Tractability Results for Automatic Contracting

Paul E. Dunne and **Michael Laurence** and **Michael Wooldridge**[1]

Abstract. Automated negotiation techniques have received considerable attention over the past decade, and much progress has been made in developing negotiation protocols and strategies for use by software agents. However, comparatively little effort has been devoted to understanding the computational complexity of such protocols and strategies. Building on the work of Rosenschein, Zlotkin, and Sandholm, we consider the complexity of negotiation in a particular class of task-oriented domains. Specifically, we consider scenarios in which agents negotiate to achieve a more favourable redistribution of tasks amongst themselves, where the tasks involve visiting nodes in a graph. Focussing on a particular representation of the domain (as a spanning tree), we establish a number of complexity results pertaining to the complexity of negotiation in this scenario, with our main result to the effect that the problem of deciding whether a given deal could be reached by a chain of rational proposals is tractable.

1 Introduction

Automated negotiation has been the subject of considerable research over the past two decades [2, 4, 1]. One of the most important contributions to this research literature was the seminal work of Rosenschein and Zlotkin, who classified negotiation domains according to whether they were *task oriented*, *worth oriented*, or *state oriented* [2]. In a task oriented domain, each agent is allocated a set of tasks to perform, where each task set has some well-defined cost. Agents in a task oriented domain can mutually benefit from negotiation be *rearranging* the allocation of tasks amongst themselves, thereby reducing the overall cost of each agent's allocation. Perhaps the paradigm example of a task oriented domain is the "postman" scenario, which is defined by a weighted graph. An individual task in the postman scenario corresponds to visiting a node in the graph (to "deliver a letter"), and thus an agent's task allocation is a set of nodes in the graph; the cost of performing an allocation of tasks is then the cost of the minimal cost tour of the graph that includes all nodes in the allocation. Negotiation can be mutually beneficial because agents can reallocate tasks so that they are required to visit nodes in the same region of the graph, thereby reducing the cost of the minimal cost tour that includes the nodes in their allocation.

Although task oriented domains have influenced subsequent research enormously, a number of issues have prevented their wider implementation and take-up. Chief among these is that the protocols and strategies for negotiation in task oriented domains have a high computational complexity. For example, Rosenschein and Zlotkin point out that implementing the basic "Zeuthen strategy" in task oriented domains requires $O(2^n)$ computations of the task cost function [2, p.49]. But a closer analysis shows that the situation may be much worse than this in many cases: simply computing the cost of a set of tasks in the delivery domain, for example, implies solving an NP-hard optimisation problem (the synthesis of a minimal cost tour). As a consequence, it is of great importance to gain a proper understanding of (i) the precise computational complexity of negotiation in task oriented domains, and (ii) the cases in which such negotiation is tractable.

In short, the present paper contributes to this understanding. Focussing on a particular representation of the domain (as a spanning tree), we establish a number of complexity results pertaining to the complexity of negotiation in the "postman" scenario, with our main result to the effect that the problem of deciding whether a given deal could be reached by a chain of rational proposals is tractable.

[1] Department Of Computer Science, University of Liverpool, Liverpool L69 7ZF. email: mjw@csc.liv.ac.uk

2 Preliminary Definitions

We let $V = \{v_1, v_2, \ldots, v_n\}$ denote a set of n *cities* (these are the nodes in the graph that agents must visit), and let $M = [b_{i,j}]$ be an $n \times n$ (symmetric) matrix of rationals with $b_{i,j}$ being the cost of linking v_i and v_j; we assume $b_{i,i} = 0$ and $b_{i,j} = b_{j,i}$ but do not insist on the triangle inequality $b_{i,j} \leq b_{i,k} + b_{j,k}$. Thus $b_{i,j}$ is the cost of moving from city i to city j.

For $S = \{s_1, \ldots, s_p\} \subseteq V$, M_S is the $p \times p$ submatrix of M induced by including only those rows and columns indexed by elements of S. A *spanning tree* of S is formed by any set of edges $E_S \subset \{\{i,j\} : i,j \in S\}$ such that $|E_S| = |S| - 1$, and for any pair v, w of cities in S there is a path formed from edges in E_S linking v and w. The *weight* of any spanning tree $T(S,E)$ of S is the sum of the individual edge weights b_{v_i,v_j} with $\{v_i, v_j\} \in E$, with this weight denoted $w(T)$. The *cost* of S is

$$u(S) = \min \{w(T(S,E)) : T(S,E) \text{ is a spanning tree of } S\}.$$

We note that in defining $u(S)$ a minimum spanning tree is not permitted to contain locations other than those specified in S. There do in fact exist cost matrices giving rise to sets S, S' such that $S \subseteq S'$ and $u(S') < u(S)$. 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. An* allocation *to $\mathcal{A}$ is a partition $P = \langle P_1, P_2, \ldots, P_n \rangle$ of V.*

We now state a basic result with respect to these structures.

Definition 2 *The* MST Allocation Problem *(*MSTAP*) takes as an instance a triple of the form $\langle V, [b_{i,j}], k \rangle$ where $V = \{v_1, \ldots, v_n\}$ is a set of locations, $[b_{i,j}]$ a $n \times n$ cost matrix and k a positive integer ($k \geq 2$). The output given such an instance is a partition of V as $P = \langle P_1; P_2; \ldots, P_k \rangle$ for which $\sigma(P) = \sum_{i=1}^{k} u(P_i)$ is minimised.*

Theorem 1 *Given* $\langle V, [b_{i,j}], k]\rangle$, MSTAP *can be solved in* $O(n\beta_M \log n)$ *steps, where* β_M *is the maximum number of bits used to encode any value in* M.

Starting from some initial allocation – P_0 – individual agents negotiate in an attempt to improve the utility of their holding, i.e. *reduce* the cost of forming a spanning tree of their assigned locations. 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 $P_i, Q_i \subseteq V$ and $P_i \neq Q_i$) there are three possible outcomes for the agent A_i: $u(P_i) < u(Q_i)$, i.e. A_i values the allocation P_i as superior to Q_i since the cost of spanning P_i is less than that of spanning Q_i; $u(P_i) = u(Q_i)$, i.e. A_i is indifferent between P_i and Q_i; and $u(P_i) > u(Q_i)$, i.e. A_i is better 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 first 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 3 *Let* $\mathcal{A}$ *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* V. *The effect of implementing the deal* $\langle P, Q\rangle$ *is that the allocation of cities specified by* P *is replaced with that specified by* Q. *A deal* $\langle P, Q\rangle$ *is said to be* individually rational *(IR) if* $\sum_{i\leq n} u(Q_i) < \sum_{i\leq n} u(P_i)$.
Let $t \geq 1$. *A* t*-contract 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* V, *such that for some* $i, j \leq n$, $Q_i = P_i \cup X$, $Q_j = P_j - X$, $|X| \leq t$ *and* $P_k = Q_k$ *for all* $k \notin \{1, 2\}$, *and* $u(Q_i) + u(Q_j) < (P_i) + u(P_j)$. *A sequence of deals* $\langle\langle Q_0, Q_1\rangle, \ldots, \langle Q_{m-1}, Q_m\rangle\rangle$ *is called a* t-contract path *if each pair* $\langle Q_{i-1}, Q_i\rangle$ *is a* t*-contract.*

3 Tree Structures

It appears to be very difficult to prove strong complexity results for arbitrary cost matrices, so we define a subclass of such matrices.

Definition 4 (The Tree Structure restriction) A *tree structure* is a tuple $\langle V, E, cost\rangle$ such that V is a finite set, E is a set of edges of V such that (V, E) is an undirected tree and $cost : V \times V \to \mathbb{Q}^{\geq 0}$ is a function such that $cost(v, w) > 0$ if and only if $v \neq w$, $cost(u, v) + cost(v, w) \geq cost(u, w)$ for all $u, v, w \in V$ and if E' is a set of edges of V such that (V, E') is a tree then $\sum_{z\in E} cost(z) \leq \sum_{z\in E'} cost(z)$ holds; that is, (V, E) is a minimal-cost spanning tree with respect to V and the function $cost$. Given a tree structure tuple $T = \langle V, E, cost\rangle$, and $W \subseteq V$, we define $expense_T(W) = \sum_{z\in F} cost(z)$, where F is any set of edges of W such that (W, F) is a tree and the sum $\sum_{z\in F} cost(z)$ is minimal for all such edge sets.
Let $T = \langle V, E, cost\rangle$ be a tree structure. Then T is said to be *sensible* if the following holds; if $v, w, w' \in V$ and the reduced path in the tree (V, E) from v to w' passes through w, then $cost(v, w) \leq cost(v, w')$. Also T is said to be *maximal* if for all $v, w, \in V$, $cost(v, w)$ is the sum of the costs of all the edges in the path through (V, E) joining v to w.

Clearly, maximal tree structures are sensible. The justification for this terminology is that if T is maximal, then for all $v, w \in V$, $cost(v, w)$ has the maximal value compatible with the triangle inequality which the function $cost$ must satisfy. We now state some results of maximal tree structures: Theorems 2 and 3 are the main results of this section, establishing that, for maximal trees, it is decidable whether any allocation is achievable by an IR contract path (Theorem 2), and that, in addition, it is always possible to construct optimal allocations that cannot be reached by any IR $C(t)$ contract path (Theorem 3).

Theorem 2 *Let* $T = \langle V, E, cost\rangle$ *be a maximal tree structure. If* $Q = \langle S_1; S_2\rangle$ *is any allocation and* $1 \leq t < \min\{|S_1|, |S_2|\}$, *then it is decidable in polynomial time in* $\langle T, Q, t\rangle$ *whether there is a* t*-contract path from* $P_{init} = \langle V; \emptyset\rangle$ *to* Q,

Theorem 3 *For any* $k \geq 2$, *for any fixed* $t \geq 1$, *the following holds.*

a) There are instances, $\langle V, [b_{i,j}], k\rangle$ *of* MSTAP *such that with* $P_{init} = \langle V; \emptyset; \ldots; \emptyset\rangle$ *the initial allocation and* P_{opt} any *optimal allocation, there exists a* $t + 1$*-contract path realising the deal* $\langle P_{init}, P_{opt}\rangle$, *but there does* not *exist any* t*-contract path to realise the deal* $\langle P_{init}, P_{opt}\rangle$.
b) There is a maximal tree structure $T = \langle V, E, cost\rangle$ *such that with* $P_{init} = \langle V; \emptyset; \ldots; \emptyset\rangle$ *the initial allocation and* P_{opt} any *optimal allocation, there does* not *exist any* t*-contract path to realise the deal* $\langle P_{init}, P_{opt}\rangle$.

4 The Main Theorems

Theorem 4 *Let* $T = \langle V, E, cost\rangle$ *be a sensible tree structure and let* $\langle V_0 = V, \emptyset\rangle$, $\langle V_1, V - V_1\rangle, \ldots, \langle V_n, V - V_n\rangle$ *be a* 1*-contract path. Then each set* V_{i+1} *has one fewer elements than* V_i.

Theorem 5 *Let* $T = \langle V, E, cost\rangle$ *be a sensible tree structure and let* $\langle V_0 = V, \emptyset\rangle$, $\langle V_1, V - V_1\rangle, \ldots, \langle V_n, V - V_n\rangle$ *be a* 1*-contract path. Let* $\langle U_0 = V, \emptyset\rangle$, $\langle U_1, V - U_1\rangle, \ldots, \langle U_m, V - U_m\rangle$ *be another* 1*-contract path with* $U_1 = V_1$. *Suppose* $U_m \supseteq V_n$. *Then this second sequence can be continued monotonically using* 1*-contracts until it reaches* $\langle V_n, V - V_v\rangle$.

Theorem 6 *Let* $T = \langle V, E, cost\rangle$ *be a sensible tree structure and let* $\langle V', V - V'\rangle$ *be any (not necessarily minimal-cost) allocation. It is decidable in polynomial time whether* $\langle V', V - V'\rangle$ *is reachable from* $\langle V, \emptyset\rangle$ *by a* 1*-contract path.*

Acknowledgments

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

REFERENCES

[1] S. Kraus. *Strategic negotiation in multiagent environments*. MIT Press, 2001.
[2] J. S. Rosenschein and G. Zlotkin. *Rules of Encounter*. MIT Press, 1994.
[3] T. W. Sandholm. Contract types for satisficing task allocation: I theoretical results. In *AAAI Spring Symposium: Satisficing Models*, 1998.
[4] T. W. Sandholm. Distributed Rational Decision Making. In *Multiagent Systems* (Editor: G. Weiß) pages 201–258, MIT Press, 1999.