Results and Open Questions in Distributed Resource Allocation

Yann Chevaleyre[*], Ulle Endriss [†] and Nicolas Maudet[*]

[*] Paris-Dauphine, Place du Maréchal de Lattre de Tassigny, 75775 Paris cedex 16, France.
[†] Department of Computing, Imperial College London, 180 Queen's Gate, London SW7 2AZ, UK.

Abstract When rational but myopic agents negotiate over the exchange of indivisible resources, any restriction to the negotiation protocol or to the preference structure of the agents may reduce the complexity of the process, but also prevent the system from converging to a socially optimal allocation in the general case. This paper reviews some recent theoretical results addressing this issue. More precisely, It analyses how the confinement to simple deals and to some restricted classes of utility functions can enable agents to move to an optimal allocation, while reducing the overall computational complexity. The case of complex deals is also studied, and both restrictions on utility functions and specially designed protocols are proposed which drastically reduce the computational complexity of the resource allocation process.

1 Introduction

A multiagent system may be thought of as an artificial society of autonomous software agents. Negotiation over the distribution of resources (or tasks) amongst the agents inhabiting such a society is an important area of research in artificial intelligence and computer science (Chavez et al. (1997); Kraus (2001); Rosenschein and Zlotkin (1994a); Sandholm (1999)). A number of variants of this problem have been studied in the literature. Here we consider the case of an artificial society of agents where, to begin with, each agent holds a bundle of discrete (*i.e.* non-divisible) resources to which it assigns a certain utility. Agents may then negotiate with each other in order to agree on the redistribution of some of these resources to benefit either themselves or the agent society they inhabit.

Rather than being concerned with specific strategies for negotiation (as in Kraus (2001); Rosenschein and Zlotkin (1994a)) or concrete communication protocols to enable negotiation in a multiagent system (see Smith (1980)), we analyse how the redistribution of resources by means of negotiation affects the well-being of the agent society as a whole. To this end, we make use of formal tools for measuring social welfare developed in welfare economics and social choice theory (Arrow et al. (2002); Moulin (1988a)). In multiagent systems research, the concept of social welfare is usually given a utilitarian interpretation, *i.e.* whatever increases the average welfare of the agents inhabiting a society is taken to be beneficial for society as well. Of course, other notions such as egalitarian social welfare (Sen (1970)), Lorenz optimality (Moulin (1988a)), or envy-freeness (Brams and Taylor (1996)) may be usefully exploited when designing multiagent systems.

Recently, much work has focused on analysing the conditions under which the optimal social welfare could be reached. In this paper we enumerate the main theoretical results related to this issue, and raise some opened questions. The remainder of this paper is organised as follows. The next section introduces the basic negotiation framework for resource reallocation we are going to consider. Then, convergence results are shown in the setting in which only a single resource is allowed to be exchanged at each time step. This restriction is then raised, deals exchanging at most k resources at a time are studied. Finally, a further restriction on utility functions is shown to make the process even more tractable.

2 Resource Allocation by Negotiation

In this section, we introduce the framework of *resource allocation by negotiation* put forward in Endriss et al. (2003a) and recall some of the results presented there.

2.1 The Negotiation Framework

An instance of our negotiation framework consists of a finite set of (at least two) *agents* $\mathcal{A}$ and a finite set of non-divisible *resources* $\mathcal{R}$. A resource *allocation* A is a partitioning of the set $\mathcal{R}$ amongst the agents in $\mathcal{A}$. For instance, given an allocation A with $A(i) = \{r_3, r_7\}$, agent i would own resources r_3 and r_7. 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 deal may be coupled with a number of monetary side payments to compensate some of the agents involved for an otherwise disadvantageous deal. Rather than specifying for each pair of agents how much the former is supposed to pay to the latter, we simply say how much money each and every agent either pays out or receives. This can be modelled using a *payment function* p mapping agents in $\mathcal{A}$ to real numbers. 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 $-p(i)$.

2.2 Individual Rationality and Social Welfare

To measure their individual welfare, every agent $i \in \mathcal{A}$ is equipped with a *utility function* u_i mapping sets of resources (subsets of $\mathcal{R}$) to real numbers. We abbreviate $u_i(A) = u_i(A(i))$ for the utility value assigned by agent i to the set of resources it holds for allocation A.

An agent may or may not find a particular deal acceptable. In this paper, we assume that agents are *rational* in the sense of never accepting a deal that would not improve their personal welfare (see Sandholm (1998) for a justification of this approach). For deals with money, this "myopic" notion of individual rationality (IR) may be formalised as follows:

Definition 2.1 (Individual rationality). A deal $\delta = (A, A')$ with money 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 $A(i) = A'(i)$.

The notion of rationality (IR) provides a *local* criterion that ensures that negotiation is beneficial for all individual participants. For a *global* perspective, welfare economics (see e.g. Moulin (1988b)) provides tools to analyse how the reallocation of resources affects the well-being of a society of agents as a whole. Here we are going to be particularly interested in maximising *social welfare:*

Definition 2.2 (Social welfare). The social welfare $sw(A)$ of an allocation of resources A is defined as follows:

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

We should stress that this is the *utilitarian* view of social welfare; other notions of social welfare have been developed as well (see Moulin (1988b); Sen (1970)) and may be usefully exploited in the context of multiagent systems (Endriss et al. (2003b)).

Before we move on to discuss previous results for this framework, we should stress that, while the most widely studied mechanisms for the reallocation of resources in multiagent systems are *auctions*, our scenario of resource allocation by negotiation is *not* an auction. Auctions are centralised mechanisms to help agents agree on a price at which an item (or a set of items) is to be sold (Kersten et al. (2000)). In our work, on the other hand, we are not concerned with this aspect of negotiation, but only with the patterns of resource exchanges that agents actually carry out in a truly distributed manner.

2.3 Convergence Results

We recall in this section a fondamental lemma, as well as the main convergence result of the framework (Endriss et al. (2003a)), which is essentially equivalent to a result on sufficient contract types for optimal task allocations by Sandholm (1998).

The following lemma makes explicit the connection between the local decisions of agents and the global behaviour of the system.

Lemma 2.3 (Individual rationality and social welfare). *A deal* $\delta = (A, A')$ *is rational iff* $sw(A) < sw(A')$.

Theorem 2.4 (Maximising social welfare). *Any sequence of rational deals with money will eventually result in an allocation of resources with maximal social welfare.*

This means that (1) there can be no infinite sequence of deals all of which are rational, and (2) once no more rational deals are possible the agent society must have reached an allocation that has maximal social welfare. The crucial aspect of Theorem 2.4 is that *any* sequence of deals satisfying the rationality condition will cause the system to converge to an optimal allocation. That is, whatever deals are agreed on in the early stages of the negotiation, the system will never get stuck in a local optimum and finding an optimal allocation remains an option throughout.

A drawback of the general frameworks, to which Theorem 2.4 applies, is that these results only hold if deals involving any number of resources and agents are admissible (Endriss et al. (2003a); Sandholm (1998)). In some cases this problem can be alleviated by putting suitable restrictions on the utility functions agents may use to model their preferences.

3 One-resource-at-a-time negotiation

While Theorem 2.4 shows that, it is always possible to negotiate an allocation of resources that is optimal from a social point of view, deals involving any number of agents and resources may be required to do so (Sandholm (1998); Endriss et al. (2003a)). We now introduce a restriction called one-resource-at-a-time deals (or 1-deals).

Definition 3.1 (1-deals). A 1-deal is a deal $\delta = (A, A')$ resulting in the displacement of exactly one resource.

This basic type of deal, implemented in most systems realising a kind of *Contract Net* protocol (Smith (1980)), is certainly *not* sufficient for negotiation between agents that are not only rational but also myopic.

This has first been shown by Sandholm (1998) and is best explained by means of an example. Let $\mathcal{A} = \{1, 2, 3\}$ and $\mathcal{R} = \{r_1, r_2, r_3\}$. Suppose the utility functions of these agents are defined as follows (over singleton sets):

$$\begin{array}{lll} u_1(\{r_1\}) = 5 & u_1(\{r_2\}) = 1 & u_1(\{r_3\}) = 0 \\ u_2(\{r_1\}) = 0 & u_2(\{r_2\}) = 5 & u_2(\{r_3\}) = 1 \\ u_3(\{r_1\}) = 1 & u_3(\{r_2\}) = 0 & u_3(\{r_3\}) = 5 \end{array}$$

Furthermore, for any bundle R not listed above, suppose $u_i(R) = 0$ for all $i \in \mathcal{A}$. Let A with $A(1) = \{r_2\}$, $A(2) = \{r_3\}$ and $A(3) = \{r_1\}$ be the initial allocation, *i.e.* $sw(A) = 3$. The optimal allocation would be A' with $A'(1) = \{r_1\}$, $A'(2) = \{r_2\}$ and $A'(3) = \{r_3\}$, which yields a social welfare of 15. All other allocations have lower social welfare than A'. Hence, starting from A, the deal $\delta = (A, A')$ would be the only deal increasing social welfare. By Lemma 2.3, δ would also be the only rational deal. This deal, however, involves all three resources and affects all three agents. In particular, δ is not a 1-deal. Hence, if we choose to restrict ourselves to *rational* deals, then 1-deals are not sufficient to negotiate allocations of resources with maximal social welfare.

Of course, for some particular negotiation problems, rational 1-deals *will* be sufficient. The difficulty lies in recognising the problems where this is so. Closely related to this issue, Dunne et al. (2005) have shown that, given two allocations A and A' with $sw(A) < sw(A')$, the problem of checking whether it is possible to reach A' from A by means of a sequence of rational 1-deals is NP-hard, as long as the utilities are represented in a compact way.

The structural complexity of deals required to be able to guarantee socially optimal outcomes partly stems from the generality of the framework. In particular, so far we have made no assumptions on the structure of utility functions used by the agents to model their preferences. By introducing restrictions on the class of admissible utility functions, it may be possible to ensure convergence to an allocation with maximal social welfare by means of simpler deals. In the rest of the section, we will be interested in characterising more precisely those classes of utility functions that permit 1-deal negotiation.

Definition 3.2 (1-deal negotiation). A class $\mathcal{C}$ of utility functions is said to permits 1-deal negotiation iff any sequence of rational 1-deals will eventually result in an allocation of resources with maximal social welfare whenever all utility functions $\{u_1, \ldots, u_n\}$ are drawn from $\mathcal{C}$.

Under this perspective, a relevant result is due to Endriss et al. (2003a), who show that rational 1-deals are sufficient to guarantee outcomes with maximal social welfare in case all agents use *additive* utility functions.[1] We are going to prove a slight generalisation of this result in the next section.

3.1 Modular functions are sufficient

We are now going to define the class of *modular* utility functions. This is an important (see e.g. Rosenschein and Zlotkin (1994b)), albeit simple, class of functions that can be used in negotiation domains where there are no synergies (neither complementaries nor substitutables) between different resources.

Definition 3.3 (Modular utility). A utility function u is modular iff the following holds for all bundles $R_1, R_2 \subseteq \mathcal{R}$:

$$u(R_1 \cup R_2) \quad = \quad u(R_1) + u(R_2) - u(R_1 \cap R_2) \tag{3.1}$$

The class of modular functions includes the aforementioned additive functions. This may be seen as follows. Let R be any non-empty bundle of resources and let $r \in R$. Then equation (3.1) implies $u(R) = u(R \setminus \{r\}) + [u(\{r\}) - u(\{\,\})]$. If we apply this step recursively for every resource in R, then we end up with the following equation:

$$u(R) \quad = \quad u(\{\,\}) + \sum_{r \in R} [u(\{r\}) - u(\{\,\})] \tag{3.2}$$

That is, in case $u(\{\,\}) = 0$, the utility assigned to a set will be the sum of utilities assigned to its members (*i.e.* u will be additive). Clearly, equation (3.2) also implies equation (3.1), *i.e.* the two characterisations of the class of modular utility functions are equivalent.

It turns out that in domains where all utility functions are modular, it is always possible to reach a socially optimal allocation by means of a sequence of rational deals involving only a single resource each. This result, shown in (Chevaleyre et al. (2005b)), is a slight generalisation of a result proved by Endriss et al. (2003a), and both proof are close from each other.

Theorem 3.4 (Negotiation in modular domains). *The class $\mathcal{M}$ of modular utility functions permits 1-deal negotiation.*

Proof. By Lemma 2.3, any rational deal results in a strict increase in social welfare. Together with the fact that the number of distinct allocations is finite, this ensures that there can be no infinite sequence of rational deals (termination). It therefore suffices to show that for any allocation that does not have maximal social welfare there still exists a rational 1-deal that would be applicable.

We are going to use the alternative characterisation of modular utility functions given by equation (3.2). For any allocation A, let f_A be the function mapping each resource r to the

[1] A utility function is additive iff the utility assigned to a set of resources is always the sum of utilities assigned to its members.

agent i that holds r in situation A. Then, for modular domains, the formula for social welfare (see Definition 2.2) can be rewritten as follows:

$$sw(A) \quad = \quad \sum_{i \in \mathcal{A}} u_i(\{\,\}) + \sum_{r \in \mathcal{R}} u'_{f_A(r)}(\{r\})$$

with $u'_i(R) = u_i(R) - u_i(\{\,\})$. Now assume we have reached an allocation of resources A that does not have maximal social welfare, *i.e.* there exists another allocation A' with $sw(A) < sw(A')$. Considering the above definition of social welfare and observing that $\sum_{i \in \mathcal{A}} u_i(\{\,\})$ is a constant that is independent of the current allocation, this implies that at least one resource r must satisfy the inequation $u'_{f_A(r)}(\{r\}) < u'_{f_{A'}(r)}(\{r\})$, *i.e.* the agent owning r in allocation A values that resource less than the agent owning it in allocation A'. But then the 1-deal consisting of passing r from agent $f_A(r)$ to agent $f_{A'(r)}$ would already increase social welfare and thereby be rational. □

Like Theorem 2.4, the above establishes an important convergence result towards a global optimum by means of decentralised negotiation between self-interested agents. In addition, provided all utility functions are modular, convergence can be guaranteed by means of a much simpler negotiation protocol, which only needs to cater for agreements on 1-deals (rather than multilateral deals over sets of resources).

3.2 Modular functions are not necessary

In the previous section we have introduced a class of utility functions (namely modular functions) such that ensuring that all utilities belong to it is *sufficient* for rational 1-deals to converge to an allocation with maximal social welfare. A natural question to ask would then be whether modularity is also a *necessary* condition.

As shown in (Chevaleyre et al. (2005b)), it turns out that this is not the case. We demonstrate this by means of the following example. Suppose $\mathcal{R} = \{r_1, r_2\}$ and there are two agents with utility functions u_1 and u_2:

$$\begin{array}{lcllcl}
u_1(\{\,\}) & = & 90 & u_2(\{\,\}) & = & 90 \\
u_1(\{r_1\}) & = & 93 & u_2(\{r_1\}) & = & 90 \\
u_1(\{r_2\}) & = & 95 & u_2(\{r_2\}) & = & 90 \\
u_1(\{r_1, r_2\}) & = & 98 & u_2(\{r_1, r_2\}) & = & 50
\end{array}$$

While u_1 is a modular function, 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 modular.

In fact, it is possible to show that there can be no class of utility functions that would be both sufficient and necessary in this sense. It suffices to produce two concrete utility functions u_1 and u_2 such that (i) both of them would guarantee convergence if all agents were using them, and (ii) there is a scenario where some agents are using u_1 and others u_2 and convergence is not guaranteed. This is so, because assuming that a necessary and sufficient class exists, (i) would imply that both u_1 and u_2 belong to that class, while (ii) would entail the contrary. We give two such functions for the case of two agents and two resources (the argument is easily augmented to the general case):

$$\begin{array}{lclllcl} u_1(\{\,\}) & = & 0 & & u_2(\{\,\}) & = & 0 \\ u_1(\{r_1\}) & = & 1 & & u_2(\{r_1\}) & = & 5 \\ u_1(\{r_2\}) & = & 2 & & u_2(\{r_2\}) & = & 5 \\ u_1(\{r_1, r_2\}) & = & 3 & & u_2(\{r_1, r_2\}) & = & 5 \end{array}$$

The function u_1 is modular, *i.e.* all agents using that function is a sufficient condition for guaranteed convergence to an optimal allocation by means of rational 1-deals (Theorem 2.4). Clearly, convergence is also guaranteed if all agents are using u_2. However, if the first agent uses u_1 and the second u_2, then the allocation A with $A(1) = \{r_1\}$ and $A(2) = \{r_2\}$ is not socially optimal and the only deal increasing social welfare (and thereby, the only rational deal) would be to swap the two resources simultaneously. Hence, no condition on utility functions that is sufficient to guarantee convergence to an optimal allocation by means of rational 1-deals alone is also a necessary condition for this to be the case.

Our argument for the inexistence of any such necessary and sufficient condition has directly exploited the fact that we were looking for a *single* condition to be met by the utility functions of *all* agents. The problem could be circumvented by looking for suitable conditions on negotiation problems as a whole, where different utility functions may belong to different classes. Clearly, such a condition *does* exist. However, the aforementioned result of Dunne et al. (2005) on the NP-hardness of checking whether there exists a path of rational 1-deals between two given allocations immediately suggests that verifying whether a given negotiation problem meets any such condition would be highly intractable.

3.3 The modular class is maximal

We are now going to show the main result of this section, namely the surprising fact that the class of modular utility functions is not only sufficient for 1-deal negotiation but also *maximal* in the sense that no class of utility functions strictly including the modular functions would still be sufficient for 1-deal negotiation. The significance of this result which was also shown by Chevaleyre et al. (2005b) can only be fully appreciated when considered together with the "negative" result on necessary and sufficient conditions discussed in the previous section.

Before stating the main result, we prove the following auxiliary lemma:

Lemma 3.5 (Alternative characterisation of modularity). *A utility function u is modular iff the following holds for all $R \subseteq \mathcal{R}$ and all $r_1, r_2 \in \mathcal{R}$ with $r_1, r_2 \notin R$ and $r_1 \neq r_2$:*

$$u(R \cup \{r_1, r_2\}) = u(R \cup \{r_1\}) + u(R \cup \{r_2\}) - u(R) \tag{3.3}$$

Proof. To show this, let us recall elementary facts about submodular functions. A function $v : \mathcal{R} \to \mathbb{R}$ is submodular iff $\forall R_1, R_2 \subseteq \mathcal{R},\ v(R_1) + v(R_2) \geq v(R_1 \cup R_2) + v(R_1 \cap R_2)$. It is also known that v is submodular iff $v(R \cup \{r_1\}) + v(R \cup \{r_2\}) \geq v(R \cup \{r_1, r_2\}) - v(R)$ for any $R \subseteq \mathcal{R},\ r_1, r_2 \in \mathcal{R} \backslash R$, with $r_1 \neq r_2$ (Nemhauser and Wolsey, 1988, p.662). Because a function u is modular iff both u and $-u$ are submodular, the lemma holds. □

We are now in a position to show a theorem on the maximality of the class of modular utility functions with respect to rational negotiation over one resource at a time:

Theorem 3.6 (Maximality). *Let $\mathcal{M}$ be the class of modular utility functions. Then for any class of utility functions $\mathcal{F}$ such that $\mathcal{M} \subset \mathcal{F}$, $\mathcal{F}$ does not permit 1-deal negotiation.*

The proof, which can be found in Chevaleyre et al. (2005b), is constructive. It consists in showing that for any non-modular utility function u_1 on m resources, it is possible to construct a modular utility function u_2 (with $u_i \equiv 0$ for all other agents i) and an initial allocation such that no optimal allocation can be reached by means of 1-deals. This implies that $\mathcal{M} \cup \{u_1\}$ does not permit 1-deals.

Why is this result significant? As argued earlier, while the full abstract negotiation framework introduced at the beginning of this paper would be difficult to implement, designing a system that only allows for pairs of agents to agree on deals over one resource at a time is entirely feasible. As we would like to be able to guarantee socially optimal outcomes in as many cases as possible, also for such a restricted negotiation system, we would like to be able to identify the largest possible class of utility functions for which such a guarantee can be given. However, our discussion in Section 3.2 has shown that there can be no class of utility functions that *exactly* characterises the class of negotiation problems for which negotiating socially optimal allocations by means of rational 1-deals is always possible. Still, there *are* classes of utility functions that permit 1-deal negotiation. As shown by Theorem 3.4, the class of modular functions is such a class and it is a very natural class to consider. An obvious question to ask is therefore whether this class can be enlarged in any way without losing the desired convergence property.

Theorem 3.6 settles this question by giving a negative answer: For any agent with a non-modular utility function there exist modular utility functions (for the other agents) and an initial allocation such that rational 1-deals alone do not suffice to negotiate an allocation of resources with maximal social welfare. There may well be further such classes (that are both sufficient and maximal), but the class of modular functions is likely to be the one that is the most natural and the most useful for modelling agent preferences in practice.

4 Many resources-at-a-time negotiation

We will now study under which conditions reaching the optimal allocation can be guaranteed, when at most k resources are exchanged at a time. As discussed after theorem 2.4, at worst, reaching the optimal allocation requires deals involving all resources of $\mathcal{R}$. Thus, we again need to restrain utility functions to reach the optimal allocation.

In this section, we present the class of *additively k-separable* utility functions, introduced in (Chevaleyre et al. (2005a)), 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.

Definition 4.1 (Additive k-separability). Let $P = \langle R_1, \ldots, R_q \rangle$ 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 $R \subseteq \mathcal{R}$:

$$u(R) \;=\; u(\{\,\}) + \sum_{j=1}^{q} [u(R \cap R_j) - u(\{\,\})] \tag{4.1}$$

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 (4.1) 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.

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 $\langle R_1, \ldots, R_q \rangle$. Now define a *local* utility function $u^j : 2^{R_j} \to \mathbb{R}$ 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).

4.1 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 4.2 (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.*

The proof, which can be found in Chevaleyre et al. (2005a), uses the same trick as the proof of theorem 3.4 of the present paper.

What are the crucial properties of additively k-separable domains that bring about the convergence result of Theorem 4.2? Firstly, in *separable* domains, the marginal utility gains associated with the resources belonging to different topics are entirely independent. Therefore, we could choose to alway 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.

5 Tractable negotiation with k-deals

It has been shown the previous section that with k-separable domains, k-deals were sufficient to garantee us to reach the optimal social welfare. However, finding individually rational k-deals is intractable even with reasonable values of k, because of the high number of possible k-deals. In this section, we introduce a further restriction of k-separable utilities. This restriction, denoted " k-additive tree-structured utilities", is a natural restriction for many application, which will allow us to reduce drastically the search space without loosing the convergence properties.

5.1 k-additive utilities

The succinctness of the k-separable 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 (Grabisch (1997)). A utility function is called *k-additive* iff the utility assigned to a bundle R can be represented as the sum of basic utilities to subsets R with cardinality $\leq k$. More formally, a k-additive utility function can be written as follows:

$$u_i(R) \quad = \sum_{T \subseteq \mathcal{R},\, |T| \leq k} \alpha_i^T \times I_R(T) \quad \text{with } I_R(T) = \begin{cases} 1 & \text{if } T \subseteq R \\ 0 & \text{otherwise} \end{cases}$$

That is, the utility function of agent i is characterised by the coefficients α_i^T for bundles of resources $T \subseteq \mathcal{R}$ with at most k elements. Agent i enjoys an increase in utility of α_i^T when it owns all the items in T *together*, *i.e.* α_T^i represents the synergetic value of this bundle. An example for a 2-additive utility function would be $u_i(R) = 3 \times I_R(\{r_1\}) - 2 \times I_R(\{r_2, r_3\})$. For the sake of simplicity, we are going to omit the indicator function I_R as well as the explicit mentioning of the bundle variable R when defining concrete k-additive utility functions. Using this simplified notation, the above function becomes $u_i = 3.r_1 - 2.r_2.r_3$. In this equation, r_1 and $r_2.r_3$ are called *terms*.

Agent i enjoys an utility of α_i^T when owning resources of term T *together* (that is, α_i^T represents the synergetic value of this resources held together). When then value of this term is positive, resources are said to be *complementary*, when it is negative, they are *substituable*.

In the field of *combinatorial auctions*, several other bidding languages have also been introduced and studied (see Nisan (2000)). There are links with the representation discussed above. The details are of these connections are currently under investigation.

5.2 Tree-structured Domains

To formulate this new restriction on preferences, utilities are needed to be represented in k-additive form. Thus, from now on, all utilities will be represented this way. Let $\mathcal{R}$ be the set of resources $r_1 ... r_m$, and $u_1 \dots u_n$ a set of utility funtions. $\mathcal{T}$ will denote the set of terms explicitely appearing in utility functions $u_1 \dots u_n$, and α_t^i will denote the coefficient of term t in u_i. Finally, $\mathcal{T}^l$,$\mathcal{T}^{\leq l}$ denote the set of terms in $\mathcal{T}$ consisting of respectively exactly l resources and at most l resources.

5.3 Tree-Structured Utilities

Intuitively, tree-structured k-additive utilities are functions in which there are no overlapping terms[2].

Definition 5.1. A set of utility functions $u_1 ... u_n$ is *tree-structured* iff utilities can be written in k-additive form, and $\forall T_1, T_2 \in \mathcal{T}$, we have either $T_1 \subset T_2$ or $T_1 \supset T_2$ or $T_1 \cap T_2 = \{\}$.

Notice that a set of utility functions is tree-structured iff the terms of $\mathcal{T}$ can be represented by a tree, in which $\mathcal{R}$ is the root, and each term is a node. Branches of the tree represent the $\subset$ relation. The following example illustrates this representation.

It is easy to show that if a utility function u is tree-structured with positive coefficients, then it is super-additive. However, the converse is not true.

More importantly, *k-additive tree-structured utilities are clearly k-separable as well*. Thus, it is garanteed that using IR k-deals with tree-structured k-additive utilities, the negotiation process will eventually converge towards optimal social welfare.

Example 5.2. Consider the three functions $u_1 = r_2 + 3.r_5$, $u_2 = 3.r_1 + 10.r_1.r_2.r_3.r_4 + 8.r_5 + 4.r_6$ and $u_3 = r_6 - r_4 + 8.r_3 r_4$. Clearly, they are 4-additive as well as tree-structured. Here, the set of terms $\mathcal{T} = \{\{r_1\}, \{r_2\}, \{r_4\}, \{r_5\}, \{r_6\}, \{r_3, r_4\}, \{r_1, r_2, r_3, r_4\}\}$ can also be reprensented by the following tree:

```
          R
         /|\
        / | \
       /  |  \
  {r5}  {r6}{r1r2r3r4}
              / | \
             /  |  \
        {r1}  {r2}{r3r4}
                    |
                  {r4}
```

Because finding IR k-deals is in general intractable, we will search for other types of deals less complex but insuring the same convergence properties.

Definition 5.3. Let $\mathcal{T}$ be a set of terms. A $\mathcal{T}$-deal is a deal with side-payment involving an entire term of $\mathcal{T}$ from one or more sender(s) to a *single* receiver.

[2] for e.g. the following utility contains two overlapping terms: $r_1 r_2 + r_2 r_3$

First, note that the number of possible $\mathcal{T}$-deals is low ($n \times |\mathcal{T}|$). Thus, the complexity of finding IR $\mathcal{T}$-deals is also very low compared to that of k-deals. However, as illustrated by the example below, simply allowing any $\mathcal{T}$-deals will not be sufficient to garantee us that the optimal social welfare will always eventually be reached.

Example 5.4. Let $u_1 = 10.r1$,$u_2 = 10.r_2$, and $u_3 = 11.r_1 r_2$, and let the initial allocation be the allocation assigning all resources to agent 3. Clearly, there are no IR $\mathcal{T}$-deal reaching optimal allocation.

5.4 Negotiation Protocol

Designing a centralized algorithm to share resources optimaly in this setting would be quite easy. However, we wish to use a decentralized algorithm where all agents are individually rational.

Thus, to allow agents to reach optimal social welfare, we need to design an ad hoc protocol. In this protocol, to prevent agents from getting stuck in local minimas, a special agent (called the Bank) will participate to the side-payments, by adding or taking money from the negotiating agents, thus modifying their rationality. Of course, this protocol will be tractable (polynomial time).

This protocol is called *Omniscient ϵ-Altruistic Tree-Climbing Protocol*, because the system (the bank agent) actively mediates the negotiation process. More precisely, this protocol is called *Omniscient* because the system needs to know precisely the utility functions of each agent, ϵ-Altruistic because the system must be prepared to share its money with agents in order to reach optimal social welfare, and that the amount of money it shares depends on ϵ.

In the following, $sw^{\leq l}$ denotes the social welfare computed over k-additive utilities from which all terms of more than l resources have been removed.

[h!] Omniscient ϵ-Altruistic Tree-Climbing Protocol

Require: n agents with utilities $u_1 ... u_n$ identically tree-structured. A parameter $\epsilon \in [0, 1]$. Let $\mathcal{T}$ the set of all terms.

All agents transmit to the *bank agent* their utility function.
$l \leftarrow 1$.
repeat
    Restrict allowed deals to $\mathcal{T}^l$-deals.
    For each $\mathcal{T}^l$-deal $\delta = (A, A')$, let the associated side-payment be defined as follows:
    For all the n' agents involved in the deal,
    $p_i = u_i(A') - u_i(A) - \frac{\epsilon}{n'} \left(sw^{\leq l}(A') - sw^{\leq l}(A) \right)$
    For the system (the bank agent), $p_{bank} = - \sum p_i$
    Let agents do all their deals in a IR with side payments way.
    if no more IR deal can be conducted **then**
        $l \leftarrow l + 1$
    end if
until $l > | \mathcal{R} |$

Note that in the Oϵ-ATCP protocol, the *bank agent* sometimes gives or receives money from agents. The amount of money shared with other agents depends on the value of the parameter ϵ.

If ϵ is close to one, then the bank agent in not garanteed to globally earn money. In that case, the bank agent is said to be *altruistic*, because negotiating agents will benefit from this by earning a lot.

If ϵ is close to 0, during each deal, the bank agent takes from other agents as much money as possible. Thus, the negotiating agents will earn the minimum amount of money necessary to reach the optimal allocation, and the bank agent will earn a lot more (not altruistic at all in that case).

For the sake of fairness, this protocol could be extended in such a way that the money earned by the bank agent would be equally redistributed at the end of the process. The following results show that this protocol converges as well.

Theorem 5.5. *With the Omniscient ϵ-Altruistic Tree-Climbing protocol, the optimal social welfare is eventually reached if $\epsilon \in [0, 1]$. In addition, the bank agent will globally earn money during the entire negotiation is ϵ is sufficiently small.*

The proof can be found in (Chevaleyre et al. (2005c)), as well as arguments showing that this protocol is tractable (quadratic). However, a strong drawback of this protocol is the need for the system to know the utilities of each agents. As is, this makes it quite unrealistic. This will be discussed in opened questions.

6 Conclusion and Open Questions

This paper reviews some recent results concerning theoretical analysis of an abstract negotiation framework. Its main results are a negative result for one-resource-at-a-time negotiation (maximality) and two positive results (convergence for separable domains, and tractable convergence for tree structured utilities).

The negative result has a number of practical implications. From a designer perspective, it means for instance that a negotiation platform designed for modular domains cannot accept an agent not implementing a modular utility function without jeopardising the outcome of the process. Note however that it does *not* imply that a system containing non modular agents would never guarantee an optimal outcome.

The positive results, which can be viewed as a generalization of Theorem 2.4 show that rational deals involving at most k deals are sufficient for convergence to an optimal allocation in case all utility functions are *additively separable* with respect to a common partition of $\mathbb{R}$ (*i.e.* synergies across different parts of the partition are not possible and overall utility is defined as the sum of utilities for the different sets in the partition (Fishburn (1970))), and each set in this partition has at most k elements. In addition, it will also be tractable when utilities are tree-structured.

There are many important opened questions yielded by these results.

- **Maximality for k-deals**. The maximality property is a strong negative result for 1-resource at a time negotiation. Can this result be extended to k-resource at a time negotiation ? More specifically, for a given partition, is the corresponding k-separable domain maximal, in the same sense as in section 3 ? Our intuition is that it is not.

- **Negotiating without money**. All the results presented here concern negotiation with side-payments. What is money is not allowed ? A preliminary result has been shown for negotiating with k resources at a time in Chevaleyre et al. (2005a), but could probably be widely improved.
- **Other social welfares**. All the results shown here deal with the utilitarian social welfare. Can similar results be derived for other measures, such as egalitarian social welfare, pareto optimality, etc ?
- **Complexity issues**. In this paper, two compact representations of utility functions have been used. How complex are usual algorithmic problems (e.g. finding optimum) using these representations ?
- **Nominal vs Cardinal preferences**. In our setting, all preferences are represented with cardinal utility functions. What if we choose to use ordinal preferences ? In this new setting, side-payments would no longer be usable. What kind of result could we get ?
- **Truthfull tractable negotiation**. In the last section, the protocol which allows efficient negotiation requires agents to reveal their preferences. As agents are selfish, they may lie to increase their profit. Mechanism design theory studies how to build truthful protocols, in which telling the truth becomes profitable for agents. An important opened question would be how to make the preference revelation truthful (or approximately truthful), if possible.

Bibliography

Kenneth J. Arrow, Amartya K. Sen, and Kotaro Suzumura, editors. *Handbook of Social Choice and Welfare*, volume 1. North-Holland, 2002.

Steven J. Brams and Alan D. Taylor. *Fair Division: From Cake-cutting to Dispute Resolution*. Cambridge University Press, 1996.

Anthony Chavez, Alexandros Moukas, and Pattie Maes. Challenger: A multi-agent system for distributed resource allocation. In *Proc. of the 1st Intl. Conf. on Autonomous Agents (Agents-1997)*. ACM Press, 1997.

Yann Chevaleyre, Ulle Endriss, Sylvia Estivie, and Nicolas Maudet. Multiagent resource allocation with k-additive utility functions. In *Proc. of the DIMACS-LAMSADE Workshop on Computer Science and Decision Theory*, 2004.

Yann Chevaleyre, Ulle Endriss, Jérôme Lang, and Nicolas Maudet. Negotiating over small bundles of resources. In *Proceedings of the 4th International Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS-2005)*, July 2005a. To appear.

Yann Chevaleyre, Ulle Endriss, and Nicolas Maudet. On maximal classes of utility functions for efficient one-to-one negotiation. In *Proceedings of the 19th International Joint Conference on Artificial Intelligence (IJCAI-2005)*, August 2005b. To appear.

Yann Chevaleyre, Ulle Endriss, and Nicolas Maudet. Protocols for tractable resource allocation with k-additive utilities. In A. Herzig and Y. Lespérance, editors, *Troisièmes Journées Francophones sur les Modèles Formels d'Interaction (MFI-2005)*. Cépaduès-Éditions, May 2005c. To appear.

Paul E. Dunne, Michael Wooldridge, and Michael Laurence. The complexity of contract negotiation. *Artificial Intelligence*, 2005. To appear.

Ulrich Endriss, Nicolas Maudet, Fariba Sadri, and Francesca Toni. On optimal outcomes of negotiations over resources. In *Proceedings of the 2nd International Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS-2003)*. ACM Press, 2003a.

Ulrich Endriss, Nicolas Maudet, Fariba Sadri, and Francesca Toni. Resource allocation in egalitarian agent societies. In A. Herzig, B. Chaib-draa, and Ph. Mathieu, editors, *Secondes Journées Francophones sur les Modèles Formels d'Interaction (MFI-2003)*, pages 101–110. Cépaduès-Éditions, May 2003b.

P. C. Fishburn. *Utility Theory for Decision Making*. John Wiley and Sons, 1970.

Michel Grabisch. k-order additive discrete fuzzy measures and their representation. *Fuzzy Sets and Systems*, 92:167–189, 1997.

Gregory E. Kersten, Sunil J. Noronha, and Jeffrey Teich. Are all e-commerce negotiations auctions? In *Proc. of the 4th Intl. Conf. on the Design of Cooperative Systems*, 2000.

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

H. Moulin. *Axioms of Cooperative Decision Making*. Cambridge University Press, 1988a.

Hervé Moulin. *Axioms of Cooperative Decision Making*. Cambridge University Press, 1988b.

G. L. Nemhauser and L.A. Wolsey. *Integer and Combinatorial Optimization*. John Wiley and Sons, 1988.

Noam Nisan. Bidding and allocation in combinatorial auctions. In *Proceedings of the ACM Conference on Electronic Commerce (EC-2000)*, 2000.

J. S. Rosenschein and G. Zlotkin. *Rules of Encounter*. MIT Press, 1994a.

Jeffrey S. Rosenschein and Gilad Zlotkin. *Rules of Encounter*. MIT Press, 1994b.

Tuomas W. Sandholm. Contract types for satisficing task allocation: I Theoretical results. In *Proceedings of the AAAI Spring Symposium: Satisficing Models*, 1998.

Tuomas W. Sandholm. Distributed rational decision making. In *Multiagent Systems: A Modern Approach to Distributed Artificial Intelligence*. MIT Press, 1999.

Amartya K. Sen. *Collective Choice and Social Welfare*. Holden Day, 1970.

R. G. Smith. The contract net protocol: High-level communication and control in a distributed problem solver. *IEEE Transactions on Computers*, C-29(12):1104–1113, 1980.