The Complexity of Contract Negotiation

Paul E. Dunne * Michael Wooldridge Michael Laurence

Dept. of Computer Science, University of Liverpool, Liverpool L69 7ZF, United Kingdom

Abstract

The use of agent systems as a means of implementing contract negotiation in e-commerce and e-trading environments has been the focus of considerable recent interest. A widely studied abstract model considers the setting in which a set of agents have some collection of resources shared out between them and attempt to construct a mutually beneficial optimal reallocation of these by trading resources. The simplest such trades are those in which a single agent transfers exactly one resource to another – so-called 'one-resource-at-a-time' or '*O-contracts*'. In this research note we consider the computational complexity of a number of natural decision problems in this setting.

Key words: Negotiation, Multiagent systems, Computational Complexity

1 Introduction

Mechanisms for negotiating allocation of resources within a group of agents form an important body of work within the study of multiagent systems. Typical abstract models derive from game-theoretic perspectives in economics and among the issues that have been addressed are strategies that agents may use to negotiate, e.g. [5,8,10], and protocols for negotiation in agent societies, e.g. [1,6].

The scenario we are concerned with is encapsulated in the following definition.

Definition 1 *Let $\mathcal{A} = \{A_1, A_2, \ldots, A_n\}$ be a set of (at least two) agents and $\mathcal{R} = \{r_1, r_2, \ldots, r_m\}$ a collection of (non-shareable) resources. A* utility function*, u, is a mapping from subsets of $\mathcal{R}$ to rational values. Each agent $A_i \in \mathcal{A}$ has associated with it a particular utility function u_i. An* allocation *of $\mathcal{R}$ to $\mathcal{A}$ is a partition $P = \langle P_1, P_2, \ldots, P_n \rangle$ of $\mathcal{R}$. The value $u_i(P_i)$ is called the* utility *of the resources assigned to A_i.*

* Corresponding author.
Email address: `ped@csc.liv.ac.uk` (Paul E. Dunne).

Starting from some initial allocation – P_0 – individual agents negotiate in an attempt to improve the utility of their holding. A number of interpretations have been proposed in order to define what constitutes a 'sensible' transfer of resource from both an individual agent's viewpoint and from the perspective of the overall allocation. Thus in negotiating a change from an allocation P_i to Q_i (with $P_i, Q_i \subseteq \mathcal{R}$ and $P_i \neq Q_i$) there are three possible outcomes for the agent A_i: $u_i(P_i) < u_i(Q_i)$, i.e. A_i values the allocation Q_i as superior to P_i; $u_i(P_i) = u_i(Q_i)$, i.e. A_i is indifferent between P_i and Q_i; and $u_i(P_i) > u_i(Q_i)$, i.e. A_i is worse off after the exchange. In a setting where agents are seen as self-interested, in order for an agent to accept an exchange with the last outcome, the notion of a *pay-off* function is used, i.e. in order to accept the new allocation, A_i receives some payment sufficient to compensate for the resulting loss in utility. Of course such compensation must be made by other agents in the system who in providing it do not wish to pay in excess of any gain in resource. In defining notions of 'pay-off' the interpretation is that in any transaction each agent A_i makes a payment, π_i: if $\pi_i < 0$ then A_i is given $-\pi_i$ in return for accepting a contract; if $\pi_i > 0$ then A_i contributes π_i to the amount to be distributed among those agents whose pay-off is negative. Formally, such a notion of 'sensible transfer' is captured by the concept of *individual rationality*.

Definition 2 *Let $\mathcal{A}$ and $\mathcal{R}$ be as in Definition 1. A* deal *is a pair $\langle P, Q\rangle$ where $P = \langle P_1, \ldots, P_n\rangle$ and $Q = \langle Q_1, \ldots, Q_n\rangle$ are distinct partitions of $\mathcal{R}$. The effect of implementing the deal $\langle P, Q\rangle$ is that the allocation of resources specified by P is replaced with that specified by Q.*

A deal $\langle P, Q\rangle$ is said to be individually rational *(IR) if there is a* pay-off vector *$\pi = \langle \pi_1, \pi_2, \ldots, \pi_n\rangle$ satisfying,*

a. $\sum_{i=1}^{n} \pi_i = 0$.
b. $u_i(Q_i) - u_i(P_i) > \pi_i$*, for each agent A_i,* except that π_i *may be* 0 *if $P_i = Q_i$, i.e. should the deal (P, Q) leave the agent A_i with no change in its resource then it is not* required *that A_i be rewarded (have $\pi_i < 0$).*

Definition 2 captures one view of a deal being 'sensible' with respect to the perspective of single agents. We require also concepts of '*global*' optimality. We consider two commonly used versions of this: Pareto Optimality and Social Welfare.

Definition 3 *Let P be an allocation of $\mathcal{R}$ among $\mathcal{A}$. The* social welfare *resulting from P, denoted $\sigma(P)$, is given by $\sum_{i=1}^{n} u_i(P_i)$.*

The allocation P is Pareto optimal *if for all allocations Q differing from P, it holds*

$$\left(\bigvee_{i=1}^{n} [u_i(Q_i) > u_i(P_i)]\right) \Rightarrow \left(\bigvee_{i=1}^{n} [u_i(Q_i) < u_i(P_i)]\right) \quad (1)$$

Thus a Pareto optimal allocation is one in which no agent can attain better than its current utility except at the cost of leaving some agent worse off.

A result linking social welfare and individual rationality, which we make frequent use of subsequently, is given in recent work of Endriss *et al.*[4, Lemma 1]

Fact 4 *A deal $\langle P, Q\rangle$ is IR if and only if $\sigma(Q) > \sigma(P)$.*

In a typical application it is unlikely that an initial allocation P_0 to $\mathcal{A}$ will either maximise social welfare or be Pareto optimal, thus the agents involved seek to find a sequence of deals that will terminate in an optimal allocation. Given the setting it is clearly the case that there are allocations P_{opt} and Q_{opt} with the properties that $\sigma(P_{opt})$ maximises social welfare and for which Q_{opt} is Pareto optimal – of course, P_{opt} and Q_{opt} may not be unique. If the object is to maximise social welfare then clearly the deal $\langle P_0, P_{opt}\rangle$ will achieve this in a single round. It is unreasonable, however, to view such a deal as a viable solution: although always IR (if it represents a strict increase of social welfare) it is questionable whether it could be identified as the *first and only* deal required. The total number of possible allocations is n^m, and so for moderately large numbers of resources (m) there are too many feasibly to enumerate (even when $n = 2$). In addition it may not be possible to implement the optimising contract in a *single* transaction even if only two agents are involved: the environment in which the trading process is implemented may not be suited to handling transactions in which large numbers of resources are involved; similarly the protocol used for negotiation and contract description may not allow arbitrarily large numbers of resources to be dealt with.

In order to develop a realistic framework for negotiation, Sandholm [9] using Smith's Contract-Net model [11] presents a number of classes of *contract type*. We review these in the definition below.

Definition 5 *([9]) Let $D = \langle P, Q\rangle$ be a deal involving an allocation of $\mathcal{R}$ among $\mathcal{A}$. We say that D is a* cluster contract *(C-contract) if there are distinct agents A_i and A_j for which,*

C1) $P_k = Q_k$ if and only if $k \notin \{i, j\}$.
C2) There is a unique (non-empty) set S for which $Q_i = P_i \cup S$ and $Q_j = P_j - S$ (with $S \subseteq P_j$) or $Q_j = P_j \cup S$ and $Q_i = P_i - S$ (with $S \subseteq P_i$)

Thus a C-contract involves one agent transferring a subset of its allocation to another agent (without receiving any subset of resources in return).

We say $\langle P, Q\rangle$ is a one contract *(O-contract) if $\langle P, Q\rangle$ is a C-contract in which S – the set of resources transferred – contains* exactly one *element.*

We recall, from Fact 4, that an X-contract $\langle P, Q\rangle$ (for $X \in \{O, C\}$) will be IR if and only if $\sigma(Q) > \sigma(P)$.

A sequence of deals $\langle\langle Q_0, Q_1\rangle, \langle Q_1, Q_2\rangle, \ldots, \langle Q_i, Q_{i+1}\rangle, \ldots, \langle Q_{t-1}, Q_t\rangle\rangle$ is called a *contract path* realising the deal $\langle Q_0, Q_t\rangle$. A contract path is an *IR-path* if every deal within it is IR.

Contract paths in which the constituent deals are O-contracts are attractive from an implementation viewpoint since these only involve agent-to-agent negotiation concerning a single resource at a time. In addition, starting from a given allocation the number of O-contracts that are consistent with it is exactly $m(n-1)$ as opposed to n^m possible allocations, thus heuristic methods may be able to find improved allocations by exploring the search space through O-contracts alone.

Appealing as the latter approach is there are, nevertheless, problems associated with it. The following results were established by Sandholm [9].

Fact 6 *Let P_0 be any initial allocation of $\mathcal{R}$ to $\mathcal{A}$, P_t any other allocation, and P_{opt} any allocation which maximises social welfare.*

a) The deal $\langle P_0, P_t\rangle$ can always be realised by a contract path in which every deal is an O-contract.
b) In some instances the choice of utility function is such that there is no IR C-contract path realising $\langle P_0, P_{opt}\rangle$.

Thus O-contracts in themselves may not suffice to form an IR contract-path realising a specific deal. In this paper we are concerned with the following decision problem:

Definition 7 *The decision problem* IR-O-path *(*IRO*) is given by*
Instance: *A 5-tuple $\langle\mathcal{A}, \mathcal{R}, \mathcal{U}, P^s, P^t\rangle$ in which $\mathcal{A}$ and $\mathcal{R}$ are as in Definition 1, P^s and P^t are allocations of $\mathcal{R}$ to $\mathcal{A}$ in which $\sigma(P^t) > \sigma(P^s)$ and $\mathcal{U} = \langle u_1, \ldots, u_n\rangle$ are the utility functions associated with $\mathcal{A}$.*
Question: *Is there an IR O-contract path that realises the deal (P^s, P^t)?*

The main result of this article is that IRO is NP–hard: thus given two distinct allocations there are unlikely to be efficient algorithms determining if one can be reached from the other through a sequence of IR O-contracts.

The next section of this article presents these results with conclusions and open questions raised in the final section.

2 Complexity Results

Before proceeding with our results we describe our representation for typical instances of IRO. The key issue here concerns the collection of utility functions $\mathcal{U}$ and how these should be encoded. A form in which the value attached to each subset of $\mathcal{R}$ is explicitly provided will result in an instance occupying space exponential in $|\mathcal{R}|$ and would not be considered reasonable in practice. On the other hand, using some encoding of $\mathcal{U}$ as a set of *Turing machine programs*, $\mathcal{M}$ say, it becomes necessary to assume certain properties in interpreting their computational behaviour, e.g. that the value of $u_i(S)$ as returned by the program M_i is defined from the content of M_i's tape after exactly some specified number of moves such as $|\mathcal{R}|$ since without such it would not be possible to establish membership in NP (or, indeed, any other complexity class). We represent each member of $\mathcal{U}$ in a manner that does not *require* explicit enumeration of each subset of $\mathcal{R}$, uses a 'program' form whose syntactic correctness can be efficiently verified and for which termination in time linear in the program length is guaranteed. The class of programs employed are the so-called *straight-line programs* which have a natural correspondence with combinational logic networks [2].

Definition 8 *An* (m,s)–combinational network *C is a directed acylic graph in which there are m* input nodes*, Z_m, labelled $\langle z_1, z_2, \ldots, z_m\rangle$ all of which have in-degree* 0*. In addition, C has s* output nodes*, called the* result vector*. These are labelled $\langle t_{s-1}, t_{s-2}, \ldots, t_0\rangle$, and have out-degree 0. Every other node of C has in-degree at most 2 and out-degree at least 1. Each non-input node (*gate*) is associated with a Boolean operation of at most two arguments.* [1] *We use $|C|$ to denote the number of* gate *nodes in C. Any Boolean instantiation of the input nodes to $\alpha \in \langle 0,1\rangle^m$ naturally induces a Boolean value at each gate of C: if h is a gate associated with the operation θ, and $\langle g_1, h\rangle$, $\langle g_2, h\rangle$ are edges of C then the value $h(\alpha)$ is $g_1(\alpha)\theta g_2(\alpha)$. Hence α induces some s-tuple $\langle t_{s-1}(\alpha), \ldots, t_0(\alpha)\rangle \in \langle 0,1\rangle^s$ at the result vector. For the (m,s)-combinational network C and $\alpha \in \langle 0,1\rangle^m$, this s-tuple is denoted by $C(\alpha)$.*

Although often considered as a model of parallel computation, (m,s)-combinational networks yield a simple form of sequential program – straight-line programs – as follows. Let C be an (m,s)-combinational network to be transformed to a straight-line program, $\mathrm{SLP}(C)$, that will contain exactly $m + |C|$ lines. Since C is directed and acyclic it may be topologically sorted, i.e. each gate, g, given a unique integer label, $\tau(g)$ with $1 \leq \tau(g) \leq |C|$ so that if $\langle g, h\rangle$ is an edge of C then $\tau(g) < \tau(h)$. The line l_i of $\mathrm{SLP}(C)$ evaluates the input z_i if $1 \leq i \leq m$ and the gate for which $\tau(g) = i - m$ if $i > m$. The gate labelling means that when g with inputs g_1 and g_2 is evaluated at $l_{m+\tau(g)}$ since g_i is either an input node or another gate its value will have been determined at l_j with $j < m + \tau(g)$.

Definition 9 *Let $\mathcal{R}$ be as previously with $|\mathcal{R}| = m$, and u a mapping from subsets of $\mathcal{R}$ to rational values, i.e. a utility function. The (m,s)–network C^u is said to* realise *the utility function u if: for every $S \subseteq R$ with α_S the instantiation of Z_m by*

[1] In practice, we can restrict the Boolean operations employed to those of binary conjunction ($\wedge$), binary disjunction ($\vee$) and unary negation ($\neg$).

$z_i = 1 \Leftrightarrow r_i \in S$, *it holds*

$$u(S) = \frac{val(C(\alpha_S))}{m}$$

where for $\beta = \langle \beta_{s-1}, \beta_{s-2}, \ldots, \beta_0 \rangle \in \langle 0, 1 \rangle^s$, $val(\beta)$ *is the whole number*[2] *whose s-bit binary expansion is* β, *i.e.* $val(\beta) = \sum_{i=0}^{s-1} \beta_i * 2^i$, *where* β_i *is regarded as the appropriate integer value from* $\{0, 1\}$.

These ideas allow any utility function u_i in $\mathcal{U}$ to be encoded using an appropriate (m, s_i)-combinational network, $C^{(i)}$ in such a way that $u_i(S)$ can be evaluated in time linear in the number of nodes in $C^{(i)}$ by determining the value of each gate under the related instantiatiation α_S and then dividing this value by m.

The decision problem IRO is concerned with the existence of a suitable contract path from one allocation to another having greater social welfare. For completeness it is useful to present three easy proofs concerning the existence of resource allocations meeting particular criteria. These problems are respectively,

Welfare Improvement (WI)
Instance: A tuple $\langle \mathcal{A}, \mathcal{R}, \mathcal{U}, P \rangle$ where $\mathcal{A}$, $\mathcal{R}$ and $\mathcal{U}$ are as before, and P is an allocation.
Question: Is there an allocation Q for which $\sigma(Q) > \sigma(P)$?

Welfare Optimisation (WO)
Instance: A tuple $\langle \mathcal{A}, \mathcal{R}, \mathcal{U}, K \rangle$ where $\mathcal{A}$, $\mathcal{R}$ and $\mathcal{U}$ are as before, and K is a rational number.
Question: Is there an allocation P for which $\sigma(P) \geq K$?

Pareto Optimal (PO)
Instance: A tuple $\langle \mathcal{A}, \mathcal{R}, \mathcal{U}, P \rangle$ as for WI.
Question: Is the allocation P Pareto optimal?

Kraus[5, p. 43] proves NP-hardness of a weaker form of the problem WO whereby in addition to the total social welfare having to attain some specified value the allocation must be such that each agent accrues some designated guaranteed utility.

Theorem 10 *Even if n (the number of agents) is* 2 *and the utility functions are monotonic, i.e.* $u(S) \geq u(R)$ *whenever* $R \subseteq S$,

a) WI *is* NP*–complete.*

[2] Although this definition assumes utility functions to have non-negative values, were it the case that some function with $u(S) < 0$ was to be represented we can achieve this by using an additional output bit, $t_\pm$ to flag whether $val(C(\alpha))$ should be treated as positive ($t_\pm = 0$) or negative ($t_\pm = 1$).

b) WO *is* NP*–complete.*
c) PO *is* CO-NP*–complete.*

Proof: We first demonstrate that the three problems are in the classes stated, recalling that the utility functions $\mathcal{U}$ are encoded by (m, s_i)-combinational networks $C^{(i)}$ as described in Defn 9. For (a), given an instance $\langle \mathcal{A}, \mathcal{R}, \mathcal{U}, P\rangle$ of WI simply non-deterministically guess an allocation $Q = \langle Q_1, \ldots, Q_n\rangle$ and compute

$$\sigma(Q) = \sum_{i=1}^{n} \frac{val(C^{(i)}(\alpha_{Q_i}))}{|\mathcal{R}|}$$

accepting if this exceeds $\sigma(P)$. For (b) a similar approach is used with an instance accepted if the guessed allocation Q has $\sigma(Q) \geq K$. Finally, for (c) we may use a CO-NP algorithm to check that for all allocations Q the Pareto Optimality condition given in Definition 3(1) holds.

We complete the proof by establishing NP–hardness for WI, WO and CO-NP-hardness for PO.

For part (a) we use a reduction from 3-SAT, instances of which are propositional formulae $\Phi(X_n)$ in conjunctive normal form with each clause of Φ defined by exactly three literals. We further restrict instances to formulae that contain *exactly* n clauses, a variant shown to be NP–complete in [3, Theorem 2(b)]. Let

$$\Phi(X_n) = \bigwedge_{i=1}^{n} C_i \;=\; \bigwedge_{i=1}^{n} (y_{i,1} \;\vee\; y_{i,2} \vee\; y_{i,3})$$

be an instance of this problem, where $y_{i,j}$ is some literal x_k or $\neg x_k$.

Given $\Phi(X_n)$ we construct an instance $\langle\{A_1, A_2\}, \mathcal{R}, \langle u_1, u_2\rangle, P\rangle$ in which

a) $\mathcal{R} = \{x_1, x_2, \ldots, x_n, \neg x_1, \ldots, \neg x_n, C_1, \ldots, C_n\}$
b) $P = \langle \emptyset\; ;\; \mathcal{R}\rangle$.

We say that a subset S of $\mathcal{R}$ is *useful* if the set

$$Lits(S) = S \cap \{x_1, x_2, \ldots, x_n, \neg x_1, \ldots, \neg x_n\}$$

satisfies

1. For each $1 \leq k \leq n$, $Lits(S)$ contains *at most one* of the literals x_k, $\neg x_k$.
2. The instantiation by which each $y \in Lits(S)$ is assigned true is a satisfying assignment of $\Phi(X_n)$.

For utility functions we use,

$$u_1(S) = \begin{cases} 0 & \text{if } \ S = \emptyset \\ \frac{|S|+1}{3n} & \text{if } \ S \text{ is useful} \\ \frac{|S|}{3n} & \text{if } \ S \text{ is } \textit{not} \text{ useful} \end{cases}$$

$$u_2(S) = \begin{cases} 2 & \text{if } \ S = \mathcal{R} \\ 1 + \frac{|S|}{3n} & \text{if } \ \mathcal{R} - S \text{ is useful} \\ 1 + \frac{|S|-1}{3n} & \text{if } \ \mathcal{R} - S \text{ is not useful} \end{cases}$$

Both of these function are monotonic in the sense given by the theorem statement. Furthermore given $\Phi(X_n)$ we may construct the combinational networks $C^{(1)}$ and $C^{(2)}$ as follows. Let the inputs for each network be $\langle z_1, \ldots, z_{3n} \rangle$ with z_i set to represent the presence of x_i (if $i \leq n$), the presence of $\neg x_{i-n}$ (if $n < i \leq 2n$) and the presence of C_{i-2n} if ($2n < i \leq 3n$).

For $C^{(1)}$ we simply use a combinational network that computes the binary representation of $\mathit{Useful}(Z_{2n}) + \sum_{i=1}^{3n} z_i$ where

$$\mathit{Useful}(Z_{2n}) = \bigwedge_{i=1}^{n} (\neg z_i \vee \neg z_{n+i}) \ \wedge \ \bigwedge_{i=1}^{n} (z_{i,1} \vee z_{i,2} \vee z_{i,3})$$

Here, $z_{i,j}$ is the variable from $\{z_1, \ldots, z_{2n}\}$ corresponding to the literal $y_{i,j}$ of clause C_i.

The summation to compute the binary representation of the number of bits set to 1 within Z_{3n} can be carried out using the using the schema of Muller and Preparata[7], see e.g. [2, pp. 112–4]. The whole number $\mathit{val}(C^1(\alpha_S))$ computed will be $|S|$, i.e. the number of variables set to 1 in α_S, if S is empty or not useful; and $|S| + 1$ if S is useful.

For $C^{(2)}$, a combinational network computes the binary representation of

$$\sum_{i=1}^{3n-1} 1 \ + \ \bigwedge_{i=1}^{3n} z_i \ + \ \sum_{i=1}^{3n} z_i \ + \ \mathit{Useful}(\neg z_1, \ldots, \neg z_n, \neg z_{n+1}, \ldots, \neg z_{2n})$$

For $S \subseteq \mathcal{R}$, this will return $val(C^{(2)}(\alpha_S))$ as

$$\begin{array}{lll} 6n & = 3n-1+1+3n+0 & \text{when } S=\mathcal{R} \\ 3n+|S| & = 3n-1+0+|S|+1 & \text{when } \mathcal{R}-S \text{ is useful} \\ 3n+|S|-1 & = 3n-1+0+|S|+0 & \text{when } \mathcal{R}-S \text{ is not useful} \end{array}$$

It is clearly the case that these descriptions can be constructed in polynomial-time from the formula $\Phi(X_n)$.

Now, noting that $\sigma(\langle\emptyset;\mathcal{R}\rangle) = 2$, we claim that there is an allocation, Q, having $\sigma(Q) > 2$ if and only if $\Phi(X_n)$ is satisfiable. To see this consider any non-empty $S \subseteq \mathcal{R}$ and the allocation $\langle S, \mathcal{R}-S\rangle$ to $\langle A_1, A_2\rangle$. We have,

$$\sigma(\langle S,\mathcal{R}-S\rangle) = \begin{cases} \frac{|S|+1}{3n}+1+\frac{|\mathcal{R}-S|}{3n} & \text{if } S \text{ is useful} \\ \frac{|S|}{3n}+1+\frac{|\mathcal{R}-S|-1}{3n} & \text{otherwise} \end{cases}$$

In the former case we get, $\sigma(\langle S,\mathcal{R}-S\rangle) = 2+1/(3n)$ and, in the latter, $\sigma(\langle S,\mathcal{R}-S\rangle) = 2-1/(3n)$. Thus the allocation $\langle\emptyset,\mathcal{R}\rangle$ is welfare improvable if and only if there is an allocation S to A_1 that is useful: a condition that requires $Lits(S)$ to induce a satisfying instantiation of $\Phi(X_n)$, completing the proof that WI is NP–hard.

For part (b) we simply form the instance, $\langle\{A_1,A_2\},\mathcal{R},\langle u_1,u_2\rangle,K\rangle$ with $\mathcal{R}$, $\langle u_1,u_2\rangle$ as in part (a) and $K = 2+1/(3n)$.

Finally for part (c), we use $\mathcal{R}$ and $\langle u_1,u_2\rangle$ as previously, but set $P = \langle P_1,P_2\rangle = \langle\{C_1,\ldots,C_n\},\{x_1,\ldots,x_n,\neg x_1,\ldots,\neg x_n\}\rangle$. In this case we have $u_1(P_1) = 1/3$ and $u_2(P_2) = 1+(2n-1)/(3n)$, so that $\sigma(P) = 2-1/(3n)$. We claim that this allocation is Pareto optimal if and only if $\Phi(X_n)$ is *unsatisfiable*. First suppose $\Phi(X_n)$ is unsatisfiable. Certainly for any allocation $Q = \langle S,\mathcal{R}-S\rangle$ differing from $\langle P_1,P_2\rangle$, it must be the case that $S=\emptyset$ or S is not useful. In the former case,

$$u_1(\emptyset) = 0 < u_1(P_1) = \frac{1}{3}$$

so that the Pareto Optimality condition of Definition 3(1) holds for $\langle P_1,P_2\rangle$ with respect to $\langle\emptyset,\mathcal{R}\rangle$.

If S is non-empty then

$$\sigma(\langle S,\mathcal{R}-S\rangle) = u_1(S)+u_2(\mathcal{R}-S) = 2-\frac{1}{3n}$$

and so does not increase social welfare. It follows that, in this case,

$$([u_1(S) > u_1(P_1)] \ \vee \ [u_2(\mathcal{R} - S) > u_2(P_2)])$$
$$\Rightarrow$$
$$([u_1(S) < u_1(P_1)] \ \vee \ [u_2(\mathcal{R} - S) < u_2(P_2)])$$

Hence if $\Phi(X_n)$ is unsatisfiable then P is Pareto optimal. On the other hand suppose $\Phi(X_n)$ *is* satisfiable. We can then demonstrate that P is *not* Pareto optimal by considering any set of literals $\{y_1, \ldots, y_n\}$ whose instantiation to true satisfies Φ. With such a set consider the allocation

$$Q = \langle Q_1, Q_2 \rangle = \langle \{y_1, \ldots, y_n\}, \{\neg y_1, \ldots, \neg y_n, C_1, \ldots, C_n\} \rangle$$

Certainly Q_1 is useful, therefore

$$u_1(Q_1) = \frac{n+1}{3n} > u_1(P_1)$$

$$u_2(Q_2) = 1 + \frac{2}{3} > u_2(P_2)$$

We deduce that the allocation P is Pareto optimal if and only if $\Phi(X_n)$ is unsatisfiable. □

Our main result is,

Theorem 11 IRO *is* NP*–hard*[3] *even when n, the number of agents, equals* 2 *and the utility functions are monotonic.*

Proof: We, again, employ a reduction from 3–SAT but with no restrictions on the number of clauses. Let

$$\Phi(X_n) = \bigwedge_{i=1}^{m} C_i \ = \ \bigwedge_{i=1}^{m} (y_{i,1} \ \vee \ y_{i,2} \vee \ y_{i,3})$$

be an instance of 3–SAT, where $y_{i,j}$ is some literal x_k or $\neg x_k$.

[3] The algorithm – guess a sequence of O-contracts realising $\langle P^s, P^t \rangle$ and check whether this defines an IR contract path – may fail to be realisable by an NP computation. Even when such a sequence exists, [9, Propn. 2] gives an upper bound on its length which is exponential in $|\mathcal{R}|$. In the absence of a polynomial upper bound it is unclear whether membership in NP can be established.

We construct an instance $T_\Phi = \langle \mathcal{A}, \mathcal{R}, \mathcal{U}, P^s, P^t \rangle$ of IRO which is accepted if and only if $\Phi(X_n)$ is satisfiable. Furthermore $|\mathcal{A}| = 2$ for this instance.

Given $\Phi(X_n)$, the set $\mathcal{A}$ of agents in T_Φ contains two members which we denote by A_{eval} and A_{coll}. The set, $\mathcal{R}$, of resources allocated between these is,

$$\mathcal{R} = \{C_1, C_2, \ldots, C_m, x_1, x_2, \ldots, x_n, \neg x_1, \neg x_2, \ldots, \neg x_n\}$$

i.e. the m clauses of $\Phi(X_n)$ together with the $2n$ literals over X_n. The source and destination allocations, P^s and P^t, are

$$\begin{aligned} P^s &= \langle P^s_{eval}; P^s_{coll} \rangle = \langle \{C_1, C_2, \ldots, C_m\} \; ; \; \{x_1, \ldots, x_n, \neg x_1, \ldots, \neg x_n\} \rangle \\ P^t &= \langle P^t_{eval}; P^t_{coll} \rangle = \langle \{x_1, \ldots, x_n, \neg x_1, \ldots, \neg x_n\} \; ; \; \{C_1, C_2, \ldots, C_m\} \rangle \end{aligned}$$

It remains to define the utility functions u_{eval} and u_{coll} for each agent. The idea underlying the construction of these is that moving literals from A_{coll} to A_{eval} by O-contracts, will *only* be IR if at some stage those literals held by A_{eval} define a satisfying instantiation of $\Phi(X_n)$ (by choosing values for the variables which make the corresponding literals true). We note that any $S \subseteq \mathcal{R}$ consists of some subset Y_S of literals together with some subset C_S of clauses. We consider the functions u_{eval} and u_{coll} in terms of this partition.

If $Y_{eval} \cup C_{eval}$ is the subset of $\mathcal{R}$ allocated to A_{eval}, then $u_{eval}(Y_{eval} \cup C_{eval})$ is,

e1) If for some i, both x_i and $\neg x_i$ belong to Y_{eval} and $C_{eval} \neq \emptyset$.

$$u_{eval}(Y_{eval} \cup C_{eval}) = 1 + \frac{|Y_{eval}| - 1}{2n + m}$$

e2) Otherwise,

$$u_{eval}(Y_{eval} \cup C_{eval}) = 1 + \frac{|Y_{eval}|}{2n + m}$$

Similarly, the value of $u_{coll}(Y_{coll} \cup C_{coll})$ is,

c1) If $C_{eval} \neq \emptyset$ and the instantiation formed by setting each literal in Y_{eval} is either undefined (since both z and $\neg z$ are in Y_{eval}) or fails to satisfy each clause in C_{coll}

$$u_{coll}(Y_{coll} \cup C_{coll}) = 1 + \frac{|C_{coll}| - 1}{2n + m}$$

c2) If $C_{eval} = \emptyset$, $|Y_{eval}| \leq n$, and the instantiation formed by setting each literal in

Y_{eval} is either undefined or fails to satisfy $\Phi(X_n)$.

$$u_{coll}(Y_{coll} \cup C_{coll}) = 1 + \frac{|C_{coll}| - 1}{2n + m}$$

c3) Otherwise,

$$u_{coll}(Y_{coll} \cup C_{coll}) = 1 + \frac{|C_{coll}|}{2n + m}$$

We observe that suitable combinational networks can be constructed efficiently from $\Phi(X_n)$ in a manner similar to that used in the proof of Theorem 10.

We claim that $\Phi(X_n)$ is satisfiable if and only if there is an IR O-contract path realising the deal

$$\langle P^s, P^t\rangle = \langle\langle\{C_1, \ldots, C_m\}, \{x_1, \ldots, \neg x_n\}\rangle, \langle\{x_1, \ldots, \neg x_n\}, \{C_1, \ldots, C_m\}\rangle\rangle$$

First suppose that $\Phi(X_n)$ is satisfiable and let $\langle y_1, \ldots, y_n\rangle$ be a set of literals the instantiation of each to true satisfies $\Phi(X_n)$. Consider the following sequence of O-contracts in which, to start, $i = 1$.

1. A_{coll} transfers the resource y_i to A_{eval}
2. If there are k_i clauses in A_{eval} that contain the literal y_i, then the next k_i O-contracts involve transferring each of these distinct clauses to A_{coll}.
3. If $i < n$, the sequence continues from (1) using the literal y_{i+1}.

After all n literals have been transferred from A_{coll} to A_{eval} and the satisfied clauses transferred to A_{coll}, P_{eval} – the resources assigned to A_{eval} – will be such that $Y_{eval} = \{y_1, \ldots, y_n\}$ and $C_{eval} = \emptyset$. Similarly, P_{coll} will have $Y_{coll} = \{\neg y_1, \ldots, \neg y_n\}$ and $C_{coll} = \{C_1, \ldots, C_m\}$. The final sequence of O-contracts consists of n deals in which the literal $\neg y_i$ is transferred from A_{coll} to A_{eval}. At the end of this sequence Y_{eval} will contain $\{x_1, \ldots, x_n, \neg x_1, \ldots, \neg x_n\}$ and C_{coll} will equal $\{C_1, \ldots, C_m\}$, i.e. the allocation described by P^t in the instance. We claim that every O-contract in the sequence described above is IR for which it suffices, from Fact 4, to show that each results in a strict increase in social welfare.

Consider the value of $\sigma(P^s)$. Noting the convention that an *empty* set of clauses is always satisfied, i.e. equivalent to the Boolean value **true**,

$$\sigma(P^s) = u_{eval}(\{C_1, \ldots, C_m\}) + u_{coll}(\{x_1, \ldots, \neg x_n\}) = 1 + 1 = 2$$

We first note that the first iteration through steps (1) and (2) – the transfer of y_1 followed by the satisfied clauses – results in increased social welfare.

$$u_{eval}(\{y_1\} \cup \{C_1, \ldots, C_m\}) = 1 + \tfrac{1}{2n+m} > 1$$
$$u_{coll}(\{x_1, \ldots, \neg x_n\} - \{y_1\}) = 1$$

The next k_1 O-contracts are transfers of single clauses (containing y_1) from C_{eval} to C_{coll}: each of these leaves the utility of A_{eval} unchanged while increasing that of A_{coll} by $1/(2n+m)$, hence these O-contracts are IR.

Now consider the allocations P^i_{eval} and P^i_{coll} to A_{eval} and A_{coll} after exactly i iterations of these steps (for $i < n$). We have

$$Y_{eval} = \{y_1, \ldots, y_i\}$$
$$C_{eval} = \{C_1, \ldots, C_m\} - \{C_k \ : y_j \in C_k \text{ for some } j \leq i\}$$
$$Y_{coll} = \{x_{i+1}, \ldots, x_n, \neg x_{i+1}, \ldots, \neg x_n, \neg y_1, \ldots, \neg y_i\}$$
$$C_{coll} = \{C_k \ : y_j \in C_k \text{ for some } j \leq i\}$$

We have shown already that $\sigma(\langle P^1_{eval}, P^1_{coll}\rangle) > \sigma(P^s)$ and that each of the O-contracts involved is IR. Letting $\langle P^0_{eval}, P^0_{coll}\rangle$ denote the initial allocation, inductively assume that for some value $1 \leq i < n$.

$$\sigma(\langle P^i_{eval}, P^i_{coll}\rangle) > \sigma(\langle P^{i-1}_{eval}, P^{i-1}_{coll}\rangle)$$

and that the O-contract path realising this deal comprises only IR contracts. We now show that $\sigma(\langle P^{i+1}_{eval}, P^{i+1}_{coll}\rangle) > \sigma(\langle P^i_{eval}, P^i_{coll}\rangle)$ and that the O-contracts effected in iteration $(i+1)$ are all IR. We know that

$$\begin{aligned}
P^i_{eval} &= \{y_1, \ldots, y_i\} && \cup \{C : y_k \notin C \ \forall\ 1 \leq k \leq i\} \\
P^i_{coll} &= \{\neg y_1, \ldots, \neg y_i\} && \cup \{C : y_k \in C \text{ for some } 1 \leq k \leq i\} \\
& && \cup \{x_{i+1}, \ldots, x_n, \neg x_{i+1}, \ldots, \neg x_n\}
\end{aligned}$$

After the O-contract transferring y_{i+1} to P^i_{eval},

$$u_{eval}(P^i_{eval} \cup \{y_{i+1}\}) = 1 + \tfrac{i+1}{2n+m} = \tfrac{1}{2n+m} + u_{eval}(P^i_{eval})$$
$$u_{coll}(P^i_{coll} - \{y_{i+1}\}) = u_{coll}(P^i_{coll})$$

representing an increase in social welfare. The next k_{i+1} O-contracts are transfers of single clauses (containing y_{i+1}) from C_{eval} to C_{coll}, and as before each of these increases the utility of A_{coll} by $1/(2n+m)$ leaving that of A_{eval} unchanged. It follows that $\sigma(\langle P^{i+1}_{eval}, P^{i+1}_{coll}\rangle) > \sigma(\langle P^{i}_{eval}, P^{i}_{coll}\rangle)$ and that the O-contracts effected in this iteration are all IR.

To complete the first part part of the proof, we note that

$$P^{n}_{eval} = \{y_1, \ldots, y_n\}$$
$$P^{n}_{coll} = \{\neg y_1, \ldots, \neg y_n\} \cup \{C_1, C_2, \ldots, C_m\}$$

The concluding sequence of O-contracts transfers each of the literals $\neg y_i$ to A_{eval}: since $C_{eval} = \emptyset$ each such transfer adds $1/(2n+m)$ to the current utility enjoyed by A_{eval} without changing the value of u_{coll}. Thus the final sequence of O-contracts is IR.

It follows that if $\Phi(X_n)$ is satisfiable then the instance T_Φ of IRO defined above is accepted.

On the other hand suppose that there is an IR O-contract path realising $\langle P^s, P^t\rangle$. We need to show, in this case, that $\Phi(X_n)$ is satisfiable.

Consider any IR O-contract path, T realising the deal $\langle P^s, P^t\rangle$, with $T_0 = P^s$ and $T_r = P^t$,

$$T = \langle T_0, T_1\rangle, \langle T_1, T_2\rangle, \ldots, \langle T_k, T_{k+1}\rangle, \ldots, \langle T_{r-1}, T_r\rangle$$

Let

$$T_i = \langle Y^{i}_{eval} \cup C^{i}_{eval}, Y^{i}_{coll} \cup C^{i}_{coll}\rangle$$

We claim that for each i (with $1 \le i \le r$) the IR O-contract $\langle T_{i-1}, T_i\rangle$ satisfies the following properties:

P1) The deal $\langle T_{i-1}, T_i\rangle$ either transfers some literal $y \in Y^{i-1}_{coll}$ to Y_{eval} or transfers some clause $C \in C^{i-1}_{eval}$ to C_{coll}. We call the former a *literal transfer* and the latter a *clause transfer*.
P2) If $\langle T_{i-1}, T_i\rangle$ is a literal transfer involving y_k then the literal $\neg y_k$ is in Y^{i-1}_{eval} only if C^{i-1}_{eval} is empty.
P3) If $\langle T_{i-1}, T_i\rangle$ is a clause transfer involving C_j then the instantiation of X_n that makes each literal in Y^{i-1}_{eval} true is both well-defined (i.e. at most one of y_k and $\neg y_k$ occurs in Y^{i-1}_{eval}) and satisfies each clause in C^{i}_{coll}.
P4) $\sigma(T_i) = 2 + \frac{|Y^{i}_{eval}|+|C^{i}_{coll}|}{2n+m}$.

We establish this claim by induction on i. For the inductive base, $i = 1$, we have

$$T_0 = \langle \{C_1, \ldots, C_m\}, \{x_1, \ldots, x_n, \neg x_1, \ldots, \neg x_n\} \rangle$$

That P1 holds of the deal $\langle T_0, T_1 \rangle$ is obvious. If, however, $\langle T_0, T_1 \rangle$ is a clause transfer then $\sigma(T_1) = \sigma(T_0) = 2$: C_{coll} contains a single clause and this cannot be satisfied by an instantiation that does not assign values to any variables, i.e case (c1) applies and the deal would not be IR. Thus $\langle T_0, T_1 \rangle$ must be a literal transfer involving some $y \in Y^0_{coll}$. Since $Y^0_{eval} = \emptyset$ it is immediate that $\neg y \notin Y^0_{eval}$ and so property P2 holds. Finally,

$$\begin{aligned} &u_{eval}(\{y\} \cup \{C_1, \ldots, C_m\}) = 1 + \tfrac{1}{2n+m} \\ &u_{coll}(Y^0_{coll} - \{y\}) = 1 \end{aligned}$$

and $\sigma(T_1) = 2 + (1 + 0)/(2n + m)$ establishing P4.

For the inductive step assume that we have established P1 through P4 hold for all deals $\langle T_{j-1}, T_j \rangle$ with $1 \leq j \leq i$ for some $i < r$, we show that these continue to hold of the deal $\langle T_i, T_{i+1} \rangle$. We have,

$$\begin{aligned} P^i_{eval} &= Y^i_{eval} \cup C^i_{eval} \\ P^i_{coll} &= Y^i_{coll} \cup C^i_{coll} \end{aligned}$$

with $\sigma(T_i) = 2 + (|Y^i_{eval}| + |C^i_{coll}|)/(2n + m)$. We first show that $\langle T_i, T_{i+1} \rangle$ must be either a literal or a clause transfer, i.e. that P1 holds. Suppose this is not the case: the only alternatives are that some literal $y \in Y^i_{eval}$ is transferred to Y^{i+1}_{coll} or some clause $C \in C^i_{coll}$ is transferred to C^{i+1}_{eval}. In the former case we have

$$\begin{aligned} \sigma(T_{i+1}) &= u_{eval}(Y^i_{eval} - \{y\} \cup C^i_{eval}) + u_{coll}(Y^i_{coll} \cup \{y\} \cup C^i_{coll}) \\ &\leq 2 + \tfrac{|Y^i_{eval}| - 1 + |C^i_{coll}|}{2n+m} < \sigma(T_i) \end{aligned}$$

Thus such a deal could not increase social welfare.[4] A similar argument establishes that $\sigma(T_{i+1}) = \sigma(T_i)$ in the event of some clause $C \in C^i_{coll}$ being moved to C^{i+1}_{eval}.

We deduce that property P1 holds of the deal $\langle T_i, T_{i+1} \rangle$. Now suppose that $\langle T_i, T_{i+1} \rangle$ is a literal transfer involving $y \in Y^i_{coll}$. If $C^i_{eval} = \emptyset$ then P2 obviously holds. If

[4] The inequality in the second line is needed because this deal would *further* decrease social welfare if the literal transferred was needed to satisfy some clause in C^i_{coll}.

$C^i_{eval} \neq \emptyset$ then were it the case that $\neg y \in Y^i_{eval}$ we would get

$$\sigma(T_{i+1}) = 1 + \frac{|Y^i_{eval}|}{2n+m} + 1 + \frac{|C^i_{coll}| - 1}{2n+m} < \sigma(T_i)$$

(via e1 and c1 of the utility function definitions) contradicting the fact that the contract $\langle T_i, T_{i+1} \rangle$ is IR. It follows that property P2 holds.

To see that P4 holds also, in both cases, first consider $C^i_{eval} = \emptyset$:

$$\begin{aligned} \sigma(T_{i+1}) &= u_{eval}(Y^i_{eval} \cup \{y\}) + u_{coll}(Y^i_{coll} - \{y\} \cup \{C_1, \ldots, C_m\}) \\ &= u_{eval}(Y^i_{eval}) + \tfrac{1}{2n+m} + \tfrac{|C^i_{coll}|}{2n+m} \\ &= 2 + \tfrac{|Y^{i+1}_{eval}| + |C^{i+1}_{coll}|}{2n+m} \end{aligned}$$

establishing P4 in this case. When $C^i_{eval} \neq \emptyset$ it suffices to observe that the instantiation given by setting all literals in Y^{i+1}_{eval} must satisfy all of the clauses belonging to $C^{i+1}_{coll} = C^i_{coll}$, hence

$$\begin{aligned} \sigma(T_{i+1}) &= 2 + \tfrac{|Y^i_{eval}| + 1 + |C^i_{coll}|}{2n+m} \\ &= 2 + \tfrac{|Y^{i+1}_{eval}| + |C^{i+1}_{coll}|}{2n+m} \end{aligned}$$

It remains to show that P3 and P4 hold in the event of $\langle T_i, T_{i+1} \rangle$ being a clause transfer involving $C \in C^i_{eval}$. From the inductive hypothesis we know that the instantiation formed by setting each literal in $Y^i_{eval} = Y^{i+1}_{eval}$ to true is well-defined (since $C^i_{eval} \neq \emptyset$) and satisfies each clause in C^i_{coll}. If this instantiation does *not* satisfy C then we would have

$$\sigma(T_{i+1}) = 2 + \frac{|Y^{i+1}_{eval}| - 1 + |C^{i+1}_{coll}| - 1}{2n+m} < \sigma(T_i)$$

and $\langle T_i, T_{i+1} \rangle$ would not be IR. We note the need for condition c2 in defining u_{coll} here: if $C_{eval} = \{C\}$ and Y_{eval} constitutes a well-defined satisfying instantiation of $\{C_1, \ldots, C_m\} - \{C\}$ but *not* of $\Phi(X_n)$ then using only (c1) and (c3) to define u_{coll} would allow the clause transfer of C to be IR even though the resulting allocation does not yield a satisfying instantiation.

It follows that if $\langle T_i, T_{i+1} \rangle$ is a clause transfer then the conditions P3 and P4 are satisfied, so completing the inductive argument that these hold for each IR O-contract in the path realising the deal $\langle P^s, P^t \rangle$. We derive a satisfying assignment for $\Phi(X_n)$ simply by setting the literals in Y^k_{eval} appropriately where T_k is the first allocation within the contract path for which $C^k_{eval} = \emptyset$.

In total an IR O-contract path realising the deal $\langle P^s, P^t \rangle$ exists if and only if $\Phi(X_n)$ is satisfiable. $\square$

Theorem 11 suggests that if it is required that each deal be IR, determining if a suitable O-contract path from an initial allocation to some notionally improved one is unlikely to be decidable efficiently. We observed earlier that one motivation for considering O-contracts is that the trading environment may render negotiations concerning large numbers of resources infeasible. Suppose we relax the conditions governing contract paths to all C-contracts where some bound is placed on the size of the resource cluster (but preserving the requirement that each deal be IR). Thus, extending Definition 5, we say that a C-contract is k-bounded (kC-contract) if the number of resources involved in any single deal is at most k. Even if we allow k to be a function of the number of resources it turns out that deciding if an IR kC-contract path exists is intractable.

Definition 12 *An instance of the decision problem* IR-kC-path *(*IR^k*) is a 5-tuple* $\langle \mathcal{A}, \mathcal{R}, \mathcal{U}, P^s, P^t \rangle$ *as in* IRO. *An instance is accepted if there is an IR kC-contract path realising the deal* $\langle P^s, P^t \rangle$.

Theorem 13 IR^k *is* NP*–hard, even if k is a function of* $m = |\mathcal{R}|$ *with* $k(m) \leq m/3$.

Proof: (*outline*) We employ a reduction from 3–SAT but with instances restricted to those containing *exactly n* clauses. Letting

$$\Phi(X_n) = \bigwedge_{i=1}^{n} C_i = \bigwedge_{i=1}^{n} (y_{i,1} \vee y_{i,2} \vee y_{i,3})$$

be an instance we construct $\langle \mathcal{A}, \mathcal{R}, \mathcal{U}, P^s, P^t \rangle$ exactly as in Theorem 11 except that the utility functions u_{eval} and u_{coll} are now given by:

- e1) $u_{eval}(Y_{eval} \cup C_{eval}) = 0$ if for some i, both x_i and $\neg x_i$ belong to Y_{eval} and $C_{eval} \neq \emptyset$.
- e2) $u_{eval}(Y_{eval} \cup C_{eval}) = |Y_{eval}|$ otherwise.
- c1) $u_{coll}(Y_{coll} \cup C_{coll}) = 0$ if $C_{eval} \neq \emptyset$ and the instantiation formed by setting each literal in Y_{eval} is either undefined (since both z and $\neg z$ are in Y_{eval}) or fails to satisfy each clause in C_{coll}
- c2) $u_{coll}(Y_{coll} \cup C_{coll}) = 0$ if $C_{eval} = \emptyset$, $|Y_{eval}| \leq n$, and the instantiation formed by setting each literal in Y_{eval} is either undefined or fails to satisfy $\Phi(X_n)$.
- c3) $u_{coll}(Y_{coll} \cup C_{coll}) = |C_{coll}|$ otherwise.

We note that $|\mathcal{R}| = 3n$ so our bound on cluster size allows at most n elements from $\mathcal{R}$ to feature in a single deal. If $\Phi(X_n)$ is satisfiable then we have already seen that an IR O-contract is possible. The reverse implication is immediate from the following claim which can be established in a manner similar to the corresponding result in the proof of Theorem 11.

If T is a IR kC-contract path realising the deal $\langle P^s, P^t\rangle$, with $T_0 = P^s$ and $T_r = P^t$,

$$T = \langle T_0, T_1\rangle, \langle T_1, T_2\rangle, \ldots, \langle T_k, T_{k+1}\rangle, \ldots, \langle T_{r-1}, T_r\rangle$$

with $T_i = \langle Y^i_{eval} \cup C^i_{eval}, Y^i_{coll} \cup C^i_{coll}\rangle$ then for each deal $\langle T_{i-1}, T_i\rangle$ $(1 \leq i \leq r)$ the following are true.

P1) $\langle T_{i-1}, T_i\rangle$ either transfers $Y_S \subseteq Y^{i-1}_{coll}$ to Y^i_{eval} (referred to as a *literal set* transfer) or transfers $C_S \subseteq C^{i-1}_{eval}$ to C^i_{coll} (referred to as a *clause set* transfer).
P2) If $\langle T_{i-1}, T_i\rangle$ is a literal set transfer of Y_S then

$$\{\neg y : y \in Y_S\} \cap Y^{i-1}_{eval} \neq \emptyset \;\Rightarrow\; C^{i-1}_{eval} = \emptyset$$

P3) If $\langle T_{i-1}, T_i\rangle$ is a clause set transfer of C_S then the instantiation given by setting each literal in Y^{i-1}_{eval} is well-defined and satisfies each clause in C^i_{coll}.
P4) $\sigma(T_i) = |Y^i_{eval}| + |C^i_{coll}|$

□

3 Conclusion

We have considered a number of decision problems that naturally arise from the multiagent contract negotiation models promoted by (among others) [4,9]. In summary, if contracts are restricted to those in which a limited number of resources can be transferred from one agent to another and are required to be rational (in the sense of strictly improving overall worth of an allocation), then not only is it the case that a suitable contract-path to an optimal allocation may fail to exist (as already shown in [9]), but even *deciding* if a path from a given allocation to a specified more beneficial allocation is possible, is intractable. There are a number of directions in which the results above could be developed. The requirement for individuals deals in a contract-path to be IR could be relaxed so that a limited number of 'irrational' deals are permitted, provided that the allocation eventually reached improves upon the initial allocation. Alternatively, we could consider contracts in which deals permitting an *exchange* of resources between two agents are allowed – the so-called *swop* or S-contracts of [9]. We conjecture, however, that even these degrees of freedom will continue to yield decision questions that are intractable.

Acknowledgments

The work reported in this article was carried out under the support of EPSRC Grant GR/R60836/01.

References

[1] F Dignum and M. Greaves (editors). *Issues in Agent Communication*. Springer-Verlag, 2000.

[2] P.E. Dunne *The Complexity of Boolean Networks* Academic Press, 1988

[3] P.E. Dunne, A.M. Gibbons, and M. Zito. Complexity-theoretic models of phase-transitions in search problems. *Theoretical Computer Science*, 294:243–263, 2000.

[4] U. Endriss, N. Maudet, F. Sadri, and F. Toni. On optimal outcomes of negotiations over resources. In *to appear: Proc. Second Intnl. Joint Conf. on Autonomous Agents and Multiagent Systems, (AAMAS-2003)*, 2003.

[5] S. Kraus. *Strategic negotiation in multiagent environments*. MIT Press, 2001.

[6] P. McBurney, S. Parsons, and M. Wooldridge. Desiderata for argumentation protocols. In *Proc. First Intnl. Joint Conf. on Autonomous Agents and Multiagent Systems, (AAMAS-2002)*, pages 402–409. ACM Press, 2002.

[7] D.E. Muller and F.P. Preparata Bounds to complexities of networks for sorting and for switching. *Jnl. of the Assoc. for Comput. Mach.*, 22, pages 195–201, 1975

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

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

[10] T. W. Sandholm. Distributed Rational Decision Making In *Multiagent Systems* (Editor: G. Weiß) pages 201–258, MIT Press, 1999.

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