

On the Communication Complexity of Multilateral Trading: Extended Report *

Ulle Endriss (ue@doc.ic.ac.uk)

Department of Computing, Imperial College London, London SW7 2AZ, UK

Nicolas Maudet (maudet@lamsade.dauphine.fr)

LAMSADE, Université Paris-Dauphine, 75775 Paris Cedex 16, France

Abstract. We study the complexity of a multilateral negotiation framework where autonomous agents agree on a sequence of deals to exchange sets of discrete resources in order to both further their own goals and to achieve a distribution of resources that is socially optimal. When analysing such a framework, we can distinguish different aspects of complexity: How many deals are required to reach an optimal allocation of resources? How many communicative exchanges are required to agree on one such deal? How complex a communication language do we require? And finally, how complex is the reasoning task faced by each agent?

Keywords: Negotiation, multiagent resource allocation, social welfare, complexity

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 [3, 14, 17, 19, 21]. In this paper, we study the complexity of a negotiation framework where autonomous agents agree on a sequence of deals to exchange sets of discrete (*i.e.* non-divisible) resources. An important characteristic of the framework considered here is that agents use very simple rationality criteria to decide whether or not to accept a proposed deal, but interaction patterns may be complex. In particular, *multilateral* deals, *i.e.* deals involving more than two agents at a time, are possible.

We may think of a multiagent system as a *society of agents*. 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*. Several formal models of social optimality that are applicable to our framework have been studied in welfare economics and social choice theory [2, 18, 22]. In this paper,

* This is an extended version of a paper by the same name that appeared in the *Proceedings of the 3rd International Joint Conference on Autonomous Agents and Multiagent Systems* [7].

we are mostly concerned with maximising *utilitarian social welfare*, but also with negotiating *Pareto optimal* allocations of resources.¹

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 [8, 9, 20]. A first analysis of the complexity of certain aspects of this framework has recently been given by Dunne *et al.* [6]. In the present paper, we put particular emphasis on the *communication complexity* of multilateral trading. That is, we are more interested in the length of negotiation processes and the amount of information that needs to be exchanged between agents than in the computational complexity associated with the tasks individual agents need to carry out for negotiation to take place.

The contribution made in this paper is twofold. Firstly, we present an overview of the different types of complexity that arise in the context of distributed negotiation over resources. This analysis will be relevant to a wide range of different negotiation frameworks. Secondly, we study one type of communication complexity in detail and establish several upper bounds on the length of both the shortest and the longest path to a socially optimal allocation of resources in a framework where rational agents agree on a sequence of deals, each of which is beneficial to all participating agents.

The remainder of this paper is structured as follows. In Section 2 we review the multilateral trading framework of [9] and quote several results on the reachability of socially optimal allocations of resources by means of specific classes of deals. Section 3 identifies different aspects of the complexity of trading resources. While we take the framework of [9] as a reference model, most of these issues are likely to be relevant to any scenario where agents negotiate over resources. The first type of complexity identified in Section 3 concerns the number of deals that need to be implemented for an agent society to converge to an optimal state. In Section 4 we prove several upper bounds on this number of deals. Different results apply in different cases, depending on the class of deals considered and whether we are interested in either the number of deals in the shortest path to an optimal allocation or the number of deals in the longest possible path before any negotiation process is bound to terminate. Section 5 concludes.

¹ The utilitarian social welfare of an allocation is defined as the sum of the utilities enjoyed by the individual agents in the system for that allocation; an allocation is Pareto optimal iff there is no other allocation that would be preferred by at least one of the agents without being worse for any of the others (see Section 2).

2. Resource Allocation by Negotiation

In this section, we introduce the framework of *resource allocation by negotiation* studied in [9] 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 rational 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 $-p(i)$. We distinguish *deals with money* and *deals without money*. For the latter, $p(i)$ is required to be 0 for every agent $i \in \mathcal{A}$. For the framework without money, it would be sufficient to model an agent's preferences by means of a (not necessarily strict) total order over alternative bundles of resources. We use utility functions nevertheless, but for presentational reasons alone.

While most work on negotiation in multiagent systems has been concerned with either *bilateral* (“one-to-one”) negotiation or *auctions* [17, 19, 23, 24], we should stress that our scenario of resource allocation by negotiation explicitly addresses *multilateral* exchanges and that it is *not* an auction. Auctions are mechanisms to help agents agree on a price at which an item (or a set of items) is to be sold [13]. 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.

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 rational 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 [20] 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 $A(i) = A'(i)$.*

That is, an agent i will be prepared to accept a deal δ iff its gain in utility (or money) strictly outweighs a possible loss in money (or utility). Only agents that are not involved in a deal, *i.e.* whose bundle remains the same, do not require a positive payoff. Note that we adopt a very simple concept of money here; the utility derived from of a given side payment is always equal to its dollar value. In other words, we may think of money as a single continuous commodity that has the same marginal utility for all the agents in the system. We also assume that agents have sufficient funds to make the side payments associated with any given rational deal. These assumptions are appropriate for applications where our main interest is in the reallocation of discrete resources and we require a convenient means of implementing explicit utility transfers, but the framework could also be extended to deal with more sophisticated models of money.

Nevertheless, these assumptions will not be justified in all circumstances, which is why we have also put forward a similar negotiation framework *without money* [9]. For the framework without money, Definition 1 can be simplified to say that any rational deal should result in a strict increase in utility for all the agents involved. However, as discussed in detail in [9], it is useful to slightly weaken the notion of rationality to be able to compensate for the fact that the framework without money does not allow us to model arbitrarily small increases in utility. That is, in scenarios where side payments are not possible, agents will be required to be cooperative in the sense of also accepting deals that do at least not decrease their personal utility:

DEFINITION 2 (Cooperative rationality). *A deal $\delta = (A, A')$ without money is rational iff $u_i(A) \leq u_i(A')$ for all agents $i \in \mathcal{A}$ and this inequality is strict in at least one case.*

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.

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. [18]) 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 3 (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 [2, 18, 22] and may be usefully exploited in the context of multiagent systems [8].

A related notion is the concept of *Pareto optimality*, which may be defined as follows:

DEFINITION 4 (Pareto optimality). *An allocation A is called Pareto optimal iff there is no other allocation A' such that $sw(A) < sw(A')$ and $u_i(A) \leq u_i(A')$ for all $i \in \mathcal{A}$.*

In other words, an allocation is Pareto optimal iff there is no other allocation that is better for at least one agent without making any of the others worse off.

2.3. CONVERGENCE RESULTS

Proofs for all the theorems quoted in this section may be found in [9]. The first of these, which is equivalent to a result on sufficient contract types for optimal task allocations by Sandholm [20], 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 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 with 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 negotiation, the system will never get stuck in a local optimum and finding an optimal allocation remains an option throughout.

For the framework without money, we can only guarantee negotiation outcomes that are Pareto optimal:

THEOREM 2 (Pareto optimal outcomes). *Any sequence of rational deals without money will eventually result in a Pareto optimal allocation.*

A drawback of the general frameworks, to which Theorems 1 and 2 apply, is that these results only hold if deals involving any number of resources and agents are admissible [9, 20]. That is, any negotiation protocol that puts restrictions on the structure of deals that agents may agree on will fail to guarantee socially optimal outcomes, even when there are no constraints on either time or computational resources. 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*):

THEOREM 3 (Additive scenarios). *If all utility functions are additive, then any sequence of rational one-resource-at-a-time deals with money will eventually result in an allocation with maximal social welfare.*

If we merely wish to model whether or not an agent *needs* a particular resource, it is sufficient to use additive utility functions which assign either 0 or 1 to each single resource. If all agents use these *0-1 functions* to model their preferences, then the previous result can be further strengthened to apply also to deals without money:

THEOREM 4 (0-1 scenarios). *If all utility functions are 0-1 functions, then any sequence of rational one-resource-at-a-time deals without*

money will eventually result in an allocation with maximal social welfare.

A question that naturally arises when we consider these convergence results is, *how many* deals of the class in question (such as the class of rational one-resource-at-a-time deals without money) are actually required to reach the respective optimal allocation of resources. We are going to discuss this question in detail in Section 4. First, however, we are going to take a somewhat broader perspective and analyse 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 study the complexity of trading within the negotiation framework lined out in the previous section. As it happens, there is not just a single notion of complexity that is of relevance here. In fact, we can distinguish at least four different *aspects of complexity*. They 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 be able 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 individual 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. We are going to analyse the number of deals required to reach optimal allocations for several instances of our negotiation framework in Section 4.

At the next lower level, we have to consider the complexity of negotiating a *single* deal in such a sequence of deals converging to an optimal allocation. 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? How many different performatives are required, and how much information do we have to be able to transfer by means of a single message?

In the remainder of this section, we are going to discuss some of these issues further and point out connections to related work in the literature.

3.1. COMMUNICATION COMPLEXITY

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. In the literature on distributed computing, this term is used to refer to the number of bits that the nodes in a distributed system need to exchange in order to jointly compute the value of a given function [15]. The so-called two-party model of communication complexity introduced by Yao [27] addresses the following problem: Two agents A and B each hold an n -bit string and their goal is to communicate in order to compute the value of a (boolean) function over these two strings. The question then is: What is the minimal number of bits that need to be exchanged to be able to compute that function? In particular, the model is *not* concerned with the computational resources required by the agents, but only with the amount of communication needed. Of course, it is always possible for A to simply send its entire n -bit string; then B computes the result (because it also knows its own n -bit string); and finally B sends this result back to A (*i.e.* $n + 1$ bits are being exchanged). But for many types of functions there are more efficient protocols. Functions classically studied in the area of communication complexity are often simple functions such as equality or the greater-than relation over integers.

The communication complexity of a *protocol* is the maximal number of bits exchanged when following that protocol (in the worst case). The communication complexity of a *function* is the communication complexity of the best protocol that computes that function.

While we do not use the term communication complexity in precisely the same sense, there are a number of parallels to be observed. The communication complexity of arranging a single deal is a combination of the number of dialogue moves that need to be sent and the amount of information contained in a single message. 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 [9, 20]. 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 [11],

- 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* as the size of the group increases (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.

3.2. MINIMAL REQUIREMENTS FOR PROTOCOLS

In what follows, we briefly discuss some of the very basic considerations pertaining to our third type of complexity, *i.e.* the complexity of the communication language (including an appropriate interaction protocol) used to negotiate. While it is generally considered desirable that both dialogue moves and protocol rules are as simple as possible, it is also important to find the right balance between simplicity and expressive power. A restricted communication language may, for instance, have negative impacts on the length of a negotiation dialogue or the quality of the deal agreed upon (which in turn would negatively affect the overall number of deals required).

Any communication language for negotiation in our framework is likely to include at least performatives such as *propose*, *accept*, and *reject*, to be able to communicate a proposed deal to (a set of) potential trading partners and to either accept or reject such a proposal, respectively (for the terminology used to describe communication protocols see, for instance, [26]). Naturally, a *sophisticated* protocol would

also include performatives to enable agents to negotiate aspects of a deal step by step, but the above seem to be minimal requirements for any suitable protocol.² The content of a *propose* move would have to include a full specification of the deal in question, *i.e.* we require a *content language* that is rich enough to express which resources are to be moved from which agent to which other agent (possibly together with the specification of a payment function). Amongst other things, the complexity of this content language would depend on the number of distinct deals that are possible (or that *may* be possible according to the knowledge of the agent proposing the next deal) at any one point during negotiation.

Related to these considerations is also the complexity of the problem of checking that a given dialogue move actually *conforms* to the chosen communication protocol. This is an important requirement in open agent societies where the correct interpretation of dialogue moves relies on the availability of system-wide norms of interaction [25]. The reason why we have not listed the complexity of checking conformance to a protocol amongst our four *aspects of complexity* at the beginning of Section 3 is that this is a general issue in agent communication rather than a problem of specific relevance to our interests in negotiating socially optimal allocation of resources.

3.3. COMPUTATIONAL COMPLEXITY

The fourth type of complexity identified earlier, *i.e.* the complexity of the reasoning task faced by an agent when deciding on its next step in a negotiation dialogue, is the only kind of *computational* complexity we have considered. It is clear that our classification does not cover *all* aspects of (computational) complexity; the reasoning task faced by an agent could be further decomposed into subproblems. For instance, we may also consider the complexity of determining whether a given bundle yields higher individual welfare than the current one (although the problem of preference elicitation lies outside the scope of this paper as we take utility functions as given).

While, in this paper, we restrict ourselves to studying the complexity of problems arising when agents strictly follow the interaction patterns foreseen in our basic negotiation framework, there are also several related problems of interest. In the remainder of this section, we briefly review some of the recent complexity results in this area.

² Instead of using the *reject* performative, we could also adopt a convention whereby proposals that are not accepted explicitly count as having been rejected. Still, *conceptually*, agents do need to be able to propose, accept and reject deals.

Dunne *et al.* [6] study the computational complexity of deciding whether one-resource-at-a-time trading (with money) is sufficient to move to a given allocation with higher social welfare than the current one. This is what one may want to call the complexity of a “meta-property” of the framework. Agents engaged in negotiation are not actually going to analyse this kind of “global” question, but rather try to agree on deals at the local level. The types of complexity we have identified here all relate directly to the problems faced by agents when engaged in negotiation, while the decision problem of Dunne *et al.* is more likely to be tackled by an outside observer.³ Nevertheless, these different views on the complexity of negotiation are strongly inter-related: the complexity of negotiating an optimal allocation, in a distributed manner, by means of a sequence of one-resource-at-a-time deals is bound to be at least as high as that of the problem of deciding whether such a sequence exists in the first place (Dunne *et al.* have shown that their decision problem is NP-hard).

Several authors have addressed the problem of finding an allocation with maximal social welfare by means of a “centralised” mechanism (rather than by negotiation, which is what we are interested in here). Dunne *et al.* [6] and Chevaleyre *et al.* [4] have analysed the complexity of this problem with respect to different ways of representing utility functions, while Kraus [14] and Fargier *et al.* [10] have considered the complexity of maximising social welfare whilst also satisfying certain side conditions (such as ensuring that each agent enjoys at least a given minimal level of individual welfare). It turns out that the different variants of this problem are all closely related to well-known combinatorial optimisation problems [12]: the complexity of finding an allocation with a social welfare exceeding a given threshold K is generally NP-complete.

4. Number of Deals

In this section, we are going to address the question characterising the first type of communication complexity identified earlier: *How many deals are required to reach an optimal allocation of resources?*

We are going to study this question in the concrete context of the negotiation framework set out in Section 2 and, specifically, we are going to analyse how many deals are required to reach the optimal allocations referred to in each of the four convergence theorems quoted towards the end of that section. The class of deals considered (with or

³ A possible exception may be a scenario where an agent tries to plan a negotiation process consisting only of one-resource-at-a-time deals that could be accepted by a trading partner without risk.

without money; one-resource-at-a-time or general) as well as the type of optimality that can be achieved (maximal social welfare or Pareto optimality) differ for each of these theorems. For instance, related to Theorem 3, we are going to investigate how many rational one-resource-at-a-time deals with money are required to reach an allocation with maximal social welfare in an additive scenario.⁴

Of course, 0 is always going to be a *lower bound*: If the initial allocation of resources is itself optimal, then not a single deal will be required to reach an optimal allocation. Hence, we are only going to be interested in upper bounds. In fact, there are two types of upper bounds that one may consider: the maximal length of the *shortest path* to an optimal allocation and the maximal length of the *longest path* to such an allocation.

4.1. MAXIMISING SOCIAL WELFARE

Let us first consider scenarios where there are no restrictions on utility functions and where any rational deal with money is admissible (the framework of Theorem 1). In this context, our question reads: *How many rational deals with money are required to reach an allocation with maximal social welfare?*

The (possibly somewhat surprising) answer to this question is: “1”. To back up this claim, we require the following lemma, which has been proved in [9]:

LEMMA 1 (Rational deals with money). *A deal $\delta = (A, A')$ with money is rational iff $sw(A) < sw(A')$.*

The upper bound for the shortest path to an optimal allocation follows immediately:

THEOREM 5 (Shortest path with money). *An allocation with maximal social welfare can always be reached by means of (at most) a single rational deal with money.*

Proof. Let A be the initial allocation and suppose A does not have maximal social welfare (otherwise the theorem holds vacuously). Then, for any allocation A' with maximal social welfare, we have $sw(A) < sw(A')$. Hence, by Lemma 1, the deal $\delta = (A, A')$ must be rational, *i.e.* A' can be reached by means of a single rational deal. \square

⁴ In a recent paper, Dunne [5] addresses a related problem and analyses the number of deals meeting certain structural requirements (such as being one-resource-at-a-time deals) needed to reach a given target allocation (whenever this is possible at all) in general scenarios.

Naturally, agents would have to be very lucky to negotiate such a perfect deal in the first round. The central point of Theorem 1 is a very different one, however: even if agents are not that lucky and farsighted, they are going to reach an optimal allocation *eventually*, provided they only agree on deals that are rational. How many deals would be required in the very worst case? Lemma 1 shows that any rational deal will result in a strict increase in social welfare. Hence, certainly no allocation can be visited twice. To see whether there could be a scenario where each and every allocation gets visited once, we need to check whether it is possible that all allocations have distinct social welfare.

LEMMA 2 (Distinct welfare). *There exist utility functions such that distinct allocations have distinct social welfare.*

Proof. Let $m = |\mathcal{A}|$ be the number of agents in our society. To simplify our presentation, we identify the set of agents with an initial segment of the non-negative integers, *i.e.* $\mathcal{A} = \{0, 1, \dots, m-1\}$. Furthermore, let $n = |\mathcal{R}|$ be the number of resources in the system, *i.e.* there are 2^n different bundles an agent may hold. We first define a “base utility function” u^* that assigns to each bundle an integer between 0 and 2^n-1 , without assigning the same number to any two distinct bundles. We then define the utility function u_i of each agent $i \in \mathcal{A}$ as follows:

$$u_i(R) = u^*(R) \cdot (2^n)^i \quad (\text{for bundles } R \subseteq \mathcal{R})$$

These utility functions verify the claim of the lemma: for any two allocations A and A' , $sw(A)$ will be different from $sw(A')$ whenever $A \neq A'$. To see this, recall the definition of social welfare:

$$sw(A) = \sum_{i \in \mathcal{A}} u_i(A(i)) = \sum_{i \in \mathcal{A}} u^*(A(i)) \cdot (2^n)^i$$

This sum may be thought of as the representation of $sw(A)$ in a number system with base 2^n : $u^*(A(i))$ contributes the digit and i determines the position of that digit. If $A \neq A'$, then the bundle $A(i)$ will differ from $A'(i)$ for at least one agent $i \in \mathcal{A}$, *i.e.* $sw(A)$ will differ from $sw(A')$ in at least one position. \square

We are now ready to establish an upper bound for the length of the longest path of deals before we converge to an allocation with maximal social welfare:

THEOREM 6 (Longest path with money). *A sequence of rational deals with money can consist of up to $|\mathcal{A}|^{|\mathcal{R}|} - 1$ deals, but not more.*

Proof. There are $|\mathcal{A}|^{|\mathcal{R}|}$ different allocations of resources (each of the resources in \mathcal{R} may be owned by any of the agents in \mathcal{A}). By Lemma 2, there exist utility functions such that all allocations have distinct social welfare. If the initial allocation is the allocation with the lowest social welfare and each deal takes us to the next best allocation, then we get a sequence consisting of exactly $|\mathcal{A}|^{|\mathcal{R}|} - 1$ deals. By Lemma 1, each of these deals is rational. Furthermore, there can be no sequence consisting of more than $|\mathcal{A}|^{|\mathcal{R}|} - 1$ rational deals, because there are only $|\mathcal{A}|^{|\mathcal{R}|}$ different allocations, and every deal has to take us to an allocation with a social welfare that is higher than that of any of the previous allocations. \square

Together with Theorem 1, this means that any sequence of rational deals with money will result in an allocation with maximal social welfare after at most $|\mathcal{A}|^{|\mathcal{R}|} - 1$ steps. Furthermore, this bound is tight, *i.e.* there are cases where exactly $|\mathcal{A}|^{|\mathcal{R}|} - 1$ deals are implemented before the optimal allocation is reached.

Does this bound change if we put restrictions on the class of admissible utility functions? It appears that for many natural restrictions, even very strong ones such as additivity, the upper bound would *not* be affected. As long as it is possible to assign distinct utilities to distinct bundles and there are no restrictions on the overall range of utility values, we can emulate the construction used in the proof of Lemma 2.

4.2. PARETO OPTIMAL OUTCOMES

We now turn our attention to the framework without money. The following lemma will be useful to prove our result concerning the shortest path to a Pareto optimal allocation (the existence of which has been established by Theorem 2).

LEMMA 3 (Concatenating deals). *Let $\delta_1 = (A, A')$ and $\delta_2 = (A', A'')$ be rational deals without money. Then the deal $\delta_3 = (A, A'')$ is also rational without money.*

Proof. The claim follows immediately from Definition 2. \square

Note that an analogue result for rational deals *with* money could easily be proved by reference to Lemma 1.

By Theorem 2, if agents continue to negotiate rational deals without money, then society will eventually converge to a state with a Pareto optimal allocation of resources. The shortest path to such an allocation, again, consists of just a single deal:

THEOREM 7 (Shortest path without money). *A Pareto optimal allocation can always be reached by means of (at most) a single rational deal without money.*

Proof. Given any initial allocation A , by Theorem 2, there exist a Pareto optimal allocation A' and a finite sequence of deals $\langle \delta_1 = (A_0, A_1), \delta_2 = (A_1, A_2), \dots, \delta_n = (A_{n-1}, A_n) \rangle$ such that each of the δ_i is a rational deal without money, $A = A_0$, and $A' = A_n$. By induction over the length n of this sequence and using Lemma 3 in the induction step, it is easy to show that the single deal $\delta = (A_0, A_n)$ will also be rational without money. \square

For the framework with money, we have shown that a sequence of rational deals can consist of up to $|\mathcal{A}|^{|\mathcal{R}|-1}$ individual deals (Theorem 6). We would get the same result for the framework without money *if* it were possible to design utility functions in such a way that for any two allocations either the first is at least as good for all agents and better for some of them or vice versa (but no two allocations are incomparable in this sense). It turns out that this is not the case, *i.e.* we obtain a better bound for the longest path of rational deals in cases where side payments are not allowed:

THEOREM 8 (Longest path without money). *Any sequence of rational deals without money consists of less than $|\mathcal{A}| \cdot (2^{|\mathcal{R}|-1})$ deals.*

Proof. Observe that for any rational deal without money at least one agent needs to make a strict welfare improvement. That agent would certainly have to change the bundle of resources it holds. At no later stage, it could again hold the previous bundle (this very point is different for the framework with money!). Hence, we can compute an upper bound for the number of times any particular agent i will be the one to have the strict improvement: it will be 1 less than the number of possible bundles, *i.e.* $2^{|\mathcal{R}|-1}$. Now, even if every single agent in the system could have a strict improvement that many times, we would get $|\mathcal{A}| \cdot (2^{|\mathcal{R}|-1})$ as an upper bound. Given that for each deal at least *two* agents will change their bundle, this would be rather generous a bound, *i.e.* the maximal length of a sequence of rational deals without money must certainly be less than $|\mathcal{A}| \cdot (2^{|\mathcal{R}|-1})$. \square

The bound of Theorem 8 is not tight: there can be no actual trading scenario where a sequence of $|\mathcal{A}| \cdot (2^{|\mathcal{R}|-1})$ rational deals without money take place. Also observe that $|\mathcal{A}| \cdot (2^{|\mathcal{R}|-1})$ may in fact be more than $|\mathcal{A}|^{|\mathcal{R}|-1}$ (the bound established in Theorem 6) for very small values of

$|\mathcal{A}|$ and $|\mathcal{R}|$. In such cases, clearly, the sharper upper bound of $|\mathcal{A}|^{|\mathcal{R}|-1}$ deals applies as well.

It is possible to show that any precise upper bound for the length of the longest path of rational deals without money would have to be at least $3 \cdot 2^{|\mathcal{R}|} - 2^{|\mathcal{R}|+1-n} - n - 1$ with $n = \min\{|\mathcal{A}|-2, |\mathcal{R}|\}$. To support this claim, we consider the following scenario. Suppose agent 1 has no preferences at all and all other agents assign distinct utility values to the $2^{|\mathcal{R}|-1}$ possible bundles of resources. Furthermore, suppose agent 2 assigns maximal utility to the empty set of resources and minimal utility to the full set, while each of the remaining agents assign maximal utility to some set consisting of only a single resource and minimal utility to the empty set. If $|\mathcal{A}|-2 \leq |\mathcal{R}|$, then suppose this preferred resource is different for each one of them; otherwise suppose that the next $|\mathcal{R}|$ agents have distinct preferred resources. Finally, suppose agent 2 initially holds the full set \mathcal{R} . Define $n = \min\{|\mathcal{A}|-2, |\mathcal{R}|\}$; *i.e.* we always have $n \geq 0$.

We describe a sequence of rational deals without money consisting of $n+1$ phases. In phase 1, agent 1 and 2 implement $2^{|\mathcal{R}|-1}$ deals, each time moving to the next best bundle for agent 2. After this phase, agent 1 owns all resources in the system. The remaining n phases all have the same structure: Just before phase k (for $2 \leq k \leq n+1$), agent 1 owns $|\mathcal{R}|+2-k$ resources. Then agent 1 and agent k implement a sequence of deals such that agent k moves through all the subsets of the resources previously owned by agent 1, moving to the next best bundle in each step. This makes $2^{|\mathcal{R}|+2-k} - 1$ rational deals during phase k . Afterwards, agent 1 owns all the resources it owned at the beginning of that phase, except agent k 's most preferred item. Altogether, the number of deals in the sequence can be computed as follows:

$$\begin{aligned} & (2^{|\mathcal{R}|} - 1) + \sum_{k=2}^{n+1} (2^{|\mathcal{R}|+2-k} - 1) \\ &= (2^{|\mathcal{R}|} - 1) + 2^{|\mathcal{R}|+1} - 2^{|\mathcal{R}|+2-(n+1)} - n \\ &= 3 \cdot 2^{|\mathcal{R}|} - 2^{|\mathcal{R}|+1-n} - n - 1 \end{aligned}$$

(Note that these transformations are correct for any $n \geq 0$.) This confirms our lower bound for the length of the longest possible path of rational deals without money. It is our intuition that this may well be a closer approximation to a precise bound than the proven upper bound of Theorem 8. In particular, it appears that the number of agents in a system has only little influence on this value whenever the number of resources is sufficiently high.

Unlike for the framework with money, now restrictions on utility functions are very likely to improve the upper bound on the longest

path. For instance, a restriction to *monotonic* utility functions (that is, functions such that agents never value a set of items less than any of its subsets) will prevent an agent from accepting a deal where it does not receive at least one new item.

4.3. ADDITIVE SCENARIOS

Theorem 3 shows that, in additive scenarios, one-resource-at-a-time deals (with money) are sufficient to guarantee outcomes of rational negotiation that maximise social welfare. This is certainly a big advantage as far as agreeing on individual deals is concerned, but when restricting ourselves to one-resource-at-a-time deals, we cannot maintain the upper bound on the shortest path of Theorem 5 anymore:

THEOREM 9 (Shortest path in additive scenarios). *If all utility functions are additive, then an allocation with maximal social welfare can always be reached by a sequence of at most $|\mathcal{R}|$ rational one-resource-at-a-time deals with money.*

Proof. Suppose all utility functions are additive. Given an initial allocation A , by Theorem 3, there exists a sequence of rational one-resource-at-a-time deals leading to an allocation A' with maximal social welfare. Consider any resource r with $r \in A(i)$ and $r \in A'(j)$ for two distinct agents $i, j \in \mathcal{A}$. By Definition 1, any such resource r having been transferred must be valued higher by the agent holding it in the final allocation than by the agent holding it at the beginning, also if r has been owned by several different agents at some point during negotiation. That is, we have $u_i(r) < u_j(r)$, *i.e.* the *direct* deal of transferring r from i to j would also be rational. Hence, the number of resources owned by distinct agents in A and A' (at most $|\mathcal{R}|$) is a (tight) upper bound for the shortest path. \square

Our result for the longest path in additive scenarios follows:

THEOREM 10 (Longest path in additive scenarios). *If all utility functions are additive, then a sequence of rational one-resource-at-a-time deals with money can consist of up to $|\mathcal{R}| \cdot (|\mathcal{A}| - 1)$ deals, but not more.*

Proof. In additive scenarios, any rational one-resource-at-a-time deal must reallocate a single resource r to an agent that values r at least slightly higher than its previous owner. Hence, in the worst case, every single resource could be passed through the entire agent society, *i.e.* we obtain a tight upper bound of $|\mathcal{R}| \cdot (|\mathcal{A}| - 1)$. \square

Hence, in additive scenarios it is advantageous to restrict oneself to one-resource-at-a-time deals—also from the viewpoint of reducing the number of deals that have to be implemented in the worst case (besides the obvious advantage of simplifying the task of agreeing on a single deal).

4.4. 0-1 SCENARIOS

Finally, we consider the 0-1 scenarios of Theorem 4. Interestingly, in these scenarios the upper bounds for the shortest and the longest path to an optimal allocation coincide.

THEOREM 11 (Shortest path in 0-1 scenarios). *If all utility functions are 0-1 functions, then an allocation with maximal social welfare can always be reached by a sequence of at most $|\mathcal{R}|$ rational one-resource-at-a-time deals without money.*

Proof. By Theorem 4, an allocation with maximal social welfare can always be reached by *some* sequence of rational one-resource-at-a-time deals without money, provided all utility functions are 0-1 functions. For each $r \in \mathcal{R}$ owned by distinct agents in the initial and the final allocation, the one-resource-at-a-time deal of moving r from the agent owning it at the beginning to the one owning it in the end is rational without money. As up to $|\mathcal{R}|$ items may have to be moved, this is a tight upper bound for the shortest path. \square

THEOREM 12 (Longest path in 0-1 scenarios). *If all utility functions are 0-1 functions, then a sequence of rational one-resource-at-a-time deals without money can consist of up to $|\mathcal{R}|$ deals, but not more.*

Proof. In 0-1 scenarios, any agent receiving a resource r by means of a rational one-resource-at-a-time deal without money must assign utility 1 to that resource. That agent would never agree to give r away again (provided only one-resource-at-a-time deals without money are admissible). Hence, $|\mathcal{R}|$ must be a (tight) upper bound for the number of deals that could possibly be negotiated. \square

5. Conclusion

This paper has addressed the communication complexity of an abstract multilateral trading framework where “myopic” agents negotiate with each other over the reallocation of a number of resources until a socially optimal allocation has been reached. The overall complexity of

Table I. How many rational deals are required to reach an optimal allocation?

<i>Utility functions</i>	general	general	additive	0-1
<i>Side payments?</i>	yes	no	yes	no
<i>Deal types</i>	any	any	1-resource	1-resource
Shortest path	1	1	$ \mathcal{R} $	$ \mathcal{R} $
Longest path	$ \mathcal{A} ^{ \mathcal{R} } - 1$	$< \mathcal{A} \cdot (2^{ \mathcal{R} } - 1)$	$ \mathcal{R} \cdot (\mathcal{A} - 1)$	$ \mathcal{R} $

the framework, we have argued, relates to different (but connected) aspects: (1) the number of deals required to reach an optimal allocation, (2) the number of dialogue moves required to agree on a deal, (3) the expressiveness of the communication language used, and (4) the (computational) complexity of the reasoning task faced by an agent when deciding on its next move.

We have studied *one* of these aspects of complexity in detail, namely the number of deals required for an agent society to converge to an allocation of resources that is socially optimal. In particular, we have given upper bounds on the length of a sequence of rational deals for each of the four instances of our negotiation framework for which we have previously proved existential convergence results in [9].

Table 5 provides a summary of the results obtained in Section 4. For each of the variants of the framework considered, it shows the upper bounds for both the shortest and the longest path to an optimal allocation as a function of $|\mathcal{A}|$, the number of agents, and $|\mathcal{R}|$, the number of resources. With one exception (namely the result for the longest path in general scenarios without money), all these bounds are *tight*, *i.e.* we can find examples where the length of the respective path is exactly as shown in Table 5. As the proofs of our shortest path theorems demonstrate, for the framework with general deals with money, *any* allocation with maximal social welfare can be negotiated (provided we start out with an allocation that is not optimal already), while for the other three frameworks we can only guarantee that *some* optimal allocation is within reach.

We hope to have been able to show that multilateral trading is an exciting area of research that complements recent work on negotiation in multiagent systems, much of which has either concentrated on bilateral exchanges or on auction mechanisms. Future work in this area should also address the other aspects of complexity identified in Section 3 in more detail than has been possible in the present paper. Besides complexity-theoretic issues, there are also a large variety of other interesting research questions connected to the negotiation of

socially optimal allocations of resources that deserve the attention of the multiagent systems research community. These include the compact representation of agent preferences [4, 16]; the investigation of alternative (non-utilitarian) concepts of social welfare as a measure for overall system performance [8]; the study of connections to other combinatorial problems with the aim of designing algorithms to guide negotiation [4]; and the development of heuristics through experimentation [1].

Acknowledgements

The work reported here has been partially supported by the European Commission as part of the SOCS project (IST-2001-32530).

References

1. M. Andersson and T. Sandholm. Time-quality tradeoffs in reallocation negotiation with combinatorial contract types. In *Proc. of the 16th Nat. Conf. on Artificial Intelligence*. MIT Press, 1999.
2. K. J. Arrow, A. K. Sen, and K. Suzumura, editors. *Handbook of Social Choice and Welfare*, volume 1. North-Holland, 2002.
3. A. Chavez, A. Moukas, and P. 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.
4. Y. Chevaleyre, U. Endriss, S. Estivie, and N. Maudet. Multiagent resource allocation with k -additive utility functions. In *Proc. of the DIMACS-LAMSADE Workshop on Computer Science and Decision Theory*, 2004.
5. P. E. Dunne. Extremal behaviour in multiagent contract negotiation. *J. of Artificial Intelligence Research*, 23:41–78, 2005.
6. P. E. Dunne, M. Wooldridge, and M. Laurence. The complexity of contract negotiation. *Artificial Intelligence*, 2005. To appear.
7. U. Endriss and N. Maudet. On the communication complexity of multilateral trading. In *Proc. of the 3rd Intl. Joint Conf. on Autonomous Agents and Multiagent Systems*. ACM Press, 2004.
8. U. Endriss and N. Maudet. Welfare engineering in multiagent systems. In *Engineering Societies in the Agents World IV*. Springer-Verlag, 2004.
9. U. Endriss, N. Maudet, F. Sadri, and F. Toni. On optimal outcomes of negotiations over resources. In *Proc. of the 2nd Intl. Joint Conf. on Autonomous Agents and Multiagent Systems*. ACM Press, 2003.
10. H. Fargier, J. Lang, M. Lemaître, and G. Verfaillie. Partage équitable de ressources communes: (2) Éléments de complexité et d’algorithmique. *Technique et Science Informatique*, 23(9):1219–1238, 2004.
11. A. M. Feldman. Bilateral trading processes, pairwise optimality, and Pareto optimality. *The Review of Economic Studies*, 40(4):463–473, 1973.
12. M. R. Garey and D. S. Johnson. *Computers and Intractability: A Guide to the Theory of NP-completeness*. W. H. Freeman and Company, New York, 1979.

13. 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.
14. S. Kraus. *Strategic Negotiation in Multiagent Environments*. MIT Press, 2001.
15. E. Kushilevitz and N. Nisan. *Communication Complexity*. Cambridge University Press, 1996.
16. J. Lang. Logical preference representation and combinatorial vote. *Annals of Mathematics and Artificial Intelligence*, 42(1-3):37-71, 2004.
17. A. Lomuscio, M. Wooldridge, and N. R. Jennings. A classification scheme for negotiation in electronic commerce. In *Agent-Mediated Electronic Commerce: The European AgentLink Perspective*, LNCS. Springer Verlag, 2001.
18. H. Moulin. *Axioms of Cooperative Decision Making*. Cambridge University Press, 1988.
19. J. S. Rosenschein and G. Zlotkin. *Rules of Encounter*. MIT Press, 1994.
20. T. W. Sandholm. Contract types for satisficing task allocation: I Theoretical results. In *Proc. of the AAAI Spring Symposium: Satisficing Models*, 1998.
21. T. W. Sandholm. Distributed rational decision making. In *Multiagent Systems: A Modern Approach to Distributed Artificial Intelligence*. MIT Press, 1999.
22. A. K. Sen. *Collective Choice and Social Welfare*. Holden Day, 1970.
23. T. Suyama and M. Yokoo. Strategy/false-name proof protocols for combinatorial multi-attribute procurement auction. In *Proc. of the 3rd Intl. Joint Conf. on Autonomous Agents and Multiagent Systems*. ACM Press, 2004.
24. P. Winoto, G. McCalla, and J. Vassileva. Non-monotonic-offers bargaining protocol. In *Proc. of the 3rd Intl. Joint Conf. on Autonomous Agents and Multiagent Systems*. ACM Press, 2004.
25. M. Wooldridge. Semantic issues in the verification of agent communication languages. *J. of Autonomous Agents and Multi-Agent Systems*, 3(1):9-31, 2000.
26. M. Wooldridge. *An Introduction to MultiAgent Systems*. Wiley, 2002.
27. A. C.-C. Yao. Some complexity questions related to distributive computing (preliminary report). In *Proc. of the Annual ACM Symposium on Theory of Computing*, 1979.

