Monotonic Concession Protocols for Multilateral Negotiation

Ulle Endriss
Institute for Logic, Language and Computation
University of Amsterdam, The Netherlands
ulle@illc.uva.nl

ABSTRACT

The most natural way of thinking about negotiation is probably a situation whereby each of the parties involved initially make a proposal that is particularly beneficial to themselves and then incrementally revise their earlier proposals in order to come to an agreement. This idea has been formalised in the so-called monotonic concession protocol, a set of rules defining the range of acceptable moves during a negotiation process intended to follow this general scheme. In the case of negotiation between just two agents, the monotonic concession protocol has become a textbook example and its formal properties are well-understood. In the case of multilateral negotiation, where more than two agents need to come to an agreement, on the other hand, it is not at all clear how to set up a monotonic concession protocol. As it turns out, the design of such a protocol boils down to the question of what constitutes a multilateral concession. In this paper, we make several proposals as to what might be an appropriate definition and analyse the properties of the proposed concession criteria.

Categories and Subject Descriptors

I.2.11 [**Artificial Intelligence**]: Distributed Artificial Intelligence—*Multiagent systems*; J.4 [**Social and Behavioral Sciences**]: Economics; K.4.4 [**Computers and Society**]: Electronic Commerce

General Terms

Economics, Theory

Keywords

Negotiation, Protocols, Game Theory

1. INTRODUCTION

Negotiation between self-interested agents is a blossoming research area at the heart of multiagent systems, incorporating many ideas and models originally developed in different branches of economics (notably game theory), and at the same time benefiting significantly from the adoption of computational methods [6, 7, 10]. Much work on negotiation is either concerned with *bilateral* ("one-to-one") negotiation [10] or with *auctions* [2]. The latter are highly structured forms of negotiation between several agents.[1] Truly *multilateral* negotiation, on the contrary, where groups of more than just two agents can freely come together and agree on a deal between them has received far less attention. This is very much due to the significant technical challenges posed by such rich negotiation models (and certainly not because they would not be considered important).

For instance, in the context of distributed task and resource allocation, it is well-known that protocols only permitting two agents at a time to negotiate agreements are usually not expressive enough to reach optimal allocations [3, 11]. A typical example would be a scenario with three agents and three resources. Think of the agents as being arranged in a circle. Suppose the agents all differ in their assessment of the respective values of the three items and also suppose that, initially, each agent holds their second-most preferred item, their lefthand neighbour holds their favourite item, and their righthand neighbour holds their least-preferred item. Now any potential deal between two agents would either involve one of them losing all resources, or having to exchange their second-best item for the worst one. Neither would be considered beneficial by that agent, *i.e.* a *rational* agent would have to refuse any such deal.[2] Hence, using bilateral deals alone, it is not possible for rational agents to move away from the initial allocation and reach the ideal situation where each agent gets their preferred resource. Such limitations can be overcome by negotiation mechanisms that allow for more than just two agents to make a deal. In our example, for instance, it is clear that there is a multilateral deal between all three agents that would be considered rational by all of them.

Other examples for the need for multilateral negotiation abound (international politics being just one of them). Clearly, there are many situations where the options to agree on deals between more than just two parties at a time can greatly enrich the opportunities available to the negotiators

[1] But note that the *outcome* of an auction is usually also just a collection of separate bilateral deals between the auctioneer and different bidders.

[2] Even with monetary side payments, it is not difficult to set up an example where any bilateral deal would be considered irrational by at least one of the agents concerned [3, 11].

AAMAS'06 May 8–12 2006, Hakodate, Hokkaido, Japan.

and can lead to much better negotiation outcomes. On the downside, of course, such multilateral negotiation models are a lot more complex than bilateral negotiation (and typically also more complex than the highly structured negotiation schemes between several agents given by auctions). Acknowledging these difficulties, this paper aims at clarifying some of the more fundamental issues at stake.

Specifically, in this paper we are interested in *monotonic concession protocols* for multilateral negotiation. A negotiation *protocol* is a set of rules specifying the range of "legal" moves available to each agent at any stage of a negotiation process (in contrast to this, an agent's negotiation *strategy* is a function from the negotiation history to a follow-up move respecting the protocol). Monotonic concession protocols, in particular, formalise what might be considered the most natural way of thinking about negotiations: Initially, each of the parties involved make a proposal that is particularly beneficial to themselves, and then they incrementally revise their earlier proposals in order to eventually come to an agreement. In the case of negotiation between two agents, the monotonic concession protocol has become a textbook example for negotiation mechanisms [10, 13] and its formal properties are well-understood. In the multilateral case, as we shall see, there appear to be many different ways of generalising the (single) definition for the two-agent case. Crucially, the design of multilateral monotonic concession protocols boils down to the question of what constitutes a concession to a group of agents, rather than to a single opponent. This paper identifies (at least some of) the available options, and analyses their formal properties. We also include a preliminary discussion of the question of how to design suitable negotiation strategies for these protocols.

Paper Overview

Section 2 reviews the monotonic concession protocol for the *two-agent case* and an associated negotiation strategy, as developed in the work of Zeuthen [15], Harsanyi [4], and Rosenschein and Zlotkin [10]. Following this, Section 3 sets out the general framework of generalising the protocol to the multilateral case. Such a generalisation requires us to decide what constitutes an *agreement* between several agents and, crucially, how to characterise a *concession* to more than one opponent. These concepts are discussed in Section 4. In particular, Section 4 presents seven alternative definitions of the term *multilateral concession*. Each of these gives rise to a different variant of the monotonic concession protocol. The *properties* of these variants of the basic protocol, in particular *termination* and *deadlock-freedom*, are then analysed in Section 5. Some of the issues pertaining to the definition of a suitable *strategy* in a multilateral negotiation context are discussed in Section 6. Section 7 concludes.

2. THE TWO-AGENT CASE

In this section, we are going to describe the monotonic concession protocol for *two* agents, as presented by Rosenschein and Zlotkin [10]. We are also going to discuss a particular negotiation strategy, originally introduced by Zeuthen [15] and shown to be based on sound game-theoretical foundations by Harsanyi [4], that agents may adopt during an interaction regulated by the protocol.

Notation and Assumptions

Let $\mathcal{A} = \{1, 2\}$ be a set of two agents and let $\mathcal{X}$ be a finite set of *potential agreements* (which we are also going to refer to as *deals* or *proposals*). Each of the agents $i \in \mathcal{A}$ is equipped with a *utility function* $u_i : \mathcal{X} \to \mathbb{R}_0^+$ mapping agreements to non-negative reals. That is, we exclude any potential agreements that would yield negative utility for either one of the agents *a priori*. $\mathcal{X}$ includes a specific agreement, called the *conflict deal*, that yields utility 0 for both agents. This will be the agreement chosen in case negotiation breaks down and is assumed to be the worst possible outcome.

Monotonic Concession Protocol

The monotonic concession protocol proceeds in *rounds*. In each round, both agents make simultaneous *proposals*, *i.e.* they each propose one of the potential agreements from $\mathcal{X}$. In the first round, each agent is free to make any proposal. In any subsequent round, each agent $i \in \mathcal{A}$ has got two options (let $x_i \in \mathcal{X}$ be the most recent proposal made by i):

(1) Make a *concession* and propose a new deal x_i' that is preferable to the other agent j: $u_j(x_i) < u_j(x_i')$.

(2) Refuse to make a concession and stick to proposal x_i.

An *agreement* is found when one agent makes a proposal that its opponent rates at least as high as its own current proposal, *i.e.* when $u_1(x_1) \leq u_1(x_2)$ or $u_2(x_2) \leq u_2(x_1)$ or both. *Conflict*, on the the hand, arises when both agents refuse to make a concession during a given round. The protocol runs until either an agreement is found or conflict occurs. In case of agreement, the proposal satisfying both agents is chosen (in case both proposals are satisfiable, we flip a coin to pick one). In case of conflict, the conflict deal is the outcome of the negotiation.

The monotonic concession protocol is known to always *terminate* [10], *i.e.* there can be no infinite sequence of rounds.

Zeuthen Strategy

What would be a suitable strategy for agents to adopt when following the monotonic concession protocol? Zeuthen [15] proposed that agents should evaluate their respective *willingness to risk conflict*, and that the agent with the lower value for this measure should make the next concession (and both agents should concede in case their risk values are equal). Furthermore, a concession should be *sufficient* to ensure that the other agent will have to concede in the following round.

Zeuthen proposed to calculate an agent's willingness to risk conflict as the quotient of the loss in utility incurred in case the other agent's proposal is being adopted and the loss incurred by causing conflict (both with respect to the utility of the agent's own current proposal). As the utility of the conflict deal is 0, the willingness of agent i to risk conflict (with j being the other agent) is given by the following formula (we assume $u_i(x_i) \neq 0$ for ease of presentation):

$$Z_i = \frac{u_i(x_i) - u_i(x_j)}{u_i(x_i)}$$

Zeuthen's intuitive justification for this strategy has been complemented by Harsanyi's work [4], who showed how the same strategy can be derived from a small number of *fundamental postulates*. The most important of these is the postulate of *utility maximisation.* Harsanyi argues as follows: Suppose the two agents are playing a single-round game; each agent can either accept or reject the other's proposal. Let p_2 be the probability that agent 2 will reject x_1. Then the expected payoff for agent 1 of rejecting x_2 is $(1 - p_2) \cdot u_1(x_1)$. On the other hand, Harsanyi argues, the *certain* payoff associated with accepting x_2 is $u_1(x_2)$. Hence, agent 1 should accept iff $u_1(x_2) \geq (1 - p_2) \cdot u_1(x_1)$. This is equivalent to $\frac{u_1(x_1)-u_1(x_2)}{u_1(x_1)} \leq p_2$. Together with a *symmetry* and a *monotonicity* postulate, this leads to the Zeuthen criterion: agent 1 should concede whenever $Z_1 \leq Z_2$.[3]

The Zeuthen strategy is *efficient* in the sense of leading to a final agreement maximising the product of utilities [4, 10].[4] This follows from the fact that $Z_1 \leq Z_2$ is equivalent to $u_1(x_1) \cdot u_2(x_1) \leq u_1(x_2) \cdot u_2(x_2)$, *i.e.* in each round the agent whose current proposal yields the lower product of utilities makes a concession.

The strategy is also "almost" *stable* (*i.e.* in Nash equilibrium [9]). The only critical case is when both agents have the same willingness to risk conflict after the penultimate round. In that case, an agent could benefit from deviating and sticking to their latest proposal. This problem can be overcome by extending the Zeuthen strategy with a mixed equilibrium strategy for the two-person game induced by this "last step situation", albeit at the cost of giving up on efficiency [9, 10].

3. NEGOTIATION PROTOCOLS

In this section, we are going to discuss the design of negotiation protocols that can be used by more than just two agents.

Notation and Assumptions

Again, let $\mathcal{X}$ be a finite set of potential agreements and let $\mathcal{A}$ be a finite set of agents. Each agent $i \in \mathcal{A}$ is equipped with a utility function $u_i : \mathcal{X} \to \mathbb{R}_0^+$, *i.e.* we only admit non-negative utilities, and $\mathcal{X}$ includes the so-called conflict deal to which every agent assigns utility 0. In one case we are going to comment on the effects of only allowing strictly positive utilities (for anything but the conflict deal).

A Multilateral One-Step Protocol

We should stress that, in this paper, we are not interested in finding the "best" protocol for multilateral negotiation. Instead, we are only interested in generalising the ideas present in the two-agent monotonic concession protocol. Also in the two-agent case, it is well-known that the monotonic concession protocol is not the "best" protocol. Indeed, there is a simple alternative, termed the *one-step protocol* by Rosenschein and Zlotkin [10], that is much simpler than the monotonic concession protocol and permits a negotiation strategy that is both stable and efficient. Recall that the Zeuthen strategy cannot guarantee both stability *and* efficiency (depending on how the "last step situation" is handled, either stability or efficiency need to be sacrificed). The idea of the one-step protocol is to simply ask both agents for a proposal and to adopt the proposal that yields the higher product of utilities (and to flip a coin in case the products are the same). The best strategy that agents can follow in this protocol is to propose the agreement that is best for themselves amongst those with maximal product of utilities. This is clearly both stable and efficient.

The one-step protocol can easily be extended to the multilateral case: Collect a proposal from each agent and randomly choose an outcome amongst those proposals that maximise the product of utilities. The obvious strategy, which is the same as for the two-agent case, is also both stable and efficient. This simplistic solution not withstanding, we still remain interested in the monotonic concession protocol, because —unlike the one-step protocol— it is a direct formalisation of natural negotiation behaviour.

Multilateral Monotonic Concession Protocols

How can we generalise the monotonic concession protocol to cases where more than two agents need to come to an agreement? The general structure of the negotiation protocol will stay the same:

(1) In the first round, each agent makes an initial proposal.

(2) In each subsequent round, each agent either makes a *concession* or sticks with their current proposal.

(3) The above is being iterated until either a conflict situation arises (no agent makes a concession) or *agreement* is reached.

The two critical terms here, which we have as yet left undefined, are *concession* and *agreement.*

4. AGREEMENT AND CONCESSION

In this section, we are going to discuss the meaning of the terms *agreement* and *concession* in the context of multilateral negotiation.

Multilateral Agreements

In the case of agreement, the generalisation from the two-agent case to the general multilateral case is straightforward:

> *Agreement is reached iff one agent makes a proposal that is at least as good for each other agent as their own current proposal.*

To put it more formally, agreement is reached iff there exists an agent $i \in \mathcal{A}$ such that $u_j(x_i) \geq u_j(x_j)$ for all agents $j \in \mathcal{A}$.

Options for Defining Multilateral Concessions

In the case of concession, however, it is not at all clear what would be the "correct" way of generalising to the multilateral case. Indeed, what does it mean to make a concession to a group of opponents? Do we have to make our new proposal more palatable to everyone else, to just one of them, or do we have to propose something that would somehow increase the utility of the group as a whole (and what would that actually mean)?

---

[3]What may be considered questionable in Harsanyi's argument is that he implicitly assumes that x_2 will be selected as soon as agent 1 accepts it. This disregards the possibility that both agents accept each others proposal, in which case the choice should be random. However, interestingly, this more complex assumption leads to the same solution. In this case, the expected payoff for accepting x_2 is $p_2 \cdot u_1(x_2) + (1-p_2) \cdot \frac{1}{2}(u_1(x_1) + u_1(x_2))$. Using this in place of $u_1(x_2)$ eventually leads to the same inequation $Z_1 \leq Z_2$

[4]This also implies that the final agreement will be *Pareto optimal* (no alternative agreement would be better for one agent without decreasing utility for the other [8]).

All of the following interpretations of the term *concession* seem to have some merit:

(1) *Strong concession:* Make a proposal that is strictly better for each of the other agents.

(2) *Weak concession:* Make a proposal that is strictly better for at least one of the other agents.

(3) *Pareto concession:* Make a proposal that is at least as good for all of the other agents and strictly better for at least one of them.

(4) *Utilitarian concession:* Make a proposal such that the sum of utilities of the other agents increases (utilitarian social welfare).

(5) *Egalitarian concession:* Make a proposal such that the minimum utility amongst the other agents increases (egalitarian social welfare).

(6) *Nash concession:* Make a proposal such that the product of utilities of the other agents increases (Nash product).

(7) *Egocentric concession:* Make a proposal that is worse for yourself.[5]

We are first going to clarify the choice of name for some of these concession criteria and then turn to the analysis of their properties.

Relation to Social Welfare Concepts

Recall that, in welfare economics, a *Pareto improvement* is a move to a different agreement that harms nobody but benefits at least one member of society (and a Pareto optimal agreement is one that does not admit any further Pareto improvements). Analogously, a Pareto *concession* by agent i is a concession that results in a Pareto improvement for the group of agents $\mathcal{A} \setminus \{i\}$. The names for criteria (4)–(6) are also inspired by concepts from welfare economics. These are all concepts that are used to assess the well-being of a group (or *society*) of agents in terms of the individual utilities of its members [1, 8].

The *utilitarian social welfare* enjoyed by a group, for instance, is defined as the sum of the utilities of its members. This is a useful metric for assessing the well-being of a group if we adopt a classical utilitarian point of view (note that this is also a metric for *average* utility, *i.e.* an agent making a utilitarian concession would improve the average utility of its opponents). The *egalitarian social welfare*, on the other hand, is defined as the utility of the agent with the currently lowest utility level (*i.e.* in an egalitarian system we would measure the success of society in terms of the well-being of its poorest members). The *Nash product* is an interesting compromise; it measures social welfare in terms of the product of individual utilities. This means that both overall increases in utility and inequality-reducing utility transfers would be considered beneficial to society. Other notions of social welfare studied in the literature would give rise to further concession criteria.

[5]Here *egocentric* is not to be misread as *selfish*; it simply means that this notion of concession is based on the agent's own valuations rather than that of other agents. This makes sense in domains where a reduction in utility for one agent can generally be assumed to result in an increase in utility for (some of) the other agents.

Relation to the Two-Agent Case

If there are just two agents, then making a concession of either one of the first six types simply amounts to making a new proposal that is strictly better for your (single) opponent. Hence, criteria (1)–(6) are all faithful extensions of the two-agent case.

In fact, criterion (7) is also a faithful extension of the two-agent case, at least if we assume that an agent would never make a proposal that is dominated by another potential proposal. This is so, because moving to a new proposal that is worse for yourself but not better for your opponent than your current proposal amounts to moving to a proposal that is dominated by your current proposal.

5. PROTOCOL PROPERTIES

In this section, we are going to analyse various interesting properties for the different monotonic concession protocols entailed by our seven multilateral concession criteria.

Termination

An important property of a negotiation protocol is whether it can guarantee that any negotiation process following the protocol will eventually terminate. Each of the multilateral concession criteria we have proposed gives rise to a different negotiation protocol. The following proposition summarises some simple termination results for these protocols:

PROPOSITION 1. *Except for weak concessions, each of the seven proposed concession criteria guarantees termination of the negotiation protocol.*

PROOF. Each of the concession criteria, except that of weak concession, is defined in terms of a metric that induces a strict partial (in some cases even total) order $\prec_i$ over the space of proposals from the viewpoint of each individual agent i. For instance, in the case of a utilitarian concession by agent i, the sum of utilities for the remaining agents has to increase, *i.e.* you may define $x \prec_i y$ iff $\sum_{k\in\mathcal{A}\setminus\{i\}} u_k(x) < \sum_{k\in\mathcal{A}\setminus\{i\}} u_k(y)$. Recall that the overall number of possible proposals is assumed to be finite. Hence, each agent i can only make a finite number of concessions, as each concession will take its own current proposal higher up in the $\prec_i$-order. Hence, as the overall number of agents is assumed to be finite as well, there can only be a finite number of concessions overall, *i.e.* negotiation is bound to terminate and either conflict or an agreement will be reached eventually. □

Note that any such termination result can only ever apply to the *protocol* as such, not to negotiation processes in general. If one of the agents does not make a move during one of the rounds (because, for instance, its reasoning process for deciding the next move does not terminate), then the negotiation gets stuck. In practice, this problem can be tackled by introducing deadlines ("if you do not make a move within 1 minute, then this counts as a refusal to concede").

Clearly, in the case of weak concessions, we cannot guarantee termination (unless there are only two agents). Indeed, an agent could alternate between two proposals indefinitely and make a weak concession each time. This is possible when there are two candidate proposals, such that one of them is strictly preferred by the agent's first opponent and the other is strictly preferred by its second opponent.

Compositionality

Closely related to the question of termination is what one might want to term *compositionality*. This property holds iff the composition of two consecutive concessions each meeting a given concession criterion will always meet that same criterion as well. That is, whenever the move from x to x' is considered a valid concession, and the move from x' to x'' is as well, then the direct move from x to x'' will constitute a valid concession as well.

Clearly, all of our proposed concession criteria satisfy the compositionality property, with the sole exception of weak concessions. In the case of a Nash concession, for instance, if two consecutive concessions each increase the product of utilities for the other agents, then the product of utilities will have increased as well with respect to the entire negotiation process. For weak concessions, on the other hand, the same counterexample given for termination applies: the second concession may take us back to the initial proposal.

Deadlock-freedom

Another important property of negotiation protocols to consider is deadlock-freedom. A *deadlock* is a situation where no agent can make any further move that would be compliant with the negotiation protocol, but none of the terminal states (either agreement of conflict) of the protocol has been reached either. Of course, in the context of a monotonic concession protocol, each agent always has the move of not conceding available to them. That is, for any of our concession criteria, it will always be possible for agents to continue negotiation in a fashion deemed legal by the protocol. However, there may be situations where an agent would not be able to make a concession, even if they would be prepared to give up their current position. Naturally, we would prefer to avoid such situations and use a concession criterion satisfying the following property of deadlock-freedom:

> *We call a concession criterion deadlock-free iff it guarantees that at least one agent can make a concession satisfying the criterion at any stage during negotiation, until an agreement has been reached.*

It is not difficult to see that any of our criteria must be deadlock-free in case there are only two agents in the system:

PROPOSITION 2. *In the two-agent case, each of our seven concession criteria is deadlock-free.*

PROOF. Let x_1 and x_2 be the current proposals of agents 1 and 2, respectively. If no agreement has been reached yet, then we must have both $u_1(x_1) > u_1(x_2)$ and $u_2(x_2) > u_2(x_1)$. As argued earlier, in the two-agent case criteria (1)–(6) coincide; they are fulfilled if one of the agents makes a new proposal that is better for its opponent than its current offer. Hence, x_2 would be a legal concession for agent 1 for any of the criteria (1)–(6). Furthermore, it would also be a legal concession according to criterion (7). Hence, there can never be a deadlock in the two-agent case. □

In the general case, with potentially more than just two agents, however, the situation is different. Now just some of our criteria are deadlock-free, while others are not.

Furthermore, for one of them (the Nash concession), deadlock-freedom depends on a seemingly minor detail of the negotiation scenario. Recall that we require all utility functions to be non-negative. We say that utilities are *required to be positive* in case we do require all utilities, except for the conflict deal, to be greater than 0. We are now able to state our results regarding deadlock-freedom:

PROPOSITION 3. *The weak, the utilitarian, and the egocentric concession criteria are all deadlock-free. The Pareto, the strong, and the egalitarian concession criteria are not deadlock-free. The Nash concession criterion is deadlock-free iff utilities are required to be positive.*

PROOF. We begin with the negative results, more specifically with the case of *Pareto concessions.* Suppose there are three agents with the following current proposals (which are also the only possible outcomes) and associated utilities:

	x_1	x_2	x_3
u_1	4	2	2
u_2	3	2	2
u_3	1	2	2

It is easy to see that agreement has not been reached and that no agent is in a position to make a Pareto concession. Hence, this example constitutes a deadlock. Every *strong concession* is also a Pareto concession, *i.e.* strong concessions cannot be deadlock-free either.

To construct an example that proves that *egalitarian concessions* are not deadlock-free, let ϵ be the minimal utility assigned to any outcome by any agent. Then in any game where for each possible outcome x there are at least two agents assigning utility ϵ to x, no agent can ever make an egalitarian concession (because the minimum utility of the other agents is the same for all possible proposals).

Now let us move on to the positive results. The *weak concession* criterion is clearly deadlock-free. In fact, at any stage *each* agent i could make a weak concession (until an agreement is reached). If no agreement has been reached, then there must be at least one other agent j such that $u_j(x_j) > u_j(x_i)$. But then proposing x_j would constitute a weak concession for i. An analogous argument can be used for the case of *egocentric concessions.*

Next consider the case of *utilitarian concessions.* Let i be (one of) the agent(s) making a proposal x_i yielding a maximal sum of utilities (amongst the current proposals). If there is no agreement yet, then there must be another agent j who likes x_i less than their own proposal: $u_j(x_i) < u_j(x_j)$. But we also have $\sum_{k\in\mathcal{A}} u_k(x_i) \geq \sum_{k\in\mathcal{A}} u_k(x_j)$. Hence, $\sum_{k\in\mathcal{A}\setminus\{j\}} u_k(x_i) > \sum_{k\in\mathcal{A}\setminus\{j\}} u_k(x_j)$. That is, x_i constitutes a legal utilitarian concession for agent j and there is no deadlock.

An analogous argument can be used in the case of *Nash concessions.* However, in this case the validity of the last step in the argument relies on $u_j(x_i)$ being greater than 0. That is, this proof will only go through in case utilities can be assumed to be positive. If zero utilities are admissible, then we can use a similar argument as the one used in the case of egalitarian concessions to show that Nash concessions are *not* deadlock-free. Take a game where for each possible outcome x there are at least two agents assigning utility 0 to x. Then the product of utilities of the "other" agents will always be 0, *i.e.* no agent can ever make a Nash concession. □

Lack of deadlock-freedom is a serious deficiency of a negotiation protocol. Arguably, deadlock-freedom is even more important than termination, because a lack of termination can be addressed by refining the concession criterion in question, while deadlock-freedom could only be recovered by loosening the definition of a criterion.

Wooldridge *et al.* [14], for instance, who propose a multilateral negotiation mechanism that draws on ideas from the monotonic concession protocol for the two-agent case, require concessions to satisfy both the utilitarian and the Pareto concession criteria. They recognise that their approach is not deadlock-free (without using that expression), and suggest to overcome this problem by introducing the option of *backtracking*, *i.e.* of proposing a deal that is neither a concession nor the same deal proposed in the previous round. Arguably, such an extension of the protocol takes us outside of the realm of monotonic concession protocols.

Verifiability and Privacy

Next we turn our attention to two inter-related issues: *verifiability* and *privacy*. A negotiation protocol should be verifiable in the sense that it should be possible to check whether everybody is really following the rules laid down by the protocol. But who should check?

The first type of verifiability is verifiability by an *outside observer*. Clearly, an outside observer could verify some of the more operational issues, e.g. whether agents really make simultaneous proposals. However, an outside observer usually cannot check whether every agent who is changing their proposal is really making a true concession. This would require an accurate knowledge of the utility functions of the agents concerned and, for reasons of privacy, these agents may not be prepared to publish their utility functions (or if they do, they may not always be entirely truthful about them). Of course, if we assume the presence of an *omniscient observer*, then full verifiability is clearly given for any of the variants of the protocol considered here, but it is also not very interesting.

The issue of privacy does not just affect verifiability, but is a somewhat critical issue about the whole setup of the monotonic concession protocol. Nevertheless, the two-agent monotonic concession protocol has the very desirable property that each individual agent can verify whether their opponent is conforming to the rules of the protocol. They only need to know their *own* utility function. Whether or not they report their utility function truthfully is up to them, but they can always check whether a new proposal constitutes a concession with respect to their true utility function (and, of course, also with respect to their reported utility function). In the case of more than two agents, this is not the case anymore, simply because concession criteria are defined in terms of more than one utility function (except for the egocentric criterion).

In the multilateral case, it seems interesting to introduce a concept of *distributed verifiability*. Is it possible that all the other agents together verify whether or not the latest move by a particular agent constitutes a true concession? Of course, the other agents may also be lying to each other, but we can at least check whether a concession criterion admits distributed verifiability in the following sense:

> *A concession criterion admits distributed verification iff for any move that is not a true concession, at least one of the other agents can detect that violation.*

It turns out that only the strong concession criterion is verifiable in a distributed manner, although the weak and the Pareto concession criteria also admit a slightly weaker form of distributed verification:

(1) The *egocentric* concession criterion is not verifiable, because only the proposing agent could check.

(2) The *utilitarian*, *egalitarian*, and *Nash* concession criteria are not verifiable in the distributed sense, because no individual agent can check whether the sum/minimum/product has increased.

(3) The *strong* concession criterion is verifiable in a distributed manner: for any violation, at least one of the agents will be worse off.

(4) The *weak* and the *Pareto* concession criteria are verifiable in a distributed manner if agents can communicate (if an agent has received a strict improvement they should announce that fact). If no agent reports a strict improvement, then a violation has occurred (in addition, for a Pareto concession, each agent would have to check that they did not suffer a loss in utility).

This concludes our discussion of the properties of the seven variants of the monotonic concession protocol.

6. NEGOTIATION STRATEGIES

Our main interest in this paper has been concerned with monotonic concession *protocols* for multilateral negotiation (and thereby also the nature of multilateral concessions), but this investigation also clearly raises the question what *strategy* an individual agent participating in a negotiation regulated by such a protocol should follow. A proper game-theoretical analysis of this question would be of considerable interest, but also technically demanding. Such an analysis lies outside the scope of this paper.

Still, in this section, we want to briefly review a number of seemingly obvious choices.

Choicepoints

It seems reasonable to assume that agents always start with a proposal that is optimal for themselves. Of course, from a "psychological" point of view it may be considered advantageous to start with a more realistic proposal straight away, but we are not going to be concerned with such questions here. We are also going to ignore the question as to *which* of the optimal proposals to select in case there are several.

In every subsequent round, each agent has to answer two questions:

(1) Should I make a concession?

(2) If yes, how much should I concede?

Of course, just as it is unclear how to characterise the concept of a concession in a multilateral negotiation context, it may also seem unclear how to quantify the extent of such a concession. Nevertheless, there is in fact a rather simple answer to our second question: A concession should always be *minimal* with respect to the utility loss incurred by the agent making the concession. That is, amongst the pool of potential proposals that would satisfy the concession criterion enforced by the protocol, the agent should choose the proposal that gives the highest utility to themselves (again, in this tentative discussion, we are going to ignore the question of what to do in case there are several potential concessions yielding maximal utility).

The concept of minimal concessions is useful when discussing these questions at an abstract level. In practice, it can and should be further refined. For the Zeuthen strategy as formulated by Rosenschein and Zlotkin [10], for instance, agents are supposed to make *sufficient* rather than *minimal* concessions (sufficient in the sense of forcing the opponent to concede in the next round). Of course, for most concession criteria, a sufficient concession can always be decomposed into a series of minimal concessions; so there is no important qualitative difference at the abstract level.[6]

To summarise, the question of strategy reduces to the question of *when* to make a concession. In the two-agent case, the Zeuthen strategy settles this question by letting the agents compute the *willingness to risk conflict* for both agents and asking the agent for whom this value is lower to make the next concession (with suitable modifications in case the two values are equal). In the sequel, we are going to consider three possible generalisations of the Zeuthen strategy to the multilateral case.

Willingness to Risk Conflict

In the two-agent case, the Zeuthen strategy is built around the concept of an agent's willingness to risk conflict. This value is computed as the quotient of the utility loss incurred by accepting the opponent's proposal and the utility loss incurred by settling with the conflict deal (both with respect to the utility of the agent's own current proposal). So for agent 1, we get the following value: $\frac{u_1(x_1)-u_1(x_2)}{u_1(x_1)}$.

A natural generalisation of this criterion to the multilateral case would be to evaluate the loss in utility in case of concession assuming the worst possible outcome for the agent. That is, for agent i, we obtain the following formula for its willingness to risk conflict:

$$Z_i \quad = \quad \begin{cases} 1 & \text{if } u_i(x_i) = 0 \\ \dfrac{u_i(x_i) - \min\{u_i(x_k) \mid k \in \mathcal{A}\}}{u_i(x_i)} & \text{otherwise} \end{cases}$$

Then the agent with the lowest willingness to risk conflict should make the next concession. This criterion has also been proposed by Wooldridge *et al.* [14].

While this strategy seems to be a natural generalisation of the two-agent Zeuthen strategy, it can unfortunately lead to a deadlock. Consider the following example. There are three agents and three possible outcomes. Suppose agent 1 is currently proposing x_1, agent 2 is proposing x_2, and agent 3 is proposing x_3:

	x_1	x_2	x_3
u_1	10	8	8
u_2	9	10	5
u_3	9	5	10

$\leadsto$ $Z_1 = 0.2$, $Z_2 = 0.5$, $Z_3 = 0.5$

Note that agreement has not yet been reached. According to our strategy, agent 1 should make a concession. However, agent 1 is not *able* to make a concession that would either affect the Z-value of any of its opponents or increase its own.

[6] An exception to this rule are the weak concessions, which —as we have seen earlier— are not decomposable. That is, here choosing a certain minimal concession may block the way to a suitable sufficient concession later on.

This observation suggests that simply checking for which agent the Z-value is minimal is not a sufficiently sophisticated criterion for deciding who should make the next concession. A potential way around this problem may be to select the agent with the minimal Z-value amongst those that are actually able to make an *effective* concession.

A Product-increasing Strategy

As discussed in Section 2, Harsanyi [4] showed that the Zeuthen strategy corresponds to a product-increasing strategy. That is, the agent whose proposal yields the lower product of utilities should make the next concession. This suggests another obvious generalisation of the strategy, namely that also in the multilateral case any agent whose proposal yields a minimal product of utilities should make a concession and propose a new deal with a higher product.

$$Z'_i \quad = \quad \prod_{k \in \mathcal{A}} u_k(x_i)$$

One may be inclined to believe that this strategy is bound to result in an agreement with maximal product of utilities, *i.e.* that the product-increasing strategy should be efficient. However, it is possible to construct scenarios where this is not so. The problem is that also a proposal with non-maximal product may be found to be acceptable to everyone. Take the following example (suppose agent 1 proposes x_1, agent 2 proposes x_2, and agent 3 proposes x_3):

	x_1	x_2	x_3
u_1	5	4	4
u_2	3	2	10
u_3	3	10	2

$\leadsto$

$$u_1(x_1) \cdot u_2(x_1) \cdot u_3(x_1) = 45$$
$$u_1(x_2) \cdot u_2(x_2) \cdot u_3(x_2) = 80$$
$$u_1(x_3) \cdot u_2(x_3) \cdot u_3(x_3) = 80$$

Here x_1 is the proposal with the lowest product of utilities, but it is also the only proposal that would be acceptable to all agents (as it gives them at least as much utility as their own current proposals). Whether or not such a situation could really occur in an actual negotiation remains to be investigated (note that the example is somewhat unrealistic, because agents 2 and 3 did not propose their favourite agreements first).

Sum of Products of Pairs

Yet another potential generalisation of the Zeuthen strategy, also building on its interpretation as a product-increasing strategy, would be to compute the sum of products of pairs of utilities for every proposal and to ask the agent for whose proposal that value is minimal to make a concession.

$$Z''_i \quad = \quad \sum_{j \neq k \in \mathcal{A}} u_j(x_i) \cdot u_k(x_i)$$

This appears to be another promising approach, although further investigation is required before any kind of conclusive judgement would be possible.

Stability

We conclude our preliminary discussion of issues surrounding the definition of suitable strategies for multilateral negotiation with a general observation regarding the stability of such mechanisms. Recall that a set of strategies is said to be stable (or in Nash equilibrium) iff no individual agent would have an incentive to deviate unilaterally from their assigned strategy [9].

Any pure strategy where conflict occurs whenever a single agent deviates is stable. This follows from the fact that conflict is assumed to be the worst possible outcome for any of the participating agents. Hence, as in the two-agent case, the only critical situation is when several agents should concede in the same round (and in fact this only makes a difference when this situation occurs immediately before an agreement is found). Whatever the exact criterion may be, to guarantee stability it would have to be enhanced with a mixed equilibrium strategy for the n-person game of the final round.

7. CONCLUSION

As we have argued in the introduction, *multilateral negotiation* is often a necessity, *i.e.* it may not be possible to decompose a multilateral deal into a sequence of bilateral deals that would still be acceptable to all parties. In this paper, we have discussed the generalisation of the well-known *monotonic concession protocol* for two negotiating agents to a multilateral scenario where any number of agents should be able to forge agreements. As we have seen, the central question in this context concerns the appropriate definition of a multilateral concession. We have proposed seven alternative definitions, and also argued that there may well be further candidate definitions of interest. Further social welfare orderings, such as the *leximin* ordering for example [8], could provide useful hints for additional definitions.

We have then analysed some of the properties of the negotiation protocols induced by these different definitions: termination, compositionality, deadlock-freedom, and verifiability. All but one of the proposed concession criteria, namely weak concessions, guarantee termination; and weak concessions are also the only ones that are not composable. Our results on deadlock-freedom show that some of the proposed criteria (weak, utilitarian, and egocentric) are deadlock-free, while others are not (Pareto, strong, and egalitarian). Whether or not the Nash concession criterion is deadlock-free depends on the restrictions we put on utility functions, namely whether or not they are required to assign strictly positive utilities to all outcomes but the conflict deal. Our discussion of verifiability issues shows that only the strong concession criterion is verifiable by the participating agents in a distributed manner. The weak and the Pareto criterion can also be taken to be verifiable under certain assumptions on the ability and willingness of the agents to communicate with each other on this matter.

Finally, we have also discussed some of the issues at stake when trying to define a suitable negotiation strategy for agents following a multilateral monotonic concession protocol. In particular, we have suggested three possible generalisations of the well-known Zeuthen strategy, but also pointed out some of the problems inherent in these somewhat simplistic generalisations. A serious game-theoretical analysis of the problem promises to provide a number of new and very interesting insights.

Monotonic concession protocols are but one way of coming to an agreement within a group of agents. As discussed already at the beginning of Section 3, a simple one-step protocol that directly selects a Nash bargaining solution, *i.e.* an agreement maximising the product of individual agent utilities [9, 10], would be another option. In negotiation domains with transferable utilities (*i.e.* if agents can reimburse each other through monetary side payments, and there is global agreement on the marginal utilities associated with such payments), the well-known Shapley value [8, 9] provides yet another solution. The work of Hart and Mas-Colell [5] offers a unified view on this type of approach to multilateral negotiation.

We should stress again that we do not claim that monotonic concession protocols are the "best" way of regulating negotiation, or even that they would be appropriate in all cases. Winoto *et al.* [12], for example, argue against monotonic concessions in domains where utilities can change over time. For instance, a seller's valuation of a perishable good may be assumed to decrease over time, *i.e.* a potential buyer may be well-advised to lower (rather than improve) their offers in case negotiation continues over a significant amount of time. Our argument in favour of monotonic concession protocols is simply that they are such an immediate formalisation of our everyday-understanding of negotiation.

8. REFERENCES

[1] K. J. Arrow, A. K. Sen, and K. Suzumura, editors. *Handbook of Social Choice and Welfare.* North-Holland, 2002.

[2] P. Cramton, Y. Shoham, and R. Steinberg, editors. *Combinatorial Auctions.* MIT Press, 2006.

[3] U. Endriss, N. Maudet, F. Sadri, and F. Toni. On optimal outcomes of negotiations over resources. In *Proc. 2nd International Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS-2003).* ACM Press, 2003.

[4] J. C. Harsanyi. Approaches to the bargaining problem before and after the theory of games: A critical discussion of Zeuthen's, Hick's, and Nash's theories. *Econometrica*, 24(2):144–157, 1956.

[5] S. Hart and A. Mas-Colell. Bargaining and value. *Econometrica*, 64(2):357–380, 1996.

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

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

[8] H. Moulin. *Axioms of Cooperative Decision Making.* Cambridge University Press, 1988.

[9] R. B. Myerson. *Game Theory: Analysis of Conflict.* Harvard University Press, 1991.

[10] J. S. Rosenschein and G. Zlotkin. *Rules of Encounter.* MIT Press, 1994.

[11] T. W. Sandholm. Contract types for satisficing task allocation: I Theoretical results. In *Proc. AAAI Spring Symposium: Satisficing Models*, 1998.

[12] P. Winoto, G. I. McCalla, and J. Vassileva. Non-monotonic-offers bargaining protocol. *Journal of Autonomous Agents and Multiagent Systems*, 11(1):45–67, 2005.

[13] M. Wooldridge. *An Introduction to Multiagent Systems.* John Wiley and Sons, 2002.

[14] M. Wooldridge, S. Bussmann, and M. Klosterberg. Production sequencing as negotiation. In *Proc. 1st International Conference on Practical Application of Agents and Multi-Agent Systems (PAAM-1996)*, 1996.

[15] F. Zeuthen. *Problems of Monopoly and Economic Warfare.* Routledge, 1930.