Evenchik. Reaching agreements through argumentation: a logical model and implementation. *Artificial Intelligence*, 1-2(104):1–69, October 1998.

[KW82] D. Kreps and R. Wilson. Sequential equilibria. *Econometrica*, 50:863–894, 1982.

[KWZ95] S. Kraus, J. Wilkenfeld, and G. Zlotkin. Multiagent negotiation under time constraints. *Artificial Intelligence*, 2(75):297–345, October 1995.

[Lee96] L. Chi-Hang Lee. *Negotation Strategies and their Effect in a Model of Multi-Agent Negotiation*. PhD thesis, Department of Computer Science, University of Essex, July 1996.

[LWJ00] A. Lomuscio, M. J. Wooldridge, and N. R. Jennings. *Agent Mediated Electronic Commerce; A European Perspective*, chapter A classification scheme for negotiation in electronic commerce, pages 19–34. Springer Verlag, 2000.

[Nic66] F. Nicosia. *Consumer Decision processes: marketing and advertising implications*. Prentice Hall, 1966.

[OR94] M. J. Osborne and A. Rubinstein. *A Course in Game Theory*. MIT Press, Cambridge, Massachusetts, 1994.

[PCL] `http://www.priceline.com/`.

[PCR] `http://www.pricerunner.com/`.

[PCW] `http://www.pricewatch.com/`.

[Per] `http://www.personalogic.com/`.

[PJ96] S. D. Parsons and N. R. Jennings. Negotiation through argumentation-A preliminary report. In *Proc. Second Int. Conf. on Multi-Agent Systems*, pages 267–274, Kyoto, Japan, 1996.

[PSJ98] S. Parsons, C. Sierra, and N. Jennings. Agents that reason and negotiate by arguing. *Journal of Logic and Computation*, 8(3):261–292, June 1998.

[QXL] `http://www.qxl.com/`.

[RNSP97] J. Rodriguez, P. Noriega, C. Sierra, and J. Padget. FM96.5 A Java-based Electronic Auction House. In *Proceedings of 2nd Conference on Practical Applications of Intelligent Agents and MultiAgent Technology (PAAM)*, pages 207–224, London, UK, April 1997.

[RZ94] J. S. Rosenschein and G. Zlotkin. *Rules of Encounter: Designing Conventions for Automated Negotiation among Computers*. The MIT Press, Cambridge, MA, 1994.

[San99] T. Sandholm. Distributed rational decision making. In G. Weiss, editor, *Multiagent Systems: A Modern Introduction to Distributed Artificial Intelligence*, pages 201–258. MIT Press, 1999.

[SFJ97] C. Sierra, P. Faratin, and Nick R. Jennings. A service-oriented negotiation model between autonomous agents. In *Proc. 8th Int. Workshop on Modelling Autonomous Agents in a Multi-Agent World*, pages 17–35, Ronneby, Sweden, 1997.

[SJNP98] Carles Sierra, Nick R. Jennings, Pablo Noriega, and Simon Parsons. A framework for argumentation-based negotiation. In M. P. Singh, A. Rao, and M. J. Wooldridge, editors, *Intelligent Agents IV (LNAI Volume 1365)*, pages 177–192. Springer-Verlag, Berlin, 1998.

[SSPJ99] J. Sabater, C. Sierra, S. Parsons, and N. R. Jennings. Using multi-context systems to engineer executable agents. In *Proc. 6th Int. Workshop on Agent Theories Architectures and Languages (ATAL-99)*, pages 131–148, Orlando, FL, 1999.

[Syc87] Ekaterini P. Sycara. *Resolving Adversarial Conflicts: An Approach Integrating Case-Based and Analytic Methods*. PhD thesis, Georgia Institute of Technology, Atlanta, GA, June 1987.

[Syc88] K. Sycara. Resolving goal conflicts via negotiation. In Tom M. Smith, Reid G.; Mitchell, editor, *Proceedings of the 7th National Conference on Artificial Intelligence*, pages 245–250, St. Paul, MN, August 1988. Morgan Kaufmann.

[Syc89] K. Sycara. Argumentation: Planning other agents' plans. In *Proceedings of IJCAI-89*, pages 517–523, Detroit, Michigan, August 1989.

[Syc91] K. Sycara. Problem restructuring in negotiation. *Management Science*, 37(10), 1991.

[TAT] `http://ecommerce.media.mit.edu/Tete-a-Tete/`.

[Toh97] Fernando Tohme. Negotiation and defeasible reasons for choice. In Jon Doyle and Richmond H. Thomason, editors, *Working Papers of the AAAI Spring Symposium on Qualitative Preferences in Deliberation and Practical Reasoning*, pages 95–102, Menlo Park, California, 1997. American Association for Artificial Intelligence, American Association for Artificial Intelligence.

[Vic61] W. Vickrey. Counterspeculation, auctions, and competitive sealed tenders. *Journal of Finance*, 16:8–37, 1961.

[VJ00] N. Vulkan and N. R. Jennings. Efficient mechanisms for the supply of services in multi-agent environments. *Interanational Journal of Decision Support Systems*, 28(1-2):5–19, 2000.

[WBK96] M. Wooldridge, S. Bussmann, and M. Klosterberg. Production sequencing as negotiation. In *Proceedings of the First International Conference on Practical Applications of Intelligent Agents and Multi-Agent Technology (PAAM'96)*, pages 709–726, 1996.

[WJ95] M. Wooldridge and N. R. Jennings. Intelligent agents: theory and practice. *Knowledge Engineering Review*, 2(10):115–152, 1995.

[WWW98] P. R. Wurman, M. P. Wellman, and W. E. Walsh. The Michigan Internet AuctionBot: A configurable auction server for human and software agents. In K. P. Sycara and M. J. Wooldridge, editors, *Proceedings of the 2nd International Conference on Autonomous Agents (Agents'98)*, pages 301–308, New York, May 9–13, 1998. ACM Press.

[Zeu30] F. Zeuthen. *Problems of monopoly and economic warfare*. G. Routledge and Sons, 1930.

[ZR89] Gilad Zlotkin and Jeffrey S. Rosenschein. Negotiation and task sharing among autonomous agents in cooperative domains. In *Proceedings of the Eleventh International Joint Conference on Artificial Intelligence*, pages 912–917, August 1989.

Work	*Cardinality*	*Agent Characteristics*	*Goods*	*Events parameters*	*Info parameters*
[RZ94]	Negotiation issues: 1 Interactions: one-to-one	Agent roles: homogeneous Bidding strategy: Zeuthen Rationality: perfect Commitment: total Behaviour: self-interested Knowledge: complete	Value: public Nature: discrete	Bids validity: MCP Bidding: public Clearing: no Quotes-sched: N/A	Price-quotes: no trans hist: no
[WBK96]	Negotiation issues: 1 Interactions: many-to-many	Agent roles: homogeneous Bidding strategy: Zeuthen† Rationality: perfect Commitment: total Behaviour: self-interested Knowledge: complete	Value: Public Nature: discrete	Bids validity: MCP Bidding: public Clearing: no Quotes-sched: N/A	Price-quotes: no trans hist: no
[KWZ95]-1	Negotiation issues: 1 Interactions: one-to-one	Agent roles: homogeneous Bidding strategy: see desc Rationality: perfect Commitment: total Behaviour: self-interested Knowledge: complete	Value: public Nature: continuoum	Bids validity: any Bidding: public Clearing: no Quotes-sched: N/A	Price-quotes: no trans hist: no
[KWZ95]-2	Negotiation issues: 1 Interactions: one-to-one	Agent roles: homogeneous Bidding strategy: see desc Rationality: perfect Commitment: total Behaviour: self-interested Knowledge: complete	Value: public Nature: discrete	Bids validity: any Bidding: public Clearing: no Quotes-sched: N/A	Price-quotes: no trans hist: no
[KWZ95]-3	Negotiation issues: 1 Interactions: one-to-one	Agent roles: exclusive/waiting Bidding strategy: see desc Rationality: perfect Commitment: total Behaviour: self-interested Knowledge: complete	Value: public Nature: discrete	Bids validity: any Bidding: public Clearing: no Quotes-sched: N/A	Price-quotes: no trans hist: no
[KWZ95]-4	Negotiation issues: 1 Interactions: one-to-one	Agent roles: exclusive/waiting Bidding strategy: see desc Rationality: perfect Commitment: total Behaviour: self-interested Knowledge: incomplete	Value: public Nature: discrete	Bids validity: any Bidding: public Clearing: no Quotes-sched: N/A	Price-quotes: no trans hist: no
[KWZ95]-5	Negotiation issues: 1 Interactions: Many-to-many	Agent roles: homogeneous Bidding strategy: see desc Rationality: perfect Commitment: total Behaviour: self-interested Knowledge: complete	Value: public Nature: discrete	Bids validity: any Bidding: public Clearing: no Quotes-sched: N/A	Price-quotes: no trans hist: no

Table 1: A table summarising some of the works analysed in Section 3. See also Table 2.

† This is effectively an extension of the Zeuthen strategy as that was proposed for the one-to-one case.

Work	*Cardinality*	*Agent Characteristics*	*Goods*	*Events parameters*	*Info parameters*
[FSJ98]	Negotiation issues: many Interactions: one-to-one	Agent roles: heterogeneous Bidding strategy: see desc Rationality: bounded Commitment: partial Behaviour: self-interested Knowledge: incomplete	Value: private Nature: discrete	Bids validity: see desc Bidding: public Clearing: no Quotes-sched: N/A	Price-quotes: no trans hist: no
[PSJ98]	Negotiation issues: 1 Interactions: one-to-one	Agent roles: heterogenous Bidding strategy: see desc Rationality: perfect Commitment: total Behaviour: variable Knowledge: incomplete	Value: private Nature: discrete	Bids validity: see desc Bidding: public Clearing: no Quotes-sched: N/A	Price-quotes: no trans hist: no
[Syc91]	Negotiation issues: multi Interactions: mediated	Agent roles: heterogenous Bidding strategy: see desc Rationality: bounded Commitment: total Behaviour: variable Knowledge: incomplete	Value: private Nature: continous	Bids validity: any Bidding: private Clearing: yes Quotes-sched: yes	Price-quotes: yes trans hist: yes

Table 2: A table summarising some of the works analysed in Section 3. See also Table 1.