Protocols for tractable resource allocation with k-additive utilities

Yann Chevaleyre[⋆] chevaley@lamsade.dauphine.fr
Ulle Endriss[†] ue@doc.ic.ac.uk
Nicolas Maudet[⋆] maudet@lamsade.dauphine.fr

[⋆]LAMSADE
Université Paris-Dauphine
Paris, FRANCE

[†]Department of Computing, Imperial College London
180 Queen's Gate, London SW7 2AZ, UK

Résumé :
Cet article aborde l'allocation de ressources multiagent par la négociation. Un des problèmes majeurs de cette approche, qui rend difficile la mise en pratique des résultats théoriques, est qu'il s'avère très complexe d'identifier les échanges potentiellement acceptables pour un ensemble donné d'agents. La solution que nous envisageons ici est d'utiliser différents protocoles conçus afin d'exploiter certaines propriétés des fonctions d'utilités utilisées par les agents pour modéliser leurs préférences. Nous considérons spécifiquement les domaines où les fonctions d'utilité sont k-additives (c'est-à-dire que les synergies entre les ressources sont restreintes aux lots d'au plus k ressources), et structurées sous forme d'arbre, au sens où les lots pour lesquels il existe une synergie ne se "chevauchent" pas.

Mots-clés : Négociation, Allocation de ressources

Abstract:
Negotiation over resources in multiagent systems is a timely and fruitful area of ongoing research. However, the prohibitively high complexity of the task of identifying rational deals, *i.e.* deals that are beneficial for all participants, currently hinders the successful transfer of theoretical results to practical applications. To address this issue, we propose several protocols designed to tame the complexity of negotiation by exploting structurual properties of the utility functions used by agents to model their preferences over alternative bundles of resources. In particular, we consider domains where utility functions are k-additive (that is, synergies between different resources are restricted to bundles of at most k items) and "tree-structured" in the sense that the bundles for which there are synergies do not overlap. We show how protocols exploiting these properties can enable drastically simplified negotiation processes.

Keywords: Negotiation, Resource Allocation

1 Introduction

Negotiation in general, and the allocation of resources by means of negotiation in particular, are widely regarded as important topics in multiagent systems research. In this paper, we study a multilateral negotiation framework where autonomous agents agree on a sequence of deals to exchange sets of discrete (i.e. non-divisible) resources. While, at the local level, agents arrange deals to further their own individual goals, at the global level (say, from a system designer's point of view) we are interested in negotiation processes that lead to allocations of resources that are *socially optimal*. In this paper, we are only concerned with maximising *utilitarian* [12, 15] social welfare (this concept will be defined in Section 2).

Previous work has addressed the emergence of states that are optimal from a social point of view, depending on the kinds of acceptability criteria used by individual agents when deciding whether or not to agree to a proposed exchange of resources [7, 14]. A first analysis of the complexity of certain aspects of this framework has recently been given by Dunne *et al.* [5]. In a different perspective, a certain kind of *communication complexity* of this framework (analysing, for instance, the length of negotiation processes) has also been investigated recently [6]. However, the prohibitively high complexity of the task of identifying rational deals, *i.e.* deals that are beneficial for all participants, currently hinders the successful transfer of theoretical results to practical applications. To address this issue, we propose several protocols designed to tame the complexity of negotiation by exploting structurual properties of the utility functions used by agents to model their preferences over alternative bundles of resources. We propose two protocols that restrict negotiation to deals involving only the smallest bundles first, and then incrementally bigger bundles. However, these guided negotiation processes are not guaranted to converge to an allocation with optimal social welfare, *i.e.* protocols need to be enhanced further. The first proposed protocol is made *intrusive* in the sense that it modifies the structure of agents' utility functions. The second protocol has a centralised flavour and introduces a system agent that can *compensate* for a temporary loss in social welfare. This combination of *guidance* and *social compensation* allows to restore the convergence properties of the framework.

The remainder of this paper is structured as follows. In Section 2 we review the multilateral trading framework of [7]. Section 3 recalls different aspects of the complexity of trading resources due to [6]. A compact representation of preferences based on the notion of k-additivity (where synergies between different resources are restricted to bundles of at most k items) is introduced in Section 4. Section 5 introduces a further structural constraint on utility functions (namely, the fact that they are *tree-structured* in the sense that the bundles for which there are synergies do not overlap). Section 6 presents a number of protocols exploiting these structural properties of utility functions, and shows how they drastically reduce the complexity of the negotiation process. Section 7 concludes with a discussion of the benefits and drawbacks of the approach advocated in this paper.

2 Resource Allocation by Negotiation

In this section, we introduce the framework of *resource allocation by negotiation* put forward in [7] 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 [14] for a justification of this approach). For deals with money, this "myopic" notion of individual rationality may be formalised as follows :

Definition 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 provides a *local* criterion that ensures that negotiation is beneficial for all individual participants. For a *global* perspective, welfare economics (see e.g. [12]) 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 (Social welfare) *The social welfare $sw(A)$ of an allocation of resources A is defined as follows :*

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

We should stress that this is the *utilitarian* view of social welfare ; other notions of social welfare have been developed as well [12, 15] and may be usefully exploited in the context of multiagent systems [8].

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 [11]. 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 the main convergence result of the framework [7], which is essentially equivalent to a result on sufficient contract types for optimal task allocations by Sandholm [14]. This result links individual rationality at the local level with the global concept of social welfare :

Theorem 1 (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 1 (and the next three theorems) 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 1 applies, is that these results only hold if deals involving any number of resources and agents are admissible [7, 14]. In some cases this problem can be alleviated by putting suitable restrictions on the utility functions agents may use to model their preferences. For instance, a utility function is called *additive* iff the value ascribed to a set of resources is always the sum of the values of its members. In scenarios where utility functions may be assumed to be additive, it is possible to guarantee optimal outcomes even when agents only negotiate deals involving a single resource and a pair of agents at a time (so-called *one-resource-at-a-time deals*). More generally, one may be interested in k-*deals*. A k-deal is a deal involving at most k resources, *i.e.* a deal $\delta = (A, A')$ such that :

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

We are now going to inspect what different aspects of complexity should be considered in the context of a negotiation framework such as ours.

3 Aspects of Complexity

The aim of this paper is to tame the complexity of negotiation within the multiagent resource allocation framework lined out in the previous section. In fact, there are at least four different *aspects of complexity* to consider in this context. Our discussion of these closely follows [6]. The four different aspects of complexity we can identify are epitomised by the following questions :

(1) How many deals are required to reach an optimal allocation of resources ?
(2) How many dialogue moves need to be exchanged to agree on one such deal ?
(3) How expressive a communication language do we require ?
(4) How complex is the reasoning task faced by an agent when deciding on its next dialogue move ?

The first type of complexity takes individual deals as primitives, abstracting from their inherent complexity, and evaluates the length of a negotiation process as a whole. Following a *top-down* approach, this is the first aspect of complexity to consider. At the next lower level, we have to consider the complexity of negotiating a *single* deal in such a sequence. This issue is addressed by the second type of complexity identified above. It concerns the number of messages that need to be sent back and forth between the agents participating in negotiation before a deal can be agreed upon. At the next lower level, we have to consider the complexity of deciding *what* message to send at any given point in a negotiation process ; this is the fourth type of complexity. The third type is somewhat orthogonal to the other points as it concerns the complexity of a *language :* how rich a agent communication language do we require, for instance, to be able to specify proposals and counter-proposals ? The first three of the four questions at the beginning of the section relate to what we may call the *communication complexity* of our negotiation framework. The communication complexity of reaching an optimal allocation of resources is a combination of the number of deals required and the complexity of arranging an individual deal. Recall that our negotiation framework makes multilateral deals a *necessity* ; this is the price to pay for the simplicity of our agent model based on the notion of rationality. If agents only agree to deals that improve their own welfare (rather than being prepared to accept a temporary loss in utility in view of potential future rewards), then deals involving any number of agents as well as resources may be required to be able to guarantee socially optimal outcomes [7, 14]. Truly multilateral trading, i.e. negotiating deals that involve more than just two agents, however, is considerably more complex than the more widely studied bilateral trading. As pointed out by Feldman [9],

– if the costs of arranging a multilateral deal were proportional to the number of pairs in a group of agents, then they would rise *quadratically* with the size of the group (because there are $n \cdot (n-1)/2$ pairs in a group of n agents) ; and
– if the costs were proportional to the number of subgroups in a group, then they would rise *exponentially* (because there are 2^n subgroups).

These observations directly affect the second type of complexity, i.e. the number of dialogue moves that need to be exchanged to agree on a deal between several agents. Let us consider the complexity of agreeing on rational deals. In general, there may be up to $|\mathcal{A}|^{|\mathcal{R}|} - 1$ (potentially rational) deals to consider. Clearly, the problem of finding rational deals can quickly become intractable. For this reason, we shall put restrictions on the structure of utility functions used by agents to model their preferences, in order to reduce the number of deals to be considered. Before further detailing these restrictions, we introduce in the next section the k-additive form to represent utility functions.

4 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. The succinctness of the representation of agent preferences can be 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, 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 [1, 10]. 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 $\leq k$. In what follows, we shall use the following notation :

$$u_i(R) = \sum_{t \subseteq R} \alpha_t^i \prod_{r \in t} r$$

with $\alpha_t^i = 0$ whenever $|t| > k$.

Agent i enjoys a utility of α_t^i when owning *together* all the items r composing the bundle represented by the *term* t (that is, α_t^i represents the synergetic value of these items held together). When then value of α is positive, items are said to be *complementary*, when it is negative, they are *substituable*. Both representations are equivalent in term of expressive power, in the sense that they both can represent all utility functions [1]. As both representations are equivalent regarding expressiveness, one should consider whether one is stricly better in terms of the succinctness of representation. This proves not to be the case, *i.e* there are cases where one representation would be polynomial and the other exponential, and the other way around [1]. However, the k-additive form is typically more concise in cases where there are only limited synergies between different resources.

5 Tree-structured Domains

A first idea to circumvent the complexity of finding rational deals is to simply restrict the number of resources involved in each single deal.[1] However, finding individually rational k-deals is intractable even with reasonable values of k, because of the high number of possible k-deals. More precisely, considering that a deal involves at most k resources among $|\mathcal{R}|$, and that each resources can be transfered to n different agents, the number of possible deals is $n^k \times \frac{|\mathcal{R}|!}{(|\mathcal{R}|-k)!}$.

[1] In a companion paper [2], we study a sufficient condition (to be met by every agent in the system) guaranteeing socially optimal outcomes for such k-deal negotiation. This condition requires utility functions to be "additively k-separable", which is a generalisation of the "tree-structured" utilities we are going to introduce next.

In this section, we introduce a restriction on k-additive utilities. This restriction, denoted "tree-structured utilities", is a natural restriction for many applications, which will allow us to reduce drastically the search space without losing the convergence properties.

To formulate this restriction, we require utilities to be represented in k-additive form. 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 explicitly 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.

Intuitively, tree-structured k-additive utilities are functions in which there are no overlapping terms. [2] By extension, we define the notion of tree-structured sets of utility functions.

Definition 3 *A set of utility functions $u_1 ... u_n$ is* tree-structured *iff, when represented in k-additive form, it is the case that $\forall T_1, T_2 \in \mathcal{T}$, we have either $T_1 \subset T_2$ or $T_1 \supset T_2$ or $T_1 \cap T_2 = \emptyset$.*

Example 1 *Consider the set containing only the following utility function :*

$$u_1 = r_1 + r_2 r_3 + r_3 r_4$$

composed of exactly three terms. Then we have $\mathcal{T} = \{r_1, r_2 r_3, r_3 r_4\}$, and, for instance, $\mathcal{T}^1 = \{r_1\}$. Because $r_2 r_3$ overlaps with $r_3 r_4$, this function is not tree-structured, and of course no set containing this function would ever be.

It is also helpful to 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.

Example 2 *Consider the set composed the three following utility functions :*

$$\begin{aligned} u_1 &= r_2 + 3r_5 \\ u_2 &= 3r_1 + 10 r_1 r_2 r_3 r_4 + 8r_5 + 4r_6 \\ u_3 &= r_6 - r_4 + 8 r_3 r_4 \end{aligned}$$

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 represented by the following tree :

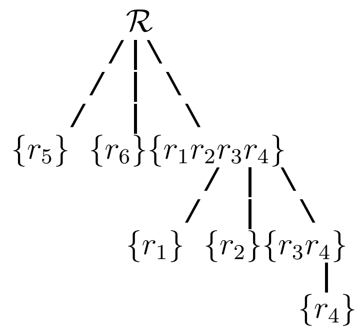

Because finding rational k-deals is in general intractable, our aim is to exploit this new restriction on utility functions to come up with other types of deals less complex, but insuring the same convergence properties.

Definition 4 *A $\mathcal{T}$-deal is a deal involving an entire term of $\mathcal{T}$ from one or more sender(s) to a* single *receiver.*

First, note that the number of possible $\mathcal{T}$-deals to consider shrinks down to ($2 \times n \times |\mathcal{R}|$). (This is so because, intuitively, the maximum number of nodes one may observe with such trees would be obtained by splitting each term in two balanced sub-nodes). In other words, the complexity of finding rational $\mathcal{T}$-deals is also very low compared to that of finding k-deals.

However, as illustrated by the example below, simply allowing any $\mathcal{T}$-deals will not be sufficient to guarantee us that the optimal social welfare will eventually be reached in all cases.

Example 3 *Let us consider the following utility functions :*

$$\begin{aligned} u_1 &= 10 r_1 \\ u_2 &= 10 r_2 \\ u_3 &= 11 r_1 r_2 \end{aligned}$$

and let the initial allocation be the allocation assigning all resources to agent 3. The (only) optimal allocation consists in allocating r_1 to agent 1 and r_2 to agent 2. Clearly, there are no rational $\mathcal{T}$-deals allowing the system to reach this optimal allocation.

[2] 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.

In order to restore the desired convergence properties of the framework, we will need to investigate more complex protocols restricting the negotiation process and exploiting through $\mathcal{T}$-deals the properties of tree-structured utilities.

6 Tree-Climbing Protocols

The basic idea of *tree-climbing protocols* is to allow $\mathcal{T}$-deals involving only the smallest bundles first, then to incrementally allow bigger bundles, until all possible $\mathcal{T}$-deals have been allowed.

Algorithm 1 Naive Tree-Climbing Protocol

Require: n agents with tree-structured utilities $u_1 ... u_n$ (with $\mathcal{T}$ the set of all terms).
$l \leftarrow 1$.
repeat
  Restrict allowed deals to $\mathcal{T}^l$-deals.
  Let agents do all their individually rational $\mathcal{T}^l$-deals.
  if no more deal can be conducted **then**
    $l \leftarrow l + 1$
  end if
until $l >\mid \mathcal{R} \mid$

Clearly, this protocol is more restrictive than the liberal "protocol" consisting of allowing any $\mathcal{T}$-deals, *i.e* it cannot guarantee an optimal outcome neither. By inspecting more carefully the scenario described in Example 3 however, we find clues indicating possible improvements of this protocol : during the first step of the protocol (where $l = 1$, and only $\mathcal{T}^1$-deals are allowed), the term $r_1 r_2$ in u_3 prevents the resources from moving away from agent 3. If, during this first step, terms containing more than one resource were *removed* from the utility functions, then the resources would be able to move towards agents 1 and 2.

6.1 Intrusive Tree-climbing protocol

Based on this idea, we design a new protocol which (at each step) l removes in each utility function terms of more than l resources. We call the following protocol *Intrusive Tree-Climbing Protocol* (ITCP) in the sense that it modifies the utility functions of each agent at each step. In the following, $u_i^{\leq l}$ is defined as u_i where all terms of size greater than l have been removed. By extension, we also note $sw^{\leq l} = \sum_i u_i^{\leq l}$ for the social welfare induced by these truncated utilities.

Algorithm 2 Intrusive Tree-Climbing Protocol.

Require: n agents with utilities $u_1 ... u_n$ identically tree-structured. Let $\mathcal{T}$ the set of all terms.
$l \leftarrow 1$.
repeat
  Restrict allowed deals to $\mathcal{T}^l$-deals,
  Enforce all agents i to change their utilities to $u_i^{\leq l}$.
  Let agents do all their rational $\mathcal{T}^l$-deals.
  if no more deal can be conducted **then**
    $l \leftarrow l + 1$
  end if
until $l >\mid \mathcal{R} \mid$

We first show that each step of the protocol will indeed lead to an allocation deemed optimal w.r.t. the truncated utility functions of agents populating the system.

Lemma 1 *If, at step l of the ITCP, an optimal allocation for $sw^{\leq l}$ is reached, then, at step $l+1$, by allowing only $\mathcal{T}^{l+1}$-deals, an optimal allocation for $sw^{\leq l+1}$ will be reached.*

Proof. Suppose at the l^{th} step, an optimal allocation A^l has been reached. At step $l + 1$, the negotiation process thus starts from A^l. Let us show that the optimal allocation w.r.t. $sw^{\leq l+1}$ will be reached. The social welfare can be written $sw^{\leq l}(A) = \sum_{t \in \mathcal{T}^{\leq l}} \beta_t^A$ where β_t^A is defined as follows :

$$\beta_t^A = \begin{cases} \alpha_t^i & \text{if an agent } a_i \text{ owns the term } t \text{ in } A \\ 0 & \text{if } t \text{ is split among several agents} \end{cases}$$

Also, $\gamma_t^A = \sum_{s \subset t} \beta_s^A$. When there is no ambiguity, the allocation will be omitted. Let $\overline{\mathcal{T}^{l+1}} = \left\{ r \in \mathcal{R} \mid \forall t' \in \mathcal{T}^{l+1}, r \notin t' \right\}$ all resources not present in $\mathcal{T}^{l+1}$. Note that $\mathcal{T}^{l+1} \cup \{\overline{\mathcal{T}^{l+1}}\}$ is a partition of $\mathcal{R}$.

$$sw^{\leq l+1}(A) \;=\; \sum_{t \in \mathcal{T}^{l+1}} (\beta_t + \gamma_t) + \sum_{t \subseteq \overline{\mathcal{T}^{l+1}}, t \in \mathcal{T}^{\leq l}} \beta_t$$

Note that the terms of the last equation are all independent. Thus,

$$\max_A \left\{sw^{\leq l+1}(A)\right\} = \sum_{t\in\mathcal{T}^{l+1}} \max_A \{\beta_t + \gamma_t\} + \max_A \left\{ \sum_{t\in\overline{\mathcal{T}^{l+1}}, t\in\mathcal{T}^{\leq l}} \beta_t \right\}$$

First note that the last term of the equation has already been maximized in A^l (which is supposed maximal w.r.t. $sw^{\leq l}$), and because it cannot be affected by $\mathcal{T}^{l+1}$-deals, this term will stay maximal and unaffected during step $l+1$. Therefore, we only need to show that terms $\beta_t + \gamma_t$ are maximized.

Let us show now that the allocation obtained after a sequence of individually rational $\mathcal{T}^{l+1}$-deals with side-payments starting from A^l will be A^{l+1}, the optimal allocation w.r.t. $sw^{\leq l+1}$. For this purpose, let us study the maximal value of $\beta_t + \gamma_t$ for each term $t \in \mathcal{T}^{l+1}$.

$$\max(\beta_t+\gamma_t) = \begin{cases} \max_i u_i(t) & \text{if } \beta_t \neq 0 \\ \max\{\gamma_t \mid \beta_t = 0\} & \text{if } \beta_t = 0 \end{cases}$$

Clearly, $\max\{\gamma_t \mid \beta_t = 0\} = \gamma_t^{A^l}$. Thus, this last equation tells that the optimal allocation for term t w.r.t. $sw^{\leq l+1}$ is either as in A^l, or t is owned by agent $i = argmax_i u_i(t)$. In the former case, there will never be any rational $\mathcal{T}^{l+1}$-deal involving t, as it is already in its optimal position in A^l. In the latter case, the deal consisting of moving t towards agent $i = argmax_i u_i(t)$ will be rational at any time during step $l+1$. Thus, allocation A^{l+1} will be reached eventually. □

We are now in a position to prove that the ITCP indeed meets the intended convergence property (note that this result still holds when utility functions contain non-positive coefficients).

Theorem 2 *For any set of tree-structured k-additive utilities, the ITCP will eventually result in an optimal social welfare.*

Proof. It has been proven elsewhere that additive utilities converge towards local optimum [7]. Thus, the property is true from the first step of the protocol. Lemma 1 shows that it remains true for the next steps. □

A possible perspective on the ITCP would be that it works with a sequence of progressively more fine-grained *approximations* to the real utility functions in such a way that structurally simple deals become individually rational (with respect to these approximations). Of course, these are approximations with respect to the *representation* of utility functions (in k-additive form), but not necessarily with respect to the values that utility functions assign to bundles of resources.

These two dimensions may, however, coincide in many practical cases. If we assume that the synergetic effect a bundle of resources reduces as the cardinality of that bundle increases, then the sequence of auxiliary utility functions used in the ITCP will indeed be progressively more accurate in terms of the values they assign to bundles as negotiation develops. In other words, if the coefficients in the k-additive form are high for single resources, and if they get smaller as the size of the respective term increases, then the utility functions used in the ITCP become more accurate in each round. This is a very reasonable assumption to make : the larger a bundle of resources is, the more difficult would it be for an agent to estimate the additional benefit incurred by owning all the resources in that bundle *together* (i.e. beyond the benefit incurred by the relevant subsets). Indeed the very same (cognitive) argument is one of the main reasons why k-additive utility functions with low values of k are not only computationally attractive but also highly relevant in practice.

Nevertheless, the ITCP *does* violate the postulate of full individual rationality for all stages of a negotiation process. Also, because of its "intrusive" aspect, this protocol is not completely satisfactory : it cannot be applied when agents are not under the control of the designer. Clearly, such a protocol would be of little use when agents are humans, or simply when privacy issues prevent the designer to directly access the agents' preferences.

6.2 A Non-Intrusive Protocol

To overcome the drawbacks of intrusive protocols, we design a third protocol in which utilities are not modified, but such that for each deal, a special side-payments function influences the rationality of each agent. In addition, the system also participates in the side-payments, by adding or taking money from the negotiating agents. The trick here is that these side-payments have exactly the same effect as a change of utility, without being intrusive (as shown in Theorem 3). Thus, this protocol could be used in real applications, even with human agents as long as utilities have been elicited.

This protocol is called *Omniscient ϵ-Altruistic Tree-Climbing Protocol* (Oϵ-ATCP), because the system 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.

Algorithm 3 Omniscient ϵ-Altruistic Tree-Climbing Protocol

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

All agents transmit to the *system 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 $\Delta_{sw^{\leq l}(A',A)} = sw^{\leq l}(A') - sw^{\leq l}(A)$
  The associated side-payment is then, for all the n' agents involved in the deal : $p_i = u_i(A') - u_i(A) - \frac{\epsilon}{n'} \times \Delta_{sw^{\leq l}(A',A)}$
  For the system agent, $p_{sys} = -\sum p_i$
  Let agents do all their rational deals.
  if no more deal can be conducted **then**
    $l \leftarrow l + 1$
  end if
until $l >\mid \mathcal{R} \mid$

Note that in the Oϵ-ATCP protocol, the *system agent* could be viewed as a kind of bank giving or receiving money from agents. The amount of money shared with other agents depends on the difference of satisfaction enjoyed by the whole society before and after the deal (that is $sw^{\leq l}(A') - sw^{\leq l}(A)$, that we shall note $\Delta_{sw^{\leq l}(A',A)}$), and on the value of a parameter ϵ.

– when ϵ is close to 1, the system agent in not guaranteed to globally earn money. In that case, the system agent is said to be *altruistic*, because negotiating agents will benefit from this by earning a lot.
– when ϵ is close to 0, then during each deal, the system 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 system 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 system agent would be equally redistributed at the end of the process.

The following results show that this protocol also converges to an allocation with optimal social welfare, as the ITCP does.

Lemma 2 *In Oϵ-ATCP at step l, a deal $\delta = (A, A')$ is rational iff $sw^{\leq l}(A') > sw^{\leq l}(A)$.*

Proof. '$\Rightarrow$'. Suppose the deal is rational. Thus, for each agents i involved, $u_i(A') - u_i(A) - p_i > 0$. Thus,

$$\sum_i (u_i(A') - u_i(A)) - \sum_i p_i =$$
$$\epsilon \times \Delta_{sw^{\leq l}(A',A)} > 0$$

'$\Leftarrow$'. Suppose $sw^{\leq l}(A') > sw^{\leq l}(A)$. Let us show that for each agent i involved in the deal, the rationality criterion is met.

$$u_i(A') - u_i(A) - p_i =$$
$$\frac{\epsilon}{n'} \Delta_{sw^{\leq l}(A',A)} > 0$$

□

This directly leads to the convergence result.

Theorem 3 *For any set of tree-structured k-additive utilities, the Oϵ-ATCP will eventually result in an optimal social welfare if $\epsilon \in]0,1]$.*

Proof. We show that deals accepted in this protocol are the same as those accepted in the ITCP protocol. Clearly, at step l, a deal $\delta = (A, A')$ is accepted in ITCP iff $sw^{\leq l}(A') > sw^{\leq l}(A)$. Lemma 2 show that O$\epsilon$-ATCP behaves the same way. □

A critical point of the Oϵ-ATCP is that it necessitates an external provider of money. However, the following result shows that the system agent globally earns money, as long as we choose a sufficiently small value for ϵ.

Theorem 4 *If ϵ is sufficiently small, the system agent will globally earn money on any Oϵ-ATCP negotiation as a whole.*

Proof. Let us compute how much is earned by the system agent each time a deal is conducted. Suppose at step l, a deal $\delta = (A, A')$ is conducted by n' agents. The system agent pays

$$\begin{aligned} p_{sys} &= -\sum_i (u_i(A') - u_i(A)) + \frac{\epsilon}{n'} \Delta_{sw^{\leq l}(A',A)} \\ &= -(sw(A') - sw(A)) + \epsilon \Delta_{sw^{\leq l}(A',A)} \end{aligned}$$

Let d^{max} be the number of deals conducted during the whole negotiation. Let p_{sys}^{total} be the total amount of money the system agent pays during the whole negociation process. Thus, $-p_{sys}^{total}$, the money earned by the system, is equal to

$$\sum_{\delta=(A,A')} \left(u_i(A') - u_i(A) - \frac{\epsilon}{n'} \Delta_{sw^{\leq l}(A',A)} \right)$$

$$\geq sw(A_{opt}) - sw(A_0) - \epsilon \times d^{max} \times \max_{A,A',l} \Delta_{sw^{\leq l}(A',A)}$$

Clearly, we can see that

$$\lim_{\epsilon \to 0} (-p_{sys}^{total}) = sw(A_{opt}) - sw(A_0)$$

and the system earns as much money as the increase of social welfare. □

The central idea in the Oϵ-ATCP is that the system agent (i.e. the mechanism) can give a "loan" to agents to allow them to accept otherwise disadvantageous deals, but that these agents repay this loan later on when they arrange deals that are overly beneficial for themselves. This general idea would, of course, be applicable also in other, more general, scenarios than we have considered here, i.e. scenarios where agents do not necessarily have tree-structured utility functions. As long a the final allocation has higher utilitarian social welfare than the initial allocation and as long as the system agent has sufficient funds to temporarily sustain allocations with very low social welfare, *any* sequence of deals can be made individually rational from the viewpoint of the negotiating agents and would still result in an eventual surplus for the system agent.

7 Conclusion

We have further analysed a negotiation framework previously studied by several authors [1, 4, 5, 6]. While most work on negotiation in multiagent systems has addressed either bilateral negotiation [13] or auctions [3], this framework is multilateral, *i.e.* deals may involve any number of agents and any number of resources. The requirement for full multilateral negotiation stems from the fact that agents are assumed to be both rational and myopic. However, the task of identifying rational deals, *i.e.* deals that are beneficial for all participants remains highly complex. In this paper, it is shown that if all agents model their preferences by means of *tree-structured* utility functions, the optimal social welfare can be reached in linear time by implementing adequate protocols. Two protocols (clearly designed to exploit the specific structure of these functions) have been proposed in this paper.

The crucial contribution in our proposal, we believe, is the *combination* of two ideas :

– The central idea in the ITCP has been to provide a protocol that *guides* negotiation. In each round, there are only a relatively small number of options as to which deal to choose (and the complexity of these deals increases stepwise, from round to round). The particular choices made in the definition of the ITCP mean that deals are not too far from truly rational deals (certainly if we make some additional assumptions on the nature of utility functions, as discussed earlier).
– The additional idea in the Oϵ-ATCP has been to introduce a system agent that can *compensate* for a temporary loss in social welfare (provided that such a loss is due to a deal sanctioned by the guidelines implemented in the ITCP). It thereby allows for a wider range of deal sequences leading to the optimal allocation, including sequences of structurally simple deals.

This combination of *guidance* and *social compensation* makes our approach work. While our current results are restricted to tree-structured domains, we believe that the central ideas are general enough to be applied to other scenarios as well.

Our proposal may be considered a hybrid between distributed and centralised approaches to multiagent resource allocation. On the one hand, allocations emerge in a distributed manner as a result of a sequence of local negotiation steps.

On the other hand, the introduction of the system agent does introduce a centralised element. One of the main arguments usually given in favour of distributed approaches over centralised ones (such as combinatorial auctions), is that it may not always be possible to find an agent that could act as the central authority (the auctioneer) executing the allocation procedure. To some degree, the same reservations may be made against our proposal. However, we believe there is at least one imortant difference. In a combinatorial auction, the auctioneer is endowed with the (very significant) computational burden of deciding the final allocation (the so-called winner determination problem [3]). In our approach, on the other hand, there are no computationally hard tasks associated with the role of the system agent. The computational burden lies with the system of agents as a whole and, indeed, this burden is substantially reduced by virtue of the simplifications made in the ITCP.

A potential problem with our approach is that individual agents may decide to stop participating in negotiation before the optimal allocation has been reached. Indeed, it may be individually rational for them to do so, because our protocols allow certain agents to temporarily enjoy higher individual welfare than they would enjoy in the socially optimal allocation. In this sense, the framework loses the "anytime character" of the original system [14], where every new allocation would be guaranteed to have higher social welfare than its predecessor. This anytime character is particularly important because of the prohibitively high complexity of the general framework [1, 5, 6]. If optimal allocations cannot be guaranteed to be found within a reasonable amount of time, then the original framework does at least guarantee some improvement over the initial allocation when time runs out. In our present proposal, the anytime character of the framework is far less important, because —as we have argued— it dramatically reduces the complexity of negotiation and thereby makes convergence to an optimal allocation within a reasonable amount of time a realistic option. Nevertheless, the issue of incentive compatibility in the Oϵ-ATCP remains an important issue that requires further investigation.

Références

[1] Y. Chevaleyre, U. Endriss, S. Estivie, and N. Maudet. Multiagent resource allocation with k-additive utility functions. In *Proceedings of the DIMACS-LAMSADE Workshop on Computer Science and Decision Theory*, volume 3 of *Annales du LAMSADE*, 2004.

[2] Y. Chevaleyre, U. Endriss, J. Lang, and N. Maudet. Negotiating over small bundles of resources. In *Proc. AAMAS-2005*. ACM Press, July 2005. To appear.

[3] P. Cramton, Y. Shoham, and R. Steinberg, editors. *Combinatorial Auctions*. MIT Press, 2005. To appear.

[4] P. E. Dunne. Extremal behaviour in multiagent contract negotiation. *Journal of Artificial Intelligence Research*, 23 :41–78, 2005.

[5] P. E. Dunne, M. Wooldridge, and M. Laurence. The complexity of contract negotiation. *Artificial Intelligence*, 2005. To appear.

[6] U. Endriss and N. Maudet. On the communication complexity of multilateral trading. In *Proc. AAMAS-2004*. ACM Press, 2004.

[7] U. Endriss, N. Maudet, F. Sadri, and F. Toni. On optimal outcomes of negotiations over resources. In *Proc. AAMAS-2003*. ACM Press, 2003.

[8] U. Endriss, N. Maudet, F. Sadri, and F. Toni. Resource allocation in egalitarian agent societies. In *Proc. MFI-2003*. Cépaduès-Éditions, 2003.

[9] A. M. Feldman. Bilateral trading processes, pairwise optimality, and Pareto optimality. *The Review of Economic Studies*, 40(4) :463–473, 1973.

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

[11] G. E. Kersten, S. J. Noronha, and J. Teich. Are all e-commerce negotiations auctions ? In *Proc. of the 4th Intl. Conf. on the Design of Cooperative Systems*, 2000.

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

[13] Jeffrey S. Rosenschein and Gilad Zlotkin. *Rules of Encounter*. MIT Press, 1994.

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

[15] A. K. Sen. *Collective Choice and Social Welfare*. Holden Day, 1970.