ON SOCIAL ATTITUDES: A PRELIMINARY REPORT

S. Kalenka and N. R. Jennings
Dept. of Electronic Eng., Queen Mary and Westfield College,
Univ. of London, Mile End Road, London E1 4NS, U.K.,
[S.Kalenka,N.R.Jennings]@qmw.ac.uk

Key words: multiple-agent communities; social interactions; loading platform scenario

ABSTRACT

This paper describes our preliminary investigations into identifying, and ultimately formalising, the wide range of social problem solving behaviour which can occur in multi-agent systems. In particular, this work examines the different agent attitudes which may prevail in social problem solving, the different types of individual and social commitment which flow from these attitudes, how and why these attitudes come to prevail, and how the attitudes lead to overall systems with radically different properties and characteristics.

1. AGENT ATTITUDES & COMMITMENTS

To clearly delimit the extent of this work, this paper concentrates on the dynamics of groups of autonomous problem solving agents who are engaged in cooperative problem solving. We assume that each agent has been assigned a particular role in the collective — the means by which this assignment is made (e.g. through centralised or decentralised planning, through peer to peer negotiation, through the establishment of long-term organisational roles, etc.) is not of concern here. Within this context, three distinct types of agent attitude can be identified: (i) *responsibility* — agents only execute the tasks directly associated with their role in the group; (ii) *helpfulness* — agents assist others when they have no responsible tasks to perform (i.e. in their spare time); and (iii) *cooperativeness* — agents assist one another in return for reciprocated support for one of their own responsible tasks.

Responsibility should not be confused with self-interest, it simply means that the agent sticks steadfastly to its own obligations within the group and that it never reconsiders this position. Thus when a group of responsible agents undertake a problem, each member performs precisely its part—if one of the other agents runs into difficulty it will receive no support from any of the others. Helpfulness and cooperativeness, on the other hand, allow an agent to be *supportive* of a group member in its pursuit of its responsible tasks. Thus if a problem is tackled by a group of supportive agents and one of them runs into difficulty with its part, then one of the others *may* come to its assistance. For this reason, supportive agents can be considered to be explicitly concerned with the wider aspects and success of the group's efforts as well as with their own individual contributions.

For helpful agents, support is unconditional in that the agent doing the helping does not expect anything in return. It helps because it has spare capacity. After allowing for all its responsible tasks, the agent has the potential to do some additional problem solving for the benefit of another agent if it so chooses. This, in turn, should improve the overall performance of the collective's endeavours. Helpful actions are undertaken because the assisting agent perceives the benefit, either to itself or to the collective, to be greater than the benefit of remaining idle. Two distinct types of help can be identified: *reliable* and *unreliable*. With reliable help, the assisting agent knows to what extent (in terms of time, resources and capability) it is able to assist another agent with its responsible tasks. With unreliable help, the assistance is less predictable — the duration, resources or capability with which the agent can help is unknown. Thus, for example, a reliable helpful agent is able to commit itself to help another agent for a set period of time, whereas an unreliable helpful agent may only be able to commit itself to help until it receives a new responsible task. In terms of obtaining assistance for one of its responsible tasks, an agent is able to predict more accurately the level of help which can be provided by a reliable agent than it can for an unreliable one. In general, such knowledge helps the agents to organise their problem solving, in terms of group creation, group evolution, and group disbanding, more effectively in uncertain and unpredictable environments.

For cooperative agents, support is conditional on receiving subsequent assistance from the agent which is being helped. Cooperation is required either because the agent is unable to fulfil its responsible tasks alone or because it decides, for whatever reason, to assist another agent and consequently may be unable to honour its own responsibilities alone. In the former case, the agent simply does not posses sufficient resources to complete the task alone and therefore it must enter into an agreement with another agent to provide the additional resources. Ideally the agent would prefer to receive this assistance without having to commit its own resources (i.e. through a helpful agent), however if such help is not forthcoming then it must entice another agent with the offer

that it will reciprocate. In the latter case, the agent offers assistance because it believes that the benefit, to itself or to the community, is greater for the other agent's tasks than it is for its own. By the very fact that it has offered to support another agent, it may no longer be able to fulfil its responsible tasks (even though it could have done so originally). In generating and evaluating cooperative proposals and agreements, agents can take one of three stances: (i) *non-risk taking:* as long as the other agent honours its commitment to provide the agreed level of support, the agent will be able to complete all of its responsible tasks; (ii) *risk taking*: an agent does not know if it will be able to perform all of its responsible tasks; (iii) *altruistic*: an agent knows that it will not be able to honour all of its responsible tasks. An agent will adopt an altruistic stance when it places greater priority on the goals of the others. It will adopt a risk taking stance when it believes the only way of achieving its own goals is to obtain assistance from others—in doing this it hopes that its own tasks will still be completed. An agent will adopt a non-risk taking stance when it places greatest priority on its own tasks, but also has a concern for the wider collective.

The main contribution of this work is to identify and characterise the different types and levels of commitment which can occur in social problem solving. At one extreme there is pure individual commitment, as exemplified by responsible agents performing their own tasks. At the next level are commitments offered by helpful agents—unreliable agents offer a one way, weak form of commitment (they renege as soon as they have any responsible tasks to perform); reliable agents have a more predictable form of one way commitment. At the highest level is the bi-directional commitments undertaken by cooperative agents—non-risk taking agents renege if their own tasks become jeopardised, risk taking agents may renege if their own tasks become jeopardised, and altruistic agents never renege. We have found that these different levels of commitment are needed if agents are to act intelligently and flexibly in dynamic and unpredictable environments.

To provide a basis for discussing these social attitudes, a loading platform scenario in which a number of fork lift trucks (agents) interact to unload lorries is used (section 2). Section 3 shows how agents with different attitudes interact to give overall system characteristics and section 4 outlines future work.

2. A MULTI-AGENT LOADING BAY SCENARIO

Lorries arrive randomly at a warehouse laden with goods which require unloading. The warehouse has a fixed number of (un)loading bays which each hold one lorry at a time. Upon arrival, lorries go to the nearest free bay. Lorries have an associated time by which they must be discharged. Additionally, the warehouse tries to ensure that all lorries are processed by some maximum time (T_{max}). Since the aim of this work is to

illustrate social problem solving behaviour, rather than to develop a real world solution, we make a number of simplifying assumptions: all lorries take the same time to unload; the time it takes one agent to unload a lorry is T_{max}; unloading time is directly proportional to the number of agents servicing a lorry (two agents will do it twice as fast as one); and there are as many fork lift trucks unloading the lorries as there are bays in the warehouse. The multi-agent system needed to solve this problem has two key objectives: (i) all lorries should be discharged just-in-time for their deadline; and (ii) all lorries should be discharged within the guaranteed maximum time (unless the specified desired deadline is greater than T_{max}).

One way of solving this problem is to assign a fork lift truck (agent) to each loading bay—meaning each agent is responsible for a specific bay. Let L_i be the lorry to be discharged, A_i be the agent responsible for L_i, and T_i be L_i's discharge deadline. This leads to four distinct cases:

- $T_i > T_{max}$: A_i can discharge L_i on its own and it has some spare time ($T_{max} - T_i$) in which it could support other agents.
- $T_i = T_{max}$: A_i can discharge L_i on its own.
- $T_i < T_{max}$: A_i needs support from other agents if it is to meet L_i's desired deadline. Without this support discharge will take T_{max}.
- $T_i = ?$: No L_i has arrived at bay i, hence A_i is free to support other agents until T_i is defined.

If all agents were simply responsible, then the multi-agent system could be guaranteed to meet its objectives as long as no lorry needed to be discharged in a time less than T_{max}. However, this situation could be improved if some of the agents adopted a helpful attitude. In this case, agents would support one another when they have a lorry with a discharge deadline greater than T_{max} or when they have no lorry at their bay. Helpfulness would ensure the agents are more heavily utilised and that more deadlines less than T_{max} are met (simply because there is more agent problem solving power available in the system).

The situation can be enhanced still further by the addition of cooperative agents. For example, consider the case where two lorries arrive simultaneously. Lorry L_1 has to be discharged in time T_{max} and lorry L_2 in time $T_{max} / 2$. If A_1 and A_2 cooperate, A_1 could assist A_2 for the first $T_{max} / 2$ units of time (meaning L_2 is discharged in time) so long as A_2 agreed to assist it for the second $T_{max} / 2$ units of time (meaning L_1 is also discharged in time). With only responsible or helpful agents, both lorries would be discharged in time T_{max} meaning L_2's desired departure time is not satisfied. Depending on the stance of A_1 and A_2, two types of commitment may be used to achieve this

form of cooperation— *one-shot* commitment in which the agents agree to assist one another with a specific set of tasks; *ongoing* commitment in which the agents agree to assist one another until a particular target state is achieved. The former involves a precise duration and level of resource, whereas the latter is open ended. If A_2 is risk-taking then it will agree to a one-shot commitment. The risk in this case, is that another lorry will arrive at its bay while it is honouring its part of the agreement in assisting A_1 with L_1. A_1 is non-risk taking in this case because L_1's desired deadline will be met. If A_2 is non-risk taking then it requires an ongoing commitment to perform this type of interaction—both agents commit to support one another while either of them are in danger of not fulfilling their responsible tasks (the target state). Thus if a lorry arrives at A_2's bay while it is assisting A_1 with L_1, then A_1 must commit itself to help A_2 clear the new lorry, and so on. In the worst case, new lorries, with deadlines less than or equal to T_{max} plus the time the responsible agent is unavailable to start work, may arrive continuously at A_1's and A_2's bays while they are supporting one another. Thus their ongoing commitment will continue until: (i) no lorries arrive before the responsible agents fulfil their duties at the other's bay; (ii) one of the lorries has a sufficiently long discharge time; or (iii) either agent receives the necessary support from some helpful agent.

3. SOCIAL INTERACTIONS & AGENT ATTITUDES

This work is concerned with the when, why, with whom, and how aspects of social interaction. We wish to examine the benefits for the individuals, for *ad hoc* cooperate groups, and for the system as a whole. The when-aspect considers the conditions under which interaction takes place; the why and to whom aspects are concerned with the motivation for interactions and the agent interrelationships and dependencies; the how-aspect covers the protocols used to govern the interactions and the nature of the commitments and attitudes adopted while performing the roles. Space restrictions mean we primarily cover the when and how aspects here, although other aspects are brought in where appropriate. In the remainder of this section we consider exemplar interactions involving the three distinct types of agent that we previously identified.

3.1 HELPFUL AGENTS OFFERING SUPPORT

Responsible agents can be helpful if they do not have a lorry at their bay or if there is a lorry with a discharge time greater than T_{max}. If an agent (H) finds itself in this situation, it may offer to help a responsible agent (R) which needs support in discharging its lorry. H will offer to help if it is concerned with the achievement of the wider group goal or if it believes that by assisting now it may receive some reciprocation in the future. To determine which agent to assist a modified form of the contract net protocol

[4] is used—in our case, however, the agent initiating the interaction is offering to assist the others rather than seeking assistance from them. H sends out a `help-request` when:

- it has finished discharging a lorry at its own bay and either no new lorry has arrived or a lorry has arrived with a discharge time greater than T_{max} (H has been acting responsibly and now wants to act helpfully);
- it has finished discharging a lorry at one of the other agent's bays and either no new lorry has arrived at its own bay or a lorry has arrived with a discharge time greater than T_{max} (H has been acting helpfully and decides to continue in this vein because there is no reason to become responsible gain).

Responsible agents unable or unsure (in that they are currently receiving unreliable help) about discharging their lorry in time will respond to this offer. If H receives more than one response it must decide whom to help. At its simplest level, H may choose to assist the agent which is in most need of its support. Alternatively more complex strategies which factor in information about the distance between R's bay and H's current location and previous experiences of social interactions with R (e.g. if R helped H before then H may decide to return this favour) may also be used. Having made its selection, H sends an `accept` message to the agent it has decided to support and `reject` messages to the others. If no R needs support to meet its deadline, H can decide to remain inactive or support an agent with a deadline which it can meet alone but which could be shortened by having more agents working on it. In the latter case, this non-crucial help may be worthwhile because discharging a lorry before its desired departure time may increase the warehouse's throughput.

3.2 RESPONSIBLE AGENTS SEEKING SUPPORT

A responsible agent (R) will actively seek assistance from other agents if it has a deadline that it cannot meet alone. R sends a `support-request` when:

- a lorry with a discharge time less than T_{max} arrives at its bay immediately after a discharged lorry left (R was acting responsibly and remains in this role);
- a lorry with a discharge time less than T_{max} arrives at its bay while it is supporting another agent (R was being unreliably helpful, but dropped this support to fulfil its responsible tasks).
- an unreliable helpful agent has withdrawn its support and the lorry in R's bay can no longer be unloaded by its desired deadline (R was being assisted with its responsible task by an agent which withdrew its commitment, R needs to find new support).
- it is engaged in an ongoing cooperative commitment and a lorry arrives at its bay with a discharge time less than T_{max} plus the time it is already committed to

support another agent (R is a cooperative agent and remains in this role with respect to the ongoing commitment).

Helpful agents respond to `support-request` messages if they are idle or engaged in non-crucial help. Responsible agents who are prepared to enter into a cooperative agreement (either one-shot or ongoing) may also respond.

If R receives multiple replies, it must decide from whom it will get the most suitable support. This can be determined by examining the following preference ordering:

i) Reliable helpful agent who can assist for exactly as long as is needed;
ii) Reliable helpful agent who can assist for longer than is needed;
iii) Several reliable helpful agents who individually have less time to support than is needed but when aggregated together are able to provide as much time as is needed (or more);
iv) Several unreliable helpful agents;
v) Cooperative agent(s).

R prefers the support of helpful agents, over cooperative ones, because it need not give a commitment to support the other agent in return. To fulfil its responsible task without having to reorganise its problem solving, R prefers the help of reliable to unreliable agents since this is a more predictable form of interaction. R prefers as few agents as possible to be in the collaboration because larger groups require greater communication and coordination overhead. Finally, R prefers the exact degree of support, rather than over support, because it is more globally beneficial.

Since agents operate in a concurrent and asynchronous environment, there are possibilities that the ordering of certain events or messages may significantly affect the behaviour of the overall system. An important instance in which this can occur is when agents have sent a `help-request` and additionally have responded to a `support-request`. A responsible agent who receives bids related to its own `support-request` and also some `help-requests` should give help requests a higher preference because these agents definitely offer unconditional help.

4. CONCLUSIONS & FUTURE WORK

This paper identified several distinct types of social attitudes with regard to their volition to cooperation. A crucial part of this consideration was the type of commitments the agents were willing or able to give one another and the way the agents reasoned about and with these commitments. Our solution aims to enable the agents to find a balance between ‘behaving too cautiously’ and ‘harmful commitment overloading’ (agents committed to something they cannot fulfil). Our classification contrasts with other work [3] which has also highlighted different levels of commitment. Two com-

munication protocols were described and their invocation was characterised by the social attitude of the participants.

The loading platform scenario requires the agents to cooperate and coordinate their behaviour to solve the overall problem in an effective manner. Our solution provides a means of maximising the predictability of the agents and the system as a whole. The scenario is powerful in that it provides ideas about cooperative scheduling and planning [1] as well as collision/conflict avoidance [2]. Its richness presents numerous possibilities for studying different patterns of individual, group and social behaviour.

For the future, we are looking to put these preliminary ideas into a more general context. We believe the agent's motivation and the priorities it attaches to individual and social goals are crucial determinants in governing individual and social problem solving. This is not only important for the individual decision making but also to predict the behaviour of others and to react accordingly. We will further distinguish different types of individual and social commitment, different strategies for preventing or minimising commitment overload, different means of handling existing commitment conflicts, different strategies for minimising the damage of dropping important commitments, and strategies for improving trustfulness and reliance. We will also describe more communication protocols to support richer forms of interaction. Finally, the performance of the various agent types will be compared by means of experimental evaluation in the loading platform scenario.

REFERENCES

[1] T. MOEHLMAN, V. LESSER, and B. BUTEAU: *Decentralized Negotiation: An Approach to the Distributed Planning Problem*, Group Decision and Negotiation, 2:161-191, 1992.

[2] J.P. MÜLLER, M. PISCHEL, and M. THIEL: *Modeling Reactive Behaviour in Vertically Layered Agent Architectures*, in: M.J. Wooldridge and N.R. Jennings (Eds.): Intelligent Agents, LNAI 890, Springer-Verlag, 1995, 261-276.

[3] T. SANDHOLM and V. LESSER: *Issues in Automated Negotiation and Electronic Commerce: Extending the Contract Net Framework*, 1st International Conference on Multi-Agent Systems, San Francisco, CA, June 1995, 328-335.

[4] R.G. SMITH: *The Contract Net Protocol: High-Level Communication and Control in a Distributed Problem Solver*, IEEE Transactions on Computers C-29 (12): 1104-113, 1980.