

Negotiating over Small Bundles of Resources

Yann Chevaleyre¹, Ulle Endriss², Jérôme Lang³ and Nicolas Maudet¹

¹LAMSADE, Université Paris-Daupine (France)

²Department of Computing, Imperial College London (UK)

³IRIT, Université Paul Sabatier, Toulouse (France)

¹{chevaley,maudet}@lamsade.dauphine.fr, ²ue@doc.ic.ac.uk, ³lang@irit.fr

ABSTRACT

When rational but myopic agents negotiate over the exchange of indivisible resources, any restriction to the negotiation protocol may prevent the system from converging to a socially optimal allocation in the general case. This paper addresses this issue by analysing how the confinement to certain classes of utility functions can enable agents to move to an optimal allocation by negotiating over small bundles of items at a time. In particular, we consider so-called k -separable domains, where the full set of resources can be divided into several preferentially independent bundles of limited cardinality.

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

Multiagent resource allocation, Negotiation, Social choice and welfare, Utility theory

1. INTRODUCTION

The allocation of indivisible goods amongst autonomous agents is an important issue in multiagent system research as well as in social choice theory [1, 4, 10, 13]. There are (at least) two different lines of work, depending on how decisions about allocations are made. In centralised approaches, at work in particular in combinatorial auctions, agents simply report their preferences and wait for the final allocation to be made by the auctioneer or some other central entity

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

AAMAS'05, July 25-29, 2005, Utrecht, Netherlands.

Copyright 2005 ACM 1-59593-094-9/05/0007 ...\$5.00.

(whether there is an initial allocation of goods, as in combinatorial exchanges, or not, as in regular combinatorial auctions). In distributed approaches, allocations evolve in a non-synchronised way, by means of local negotiation steps between agents. This line of research has recently been advocated by several authors [6, 7, 8, 14]. Negotiation steps in this framework consist of *deals*, whereby a given subset of agents agree on the reallocation of some of their goods, possibly together with side payments.

Agents are assumed to be *rational*, in the sense of only accepting deals that benefit themselves, as well as *myopic*, *i.e.* they expect a positive payoff from every single deal (rather than being prepared to accept a temporary loss in view of potential future rewards). Multiagent systems are often described as *societies of agents*, and we can use tools borrowed from welfare economics and social choice theory [1, 12], such as the *Pareto* condition or various *social welfare* measures, to assess the quality of negotiation outcomes.

In this framework, it is important to investigate the formal and computational properties of deal sequences, both in the context of implementing autonomous agents (who are expected to negotiate with other agents) and in the context of predicting outcomes reached by a set of interacting autonomous agents. So what are the relevant formal and computational aspects regarding such sequences of deals? Previous work in the area [2, 5, 6, 7, 8, 14, 15] has addressed the following issues:

- *Convergence conditions*: Is it possible to guarantee that any sequence of deals (which may be subject to some constraints) will converge to an optimal allocation (with respect to a given social criterion)?
- *Computational complexity*: How hard is it to determine whether such a sequence leading to a socially optimal allocation exists? How hard is it to find it?
- *Communication complexity*: How many deals are required to reach an optimum, if such an optimum is reachable? What is the necessary amount of information to be exchanged between agents?
- *Communication language*: How should agents encode and communicate the information that needs to be exchanged during negotiation?

These questions cannot be answered without taking the limited computational resources of agents into account. Clearly, considering any possible sequence of any kind of deals between any number of agents is not realistic when the number

of resources and/or agents is large enough. Agents may be rational only to a certain extent, due to the fact that their computational (and cognitive, in the case of human agents) resources are limited. Therefore, we need to investigate the effects of this bounded rationality assumption on the results of the above types.

One way of tackling this issue consists of assuming that agents are only able to negotiate deals over *small bundles* of items. (Clearly, there are also other ways, such as bounding the number of deals in the sequence and/or the number of agents involved in each deal.) Negotiation with a limited number of items per deal has already received some attention, especially in [6], where it is shown that the complexity of checking the existence of a sequence of deals involving one resource each is NP-hard, and in [8], where some conditions on utility functions are identified that guarantee that sequences of deals involving one resource each are sufficient to reach socially optimal allocations.

In this paper we go one step further and generalise the latter type of result by establishing a sufficient condition on utility functions for agents to be able to negotiate an allocation with maximal social welfare by means of a sequence of deals involving at most k items each. We also investigate the case where side payments are not allowed (in which case utilities do not have to be inter-comparable, and preferences can be assumed to be ordinal), and establish a sufficient condition for the existence of a sequence of deals involving k items each leading to a Pareto optimal allocation of resources. We also show that it is not feasible to give conditions of this kind that are both sufficient *and* necessary for reaching optimal outcomes by negotiating over small bundles of items at a time.

Paper overview. The remainder of this paper is organised as follows. Section 2 introduces the negotiation framework and covers relevant previous results [8, 14]. Our results on sufficient conditions for convergence for negotiation over small bundles are presented in Sections 3 and 4: *additively k -separable* utility functions permit negotiation with side payments over k items at a time, while *lexicographically k -separable* preference relations are suitable in case side payments are not possible. Section 5 discusses the problem of giving similar conditions that would be not only sufficient but also necessary, and Section 6 concludes.

2. MYOPIC NEGOTIATION

In this section we introduce the negotiation framework used throughout this paper and report a number of known technical results.

2.1 Negotiation Problems

We analyse scenarios where two or more agents belonging to a set \mathcal{A} negotiate over a number of *indivisible* resources belonging to a set \mathcal{R} . An *allocation* A is a partitioning of \mathcal{R} amongst the agents in \mathcal{A} . For instance, given an allocation A with $A(i) = \{r_5, r_6\}$, agent i owns resources r_5 and r_6 .

We consider two variants of the framework; one where agents express their preferences using utility functions mapping bundles of resources to numerical values and another one where agents only use ordinal preference relations.

DEFINITION 1 (NEGOTIATION PROBLEMS). A *negotiation problem with numerical preferences* is a tuple $\mathcal{P} =$

$\langle \mathcal{R}, \mathcal{A}, u_1, \dots, u_n, A_0 \rangle$, where

- \mathcal{R} is a finite set of indivisible resources;
- $\mathcal{A} = \{1, \dots, n\}$ is a finite set of agents ($n \geq 2$);
- for all $i \in \mathcal{A}$, u_i is a mapping from $2^{\mathcal{R}}$ to \mathbb{R} ;
- $A_0 : \mathcal{A} \rightarrow 2^{\mathcal{R}}$ is an initial allocation.

A *negotiation problem with ordinal preferences* is a tuple $\mathcal{P} = \langle \mathcal{R}, \mathcal{A}, \preceq_1, \dots, \preceq_n, A_0 \rangle$ where \mathcal{R} , \mathcal{A} and A_0 are as before and for all $i \in \mathcal{A}$, \preceq_i is a complete weak order on $2^{\mathcal{R}}$.

Every preference relation \preceq_i also gives rise to an indifference relation \sim_i and a strict preference relation \prec_i :

- $R_1 \sim_i R_2$ iff $R_1 \preceq_i R_2$ and $R_2 \preceq_i R_1$;
- $R_1 \prec_i R_2$ iff not $R_2 \preceq_i R_1$.

Furthermore, every utility function induces an ordinal preference relation: $R_1 \preceq_i R_2$ iff $u_i(R_1) \leq u_i(R_2)$. Therefore, every negotiation problem with numerical preference also induces a negotiation problem with ordinal preferences.

2.2 The Social Perspective

While preference relations and utility functions are used to model the interests of *individual* agents, as system designers, we are also interested in assessing the quality of a given allocation from a *social* point of view. Several concepts that allow us to do this have been developed in welfare economics and social choice theory [1, 12]. We write $A \preceq_i A'$ as a shorthand for $A(i) \preceq_i A'(i)$.

DEFINITION 2 (PARETO OPTIMALITY). A *resource allocation* A is *Pareto optimal* iff there is no other allocation A' such that $A \preceq_i A'$ for all $i \in \mathcal{A}$ and $A \prec_i A'$ for at least one $i \in \mathcal{A}$.

Pareto optimality is often considered a minimal requirement for socially desirable allocations. A somewhat stronger concept is that of social welfare. We abbreviate $u_i(A) = u_i(A(i))$ for the utility agent i assigns to the bundle it receives in allocation A . The following definition applies only to the framework with numerical preferences:

DEFINITION 3 (SOCIAL WELFARE). The *social welfare* $sw(A)$ of an allocation A is defined as follows:

$$sw(A) = \sum_{i \in \mathcal{A}} u_i(A)$$

This is the *utilitarian* definition of social welfare. We should stress that other notions of social welfare have been developed as well [1, 12].

2.3 Deals and Rationality

Given a particular allocation of resources, agents may agree on a (multilateral) *deal* to exchange some of the resources they currently hold. In general, a single deal may involve any number of resources and any number of agents. It transforms an allocation of resources A into a new allocation A' ; that is, we can define a deal as a pair $\delta = (A, A')$ of allocations (with $A \neq A'$).

A *k-deal* is a deal involving at most k resources, *i.e.* a deal $\delta = (A, A')$ such that:

$$|\mathcal{R} \setminus \bigcup_{i \in \mathcal{A}} (A(i) \cap A'(i))| \leq k$$

A deal may be coupled with a number of monetary side payments to compensate some of the agents involved for an otherwise disadvantageous deal. This can be modelled using a *payment function* $p : \mathcal{A} \rightarrow \mathbb{R}$. Such a function has to satisfy the side constraint $\sum_{i \in \mathcal{A}} p(i) = 0$, *i.e.* the overall amount of money in the system remains constant. If $p(i) > 0$, then agent i *pays* the amount of $p(i)$, while $p(i) < 0$ means that it *receives* the amount of $-p(i)$. We distinguish deals with side payments and deals without side payments. For the latter, $p(i)$ is required to be 0 for every agent $i \in \mathcal{A}$.

An agent may or may not find a particular deal acceptable. We shall assume that agents are *rational* in the sense of never accepting a deal that would not improve their personal welfare (see [14] for a justification of this approach). For negotiation problems with numerical preferences and deals with side payments, this myopic notion of individual rationality may be formalised as follows:

DEFINITION 4 (INDIVIDUAL RATIONALITY). *A deal $\delta = (A, A')$ with side payments is rational iff there exists a payment function p such that $u_i(A') - u_i(A) > p(i)$ for all $i \in \mathcal{A}$, except possibly $p(i) = 0$ for agents i with $A(i) = A'(i)$.*

In scenarios where side payments are not allowed, agents will be required to be cooperative in the sense of also accepting deals that do not result in a strict increase in personal welfare (see [8] for a justification):

DEFINITION 5 (COOPERATIVE RATIONALITY). *A deal $\delta = (A, A')$ without side payments is rational iff $A \preceq_i A'$ for all $i \in \mathcal{A}$ and $A \prec_i A'$ for at least one $i \in \mathcal{A}$.*

The second part of the definition ensures that at least one agent (say, the one proposing the deal) will have a strictly positive payoff for every rational deal. This condition is required to ensure the termination of negotiation processes.

For negotiation problems with numerical preferences, the following lemma establishes an important relationship between the local concept of individual rationality and the global concept of social welfare [8]:

LEMMA 1 (RATIONAL DEALS). *A deal $\delta = (A, A')$ with side payments is rational iff $sw(A) < sw(A')$.*

Note that the same is *not* true for deals without side payments. While any cooperatively rational deal will increase social welfare (for negotiation problems with numerical preferences), there are social welfare-increasing deals that can only be made rational by using side payments.

2.4 Known Convergence Results

The following theorem is due to Sandholm [14]:

THEOREM 1 (MAXIMISING SOCIAL WELFARE). *Any sequence of rational deals with side payments will eventually result in an allocation of resources with maximal social welfare.*

Endriss et al. [8] also prove the following results:

THEOREM 2 (ADDITIVE DOMAINS). *If all utility functions are additive, then any sequence of rational 1-deals with side payments will eventually result in an allocation of resources with maximal social welfare.*

THEOREM 3 (PARETO OPTIMALITY). *Any sequence of rational deals without side payments will eventually result in a Pareto optimal allocation of resources.*

What all these results have in common is that they guarantee that agents can agree on *any* sequence of deals meeting the respective rationality condition without getting stuck in a local minimum. Reaching a socially optimal allocation always remains possible and there can be no infinite sequence of rational deals. In addition to this, Theorem 2 shows that simple deals involving only a single resource (and thereby only two agents) at a time suffice to negotiate an optimal allocation in a restricted type of domain where all utility functions are additive (*i.e.* the utility assigned to a bundle is always the sum of utilities assigned to the resources in that bundle). This is not the case in general; if there are no restrictions on utility functions, then deals involving any number of agents and resources may be required [8, 14].

In this paper, we are going to prove two generalisations of the above theorems.

3. ADDITIVE k -SEPARABILITY

In this section, we introduce the class of *additively k -separable* utility functions, which are useful in domains where the full set of resources \mathcal{R} can be partitioned into several preferentially independent bundles of at most k items each. We discuss some of the properties of this class of utility functions and show that agents with additively k -separable utility functions can always negotiate an allocation with maximal social welfare even if every single deal is required to be rational with side payments and may involve at most k resources.

3.1 Definition of Additive k -Separability

Throughout this paper, let k be a (fixed) natural number.

DEFINITION 6 (ADDITIVE k -SEPARABILITY). *Let $P = \langle R_1, \dots, R_q \rangle$ be a partition of \mathcal{R} . A utility function u is called *additively k -separable* wrt. P iff (i) $|R_j| \leq k$ for all $j \in \{1..q\}$ and (ii) the following holds for all $R \subseteq \mathcal{R}$:*

$$u(R) = u(\{\}) + \sum_{j=1}^q [u(R \cap R_j) - u(\{\})] \quad (1)$$

We are going to refer to the elements R_j in the partition as *topics*. In essence, the definition says that a domain is k -separable iff each topic consists of at most k resources and there are no synergies between items belonging to distinct topics. Observe how equation (1) simplifies for $u(\{\}) = 0$, which is a reasonable assumption in most domains.

What happens when we choose extreme values for the parameter k ? Firstly, if we set $k = 1$ (in which case there is just a single possible partition of \mathcal{R}) and if we assume $u(\{\}) = 0$, then the class of additively k -separable utility functions reduces to the class of *additive* functions. Secondly, only by choosing $k = |\mathcal{R}|$ and a “partition” consisting of just a single topic can we ensure that *any* utility function is k -separable.

3.2 Compact Representation of Preferences

The “normal form” of representing utility functions, which involves listing all bundles of resources with non-zero utility, can be problematic as there may be up to $2^{|\mathcal{R}|}$ such bundles in the worst case. Utility functions that are k -separable

often permit a more compact representation than would be possible with the general notation used thus far, particularly for small values of k .

Let u be a utility function that is additively k -separable with respect to a partition $\langle R_1, \dots, R_q \rangle$. Now define a *local* utility function $u^j : 2^{R_j} \rightarrow \mathbb{R}$ for each topic R_j as follows: $u^j(R) = u(R \cap R_j) - u(\{\})$ for all $R \subseteq \mathcal{R}$. Furthermore, define $c = u(\{\})$. Then u can be written as follows:

$$u(R) = c + \sum_{j=1}^q u^j(R) \quad \text{for all } R \subseteq \mathcal{R}$$

The maximal number of non-zero values to be specified is now $1 + \sum_{j=1}^q (2^{|R_j|} - 1) \leq q \cdot 2^k$ rather than $2^{|\mathcal{R}|}$. Furthermore, specifying an additively k -separable utility function using this notation with local topic utilities will require the specification of *at most* as many non-zero values as in the normal form in the worst case (and considerably fewer in most cases).

The succinctness of the representation of agent preferences can be further improved by exploiting other structural properties of the utility functions. For instance, if synergies between different resources are restricted to bundles of at most k items (but these k items need not all belong to the same topic, as for k -separable functions), then the so-called *k-additive form* which specifies for each bundle R the marginal utility of owning all resources in R can often result in a more efficient representation [2]. For a detailed discussion of the subject of compact preference representation, particularly for ordinal preferences, we refer to the recent work of Lang [11].

3.3 Convergence Result

We call a *domain* additively k -separable iff the utility functions of all agents are additively k -separable with respect to the *same* partition of \mathcal{R} .

THEOREM 4 (ADDITIVE k -SEPARABILITY). *In additively k -separable domains, any sequence of rational k -deals with side payments will eventually result in an allocation with maximal social welfare.*

PROOF. As any rational deal with side payments results in a strict increase in social welfare (Lemma 1) and as the space of possible allocations is finite, any negotiation process where all deals are required to be rational is bound to terminate eventually. It therefore suffices to show that for any allocation A that does not have maximal social welfare there still exists a rational k -deal with side payments.

Let A be such an allocation, *i.e.* there exists a further allocation A' such that $sw(A) < sw(A')$. We are going to use this superior allocation to identify a k -deal that applies to A and that is rational with side payments. Let $\langle R_1, \dots, R_q \rangle$ be the partition of \mathcal{R} with respect to which our negotiation domain is k -separable. We rewrite equation (1) as follows:

$$u(R) = (1 - q) \cdot u(\{\}) + \sum_{j=1}^q u_i(R \cap R_j) \quad (2)$$

Next we define c_q^A as follows:

$$c_q^A = \sum_{i \in \mathcal{A}} (1 - q) \cdot u_i(\{\}) \quad (3)$$

Using equations (2) and (3), we can reformulate the definition of social welfare as follows:

$$\begin{aligned} sw(A) &= \sum_{i \in \mathcal{A}} u_i(A(i)) \\ &= c_q^A + \sum_{i \in \mathcal{A}} \sum_{j=1}^q u_i(A(i) \cap R_j) \\ &= c_q^A + \sum_{j=1}^q \sum_{i \in \mathcal{A}} u_i(A(i) \cap R_j) \end{aligned}$$

This last transformation (permutation of the sum over agents and the sum over topics) is one of the central steps in our proof. The same sequence of transformations can be applied to $sw(A')$, *i.e.* we obtain the following inequation:

$$\sum_{j=1}^q \sum_{i \in \mathcal{A}} u_i(A(i) \cap R_j) < \sum_{j=1}^q \sum_{i \in \mathcal{A}} u_i(A'(i) \cap R_j)$$

Hence, there exists a (at least one) $j^* \in \{1..q\}$ such that the following holds:

$$\sum_{i \in \mathcal{A}} u_i(A(i) \cap R_{j^*}) < \sum_{i \in \mathcal{A}} u_i(A'(i) \cap R_{j^*}) \quad (4)$$

We are now in a position to identify an allocation A^* such that the deal $\delta = (A, A^*)$ is a rational k -deal with side payments. For all agents $i \in \mathcal{A}$, define:

$$A^*(i) = [A(i) \setminus R_{j^*}] \cup [A'(i) \cap R_{j^*}]$$

That is, A^* is like A' with respect to the resources in R_{j^*} and like the current allocation A with respect to all other resources. Clearly, δ must be a k -deal, because it only involves items in R_{j^*} and $|R_{j^*}| \leq k$ (by definition of k -separability). It remains to be shown that δ is also rational.

By adding the sum $c_q^A + \sum_{j \in \{1..q\}} \sum_{i \in \mathcal{A}} u_i(A(i) \cap R_j)$ to either side of (4), we obtain the following inequation:

$$\begin{aligned} c_q^A + \sum_{j=1}^q \sum_{i \in \mathcal{A}} u_i(A(i) \cap R_j) &< \\ c_q^A + \sum_{i \in \mathcal{A}} u_i(A'(i) \cap R_{j^*}) + \sum_{j \in \{1..q\}} \sum_{i \in \mathcal{A}} u_i(A(i) \cap R_j) & \end{aligned}$$

The lefthand side is equivalent to $sw(A)$, while the righthand side is equivalent to $sw(A^*)$, *i.e.* we have shown that $sw(A) < sw(A^*)$. By Lemma 1, the latter implies that $\delta = (A, A^*)$ is a rational deal with side payments.

Hence, as long as the current allocation is not optimal, it is still possible to find a rational deal with side payments that involves at most k resources, and an optimal allocation will be reached after a finite number of such deals. \square

Observe that both Theorem 1 and Theorem 2 are in fact instances of Theorem 4. This follows from our earlier discussion of the cases $k = |\mathcal{R}|$ and $k = 1$.

What are the crucial properties of additively k -separable domains that bring about the convergence result of Theorem 4? Firstly, in *separable* domains, the marginal utility gains associated with the resources belonging to different topics are entirely independent. Therefore, we could choose to always negotiate over the resources belonging to one topic at a time. Secondly, because of the cardinality restriction of *at most* k resources per topic, k -deals are sufficient for

this kind of negotiation. Finally, the fact that, in *additively k-separable* domains, overall utility is defined as the sum of local topic utilities means that an allocation that maximises social welfare with respect to every individual topic will also maximise social welfare at the global level.

This last remark also explains why there can be no similar result for the framework without side payments, where we aim at negotiating Pareto optimal allocations. In additively *k-separable* domains, Pareto optimality with respect to individual topics does *not* entail Pareto optimality with respect to the full set of resources. Our aim for the next section will be to identify a variant of *k-separability* that *does* allow the property of Pareto optimality to transfer in this manner and that therefore allows us to prove a convergence result for rational *k-deals* without side payments.

4. LEXICOGRAPHIC *k*-SEPARABILITY

In this section, we introduce the class of *lexicographically k-separable* preference relations. As for additively *k-separable* domains, in lexicographic *k-separable* domains the set of resources \mathcal{R} can be divided into several preferentially independent topics of at most *k* resources each. In addition, agents agree on the relative importance of different topics. As we shall see, this domain restriction is sufficient to guarantee that rational *k-deals* without side payments will eventually lead to a Pareto optimal allocation.

4.1 Definition of Lexicographic *k*-separability

In this section, we are going to be concerned with negotiation problems with ordinal preferences, so we define the concept of lexicographic *k-separability* with respect to preference relations rather than utility functions.

DEFINITION 7 (LEXICOGRAPHIC *k*-SEPARABILITY). *Let $P = \langle R_1, \dots, R_q \rangle$ be an ordered partition of \mathcal{R} . A preference relation \preceq is called lexicographically *k-separable* wrt. P iff (i) $|R_j| \leq k$ for all $j \in \{1..q\}$ and (ii) there exist preference relations $\preceq_{R_1}, \dots, \preceq_{R_q}$ on $2^{R_1}, \dots, 2^{R_q}$, respectively, and an index $i \leq q$ such that the following holds for all bundles $R, R' \subseteq \mathcal{R}$:*

$$R \prec R' \quad \text{iff} \quad (R \cap R_i) \prec_i (R' \cap R_i) \text{ and} \\ (R \cap R_j) \sim_j (R' \cap R_j) \text{ for all } j < i$$

If there are no restrictions on the cardinality of topics, we also speak of lexicographically separable preference relations (without specifying a value *k*). A lexicographically *k-separable utility function* is a utility function that induces a lexicographically *k-separable* preference relation.

Intuitively, lexicographic *k-separability* is much stronger a condition than additive *k-separability*. Indeed, any lexicographically *k-separable* preference relation is *representable* by an additively *k-separable* utility function [9]. We should stress, however, that this does *not* mean that every lexicographic *k-separable* utility function is also additively *k-separable*.

4.2 Convergence Result

As indicated at the end of Section 3, a crucial prerequisite for proving a convergence result for Pareto optimality would be to show that this property transfers from individual topics to the full set of resources. This is the purpose of the following lemma:

LEMMA 2 (TRANSFER). *Let $P = \langle R_1, \dots, R_q \rangle$ be an ordered partition of \mathcal{R} and let $\mathcal{P} = \langle \mathcal{R}, \mathcal{A}, \preceq_1, \dots, \preceq_n, A_0 \rangle$ be a negotiation problem with ordinal preferences such that \preceq_i is lexicographically *k-separable* wrt. P for all $i \in \mathcal{A}$. For every $j \in \{1..q\}$, define the *j-projection* \mathcal{P}^j of \mathcal{P} as the negotiation problem $\mathcal{P}^j = \langle R_j, \mathcal{A}, \preceq_1^j, \dots, \preceq_n^j, A_0^j \rangle$, where \preceq_i^j is the restriction of \preceq_i to 2^{R_j} and for any allocation A (including A_0), A^j is defined by $A^j(i) = A(i) \cap R_j$ for all $i \in \mathcal{A}$. Then an allocation A is Pareto optimal for \mathcal{P} iff A^j is Pareto optimal for \mathcal{P}^j for all $j \in \{1..q\}$.*

PROOF. First, let A be an allocation for \mathcal{P} and suppose that for some j , A^j is not Pareto optimal for \mathcal{P}^j . Then there exists an allocation B^j of the resources in R_j such that B^j Pareto-dominates A^j for the problem \mathcal{P}^j . Consider then the allocation A' identical to A for all $l \neq j$ and where A^j is replaced by B^j , that is:

$$A'(i) = \bigcup_{l \neq j} A^l(i) \cup B^j(i) \quad \text{for all } i \in \mathcal{A}$$

B^j Pareto-dominates A^j for \mathcal{P}^j , therefore we have (a) for all i , $A^j \preceq_i^j B^j$ and (b) there is an i^* such that $A^j \prec_{i^*}^j B^j$. Furthermore, for every agent i and every $l \neq j$ we have $A^l(i) = A^l(i)$, therefore $A^l \sim_i A^l$. Now, since \preceq_i is lexicographically separable wrt. P , we get: (a') for all i , $A \preceq_i A'$ and (b') $A \prec_{i^*} A'$. This shows that A' Pareto-dominates A wrt. \mathcal{P} , therefore A is not Pareto optimal.

Conversely, suppose that A is not Pareto optimal and let then be A' such that A' Pareto-dominates A wrt. \mathcal{P} . Now, for all $i \in \mathcal{A}$, let $d(A, A', i) = \min\{l \mid A^l \prec_i A'^l\}$ and $j^* = \min\{d(A, A', i) \mid i \in \mathcal{A}\}$. The existence of j^* is guaranteed by the fact that A' Pareto-dominates A . Now, let $l < j^*$ and let i be some agent. From the definition of j^* it cannot be the case that $A^l \prec_i A'^l$; hence either $A^l \sim_i A'^l$ or $A^l \succ_i A'^l$. Now, if we had $A^l \succ_i A'^l$ then we would have $A \succ_i A'$ (just consider the smallest l such that $A^l \succ_i A'^l$ and apply the definition of lexicographic *k-separability*), which would contradict the assumption that A' Pareto-dominates A . Therefore, (a) for all $l < j^*$ and all agents i , $A^l \sim_i A'^l$. Lastly, if for some agent i we had $A^{j^*} \succ_i A'^{j^*}$ then, due to (a), we would have $A \succ_i A'$, which again contradicts the fact that A' Pareto-dominates A ; therefore, for every agent i , either $A^{j^*} \sim_i A'^{j^*}$ or $A^{j^*} \prec_i A'^{j^*}$, with $A^{j^*} \prec_i A'^{j^*}$ holding for at least one i (by definition of j^*). Therefore, A^{j^*} Pareto-dominates A'^{j^*} . In other words, there is indeed a j (namely j^*) such that A^j is not Pareto optimal for \mathcal{P}^j . \square

We are now in a position to prove the desired convergence result for Pareto optimal outcomes. A *domain* is called lexicographically *k-separable* iff the preference relations of all agents are lexicographically *k-separable* with respect to the *same* ordered partition.

THEOREM 5 (LEXICOGRAPHIC *k*-SEPARABILITY). *In lexicographically *k-separable* domains, any sequence of rational *k-deals* without side payments will eventually result in a Pareto optimal allocation of resources.*

PROOF. First observe that there can be no infinite sequence of rational deals without side payments (because in each deal at least one of the agents moves to a *strictly* preferred bundle). Therefore, it is sufficient to show that for any allocation A that is not Pareto optimal, it is still possible to implement a rational *k-deal* without side payments.

Let $P = \langle R_1, \dots, R_q \rangle$ be an ordered partition of \mathcal{R} and let $\mathcal{P} = \langle \mathcal{R}, \mathcal{A}, \preceq_1, \dots, \preceq_n, A_0 \rangle$ be a negotiation problem such that \preceq_i is lexicographically k -separable with respect to P for all $i \in \mathcal{A}$. Now, let A be an allocation that is not Pareto optimal for \mathcal{P} . By Lemma 2, there exists a $j \in \{1..q\}$ such that the allocation A^j of the j -projection \mathcal{P}^j of \mathcal{P} is not Pareto optimal either. Hence, by Theorem 3, there exists a rational deal without side payments for the negotiation problem \mathcal{P}^j , starting from allocation A^j . This deal involves at most k resources (namely only resources from R_j) and it must be rational for \mathcal{P} as well. \square

For $k = |\mathcal{R}|$ and $q = 1$, Theorem 5 reduces to Theorem 3. (The use of the latter result has simplified our proof of Theorem 5, but a direct proof is possible as well.)

5. NECESSARY CONDITIONS

We have seen earlier that additive k -separability is a *sufficient* condition for rational k -deals with side payments to converge to an allocation with maximal social welfare (Theorem 4). Is it also a *necessary* condition?

In this section, we first give a negative answer to this question by constructing a simple counterexample. We then investigate whether there exists *any* sufficient condition on utility functions that is also necessary in this sense. As we shall see, there are at least two interpretations of the concept of a necessary and sufficient condition. The first of these would be a condition on *single* functions (to be met by the utility functions of all agents). The second would be a condition on *profiles* of utility functions directly. We are going to discuss two negative results: there can be no condition of the first type and verifying the second type of condition is computationally intractable.

The discussion in this section applies to negotiation problems with numerical preferences and to deals with side payments.

5.1 Example

The following example demonstrates that there are scenarios where not all agents have additively k -separable utility functions, but negotiating an allocation with maximal social welfare by means of rational k -deals with side payments is still possible for any given initial allocation. Suppose $\mathcal{R} = \{r_1, r_2\}$ and there are two agents with utility functions u_1 and u_2 :

| | | | |
|---------------------|------|---------------------|------|
| $u_1(\{\})$ | = 90 | $u_2(\{\})$ | = 90 |
| $u_1(\{r_1\})$ | = 93 | $u_2(\{r_1\})$ | = 90 |
| $u_1(\{r_2\})$ | = 95 | $u_2(\{r_2\})$ | = 90 |
| $u_1(\{r_1, r_2\})$ | = 98 | $u_2(\{r_1, r_2\})$ | = 50 |

While u_1 is additively 1-separable, u_2 is not. The optimal allocation is the allocation where agent 1 owns both items. Furthermore, as may easily be checked, any 1-deal that involves moving a single resource from agent 2 to agent 1 is rational. Hence, rational 1-deals are sufficient to move to the optimal allocation for this scenario, despite u_2 not being 1-separable.

It is not difficult to construct examples like this. To find this particular one, we have taken another example like the above but with $u_2(\{r_1, r_2\}) = 90$ as a starting point. Now u_2 is additively 1-separable as well and Theorem 4 applies, telling us that rational 1-deals are guaranteed to suffice for any initial allocation. We have then revised the value for

$u_2(\{r_1, r_2\})$ downwards to make the worst allocation even worse. Clearly, this local change does not affect the qualitative structure of the negotiation domain, and any deal that has been rational in the original setting remains to be so after the revision. However, what local changes of this kind are possible depends on the concrete example under consideration. Therefore, this approach is not helpful in identifying an *interesting* class of utility functions that permit rational negotiation with k -deals.

5.2 Notions of Necessity

A negotiation domain is characterised by the utility functions of the agents in the system. We say that the profile of utility functions $\langle u_1..u_n \rangle$ *permits k -deal negotiation* iff any sequence of rational k -deals with side payments will eventually result in an allocation with maximal social welfare for any given initial allocation. For instance, as Theorem 4 has shown, any profile $\langle u_1..u_n \rangle$ where all u_i are additively k -separable with respect to the same partition of \mathcal{R} permits k -deal negotiation.

For the following definitions, assume the set of agents \mathcal{A} , the set of resources \mathcal{R} , and the parameter k are all fixed:

- A condition *Cond* on single utility functions is called necessary and sufficient for k -deal negotiation iff all profiles $\langle u_1..u_n \rangle$ that meet *Cond*(u_i) for all $i \in \mathcal{A}$ and only those permit k -deal negotiation.
- A condition *Cond* on profiles of utility functions is called necessary and sufficient for k -deal negotiation iff all profiles $\langle u_1..u_n \rangle$ that meet *Cond*($u_1..u_n$) and only those permit k -deal negotiation.

5.3 Conditions on Single Utility Functions

We show that it is impossible to find a necessary and sufficient condition on single utility functions for k -deal negotiation. We give the proof for $k = 1$ (the proof for $k > 1$ is similar). Of course, the number of resources needs to exceed k for this result to be valid (otherwise the tautological condition would be both necessary and sufficient).

THEOREM 6 (INEXISTENCE). *There exists no necessary and sufficient condition on single utility functions for 1-deal negotiation over more than one resource.*

PROOF. Let $\mathcal{R} = \{r_1, \dots, r_p\}$ with $p > 1$. Assume that *Cond* is a necessary and sufficient condition for 1-deal negotiation. By Theorem 2, all utility profiles consisting of additive functions permit 1-deal negotiation; therefore all additive functions must meet *Cond*. We define the class of *quasi-constant* utility functions as those satisfying $u(R) = u(R')$ for all non-empty bundles R and R' , and $u(\{\}) = 0$. Clearly, any utility profile $\langle u_1..u_n \rangle$ where each u_i is quasi-constant also permits 1-deal negotiation, *i.e.* all quasi-constant utility functions meet *Cond* as well. Now, let us consider the following two utility functions u_1 and u_2 :

- $u_1(\{\}) = 0$; $u_1(R) = 1$ for every $R \neq \{\}$;
- u_2 is the additive utility function with $u_2(\{r_1\}) = \frac{1}{4}$, $u_2(\{r_2\}) = \frac{3}{4}$, and $u_2(\{r_j\}) = 0$ for every $j > 2$.

Lastly, for each $i > 2$ let $u_i(R) = 0$ for all R . The function u_1 is pseudo-constant; u_2, \dots, u_n are additive. Therefore, u_i meets *Cond* for every $i \in \mathcal{A}$.

Now, consider an initial allocation A_0 such that $A_0(1) = \{r_2\}$ and $A_0(2) = \{r_1\}$. Clearly, no rational 1-deal is possible from A_0 . Furthermore, A_0 does not have maximal social welfare, because $sw(A_0) = \frac{5}{4}$ and we have $sw(A) = \frac{7}{4}$ for any allocation A with $A(1) = \{r_1\}$ and $A(2) = \{r_2\}$. Hence, the profile $\langle u_1..u_n \rangle$ does *not* permit 1-deal negotiation, *i.e.* $Cond$ cannot be sufficient for 1-deal negotiation. This is a contradiction to our assumptions on $Cond$, *i.e.* there can be no condition that is both necessary and sufficient for 1-deal negotiation. \square

5.4 Conditions on Profiles of Utility Functions

Clearly, there is no such problem for the second variant of defining necessary and sufficient conditions for k -deal negotiation. For any set of agents \mathcal{A} , set of resources \mathcal{R} , and choice of k , there *does* exist a condition $Cond$ on profiles of utility functions such that $Cond(u_1..u_n)$ holds iff $\langle u_1..u_n \rangle$ permits k -deal negotiation. However, here we encounter a different problem: checking whether a given profile of utility functions actually meets this condition will typically be intractable.

To motivate this statement, consider the case $k = 1$. Suppose $Cond$ is the necessary and sufficient condition on profiles of utility functions for 1-deal negotiation. That is, checking whether $Cond(u_1..u_n)$ holds is equivalent to checking whether $\langle u_1..u_n \rangle$ permits 1-deal negotiation. This problem has recently been shown to be coNP-hard by Paul E. Dunne (personal communication, March 2005). The proof explores a reduction from the complement of 3-SAT, similar to previous work by Dunne and colleagues [6].

This result strongly suggests that checking necessary and sufficient conditions for $k > 1$ will also be intractable. Only when k approaches $|\mathcal{R}|$, the number of resources in the system, the problem becomes easier again (but also irrelevant in practice). Certainly for $k = |\mathcal{R}|$, checking the condition becomes trivial as any profile of utility functions will meet it (this follows from Theorem 1). In summary, our analysis shows that the availability of a necessary and sufficient condition on *profiles* of utility functions for permitting k -deal negotiation would be only of limited use in practice. Checking whether a given profile of functions meets that condition would be intractable, even if $k = 1$.

6. CONCLUSION

We have further analysed a negotiation framework previously studied by several authors [2, 5, 6, 7, 8, 14]. While most work on negotiation in multiagent systems has addressed either bilateral negotiation [13] or auctions [4], this framework is multilateral, *i.e.* deals may involve any number of agents and any number of resources. The aim of this paper has been to investigate means to control at least one of these parameters, namely the number of resources affected by a single deal.

The requirement for full multilateral negotiation stems from the fact that agents are assumed to be both rational and myopic. This is true if agents can use arbitrary utility functions or preference relations. As we have shown, (i) if all agents use *additively k -separable* utility functions, then rational k -deals with side payments are sufficient to negotiate allocations with maximal social welfare (Theorem 4); and (ii) if all agents have *lexicographically k -separable* preferences, then rational k -deals without side payments are sufficient to reach Pareto optimal allocations (Theorem 5).

Our analysis in Section 5 has shown that it is not feasible to formulate sufficient conditions on agent preferences that are also *necessary* to permit negotiation over small bundles of items. An alternative approach would be to consider *maximality* properties with respect to the sufficient classes of preferences identified in this paper. A sufficient condition for k -deal negotiation is maximal iff no condition strictly subsumed by the former is also sufficient. In a recent paper we have addressed this issue for the case $k = 1$ [3]; a full analysis of the general case poses an interesting challenge for future work.

Future research should also be directed towards identifying properties of negotiation problems that would allow us to place an upper bound on the *number of agents* involved in each deal, as well as properties of negotiation problems that can be exploited to limit the *number of deals* in a sequence leading to an optimal allocation [5, 7].

7. REFERENCES

- [1] K. J. Arrow, A. K. Sen, and K. Suzumura, editors. *Handbook of Social Choice and Welfare*, volume 1. North-Holland, 2002.
- [2] Y. Chevaleyre, U. Endriss, S. Estivie, and N. Maudet. Multiagent resource allocation with k -additive utility functions. In *Proc. DIMACS-LAMSADE Workshop on Computer Science and Decision Theory*, Annales du LAMSADE 3, 2004.
- [3] Y. Chevaleyre, U. Endriss, and N. Maudet. On maximal classes of utility functions for efficient one-to-one negotiation. In *Proc. IJCAI-2005*. Morgan Kaufmann Publishers, 2005.
- [4] P. Cramton, Y. Shoham, and R. Steinberg, editors. *Combinatorial Auctions*. MIT Press, 2005. To appear.
- [5] P. E. Dunne. Extremal behaviour in multiagent contract negotiation. *Journal of Artificial Intelligence Research*, 23:41–78, 2005.
- [6] P. E. Dunne, M. Wooldridge, and M. Laurence. The complexity of contract negotiation. *Artificial Intelligence*, 164(1–2):23–46, 2005.
- [7] U. Endriss and N. Maudet. On the communication complexity of multilateral trading. In *Proc. AAMAS-2004*. ACM Press, 2004.
- [8] U. Endriss, N. Maudet, F. Sadri, and F. Toni. On optimal outcomes of negotiations over resources. In *Proc. AAMAS-2003*. ACM Press, 2003.
- [9] P. C. Fishburn. *Utility Theory for Decision Making*. John Wiley and Sons, 1970.
- [10] S. Kraus. *Strategic Negotiation in Multiagent Environments*. MIT Press, 2001.
- [11] J. Lang. Logical preference representation and combinatorial vote. *Annals of Mathematics and Artificial Intelligence*, 42(1–3):37–71, 2004.
- [12] H. Moulin. *Axioms of Cooperative Decision Making*. Cambridge University Press, 1988.
- [13] J. S. Rosenschein and G. Zlotkin. *Rules of Encounter*. MIT Press, 1994.
- [14] T. W. Sandholm. Contract types for satisficing task allocation: I Theoretical results. In *Proc. AAAI Spring Symposium: Satisficing Models*, 1998.
- [15] M. Wooldridge and S. Parsons. Languages for negotiation. In *Proc. ECAI-2000*. IOS Press, 2000.