Identifying opportunities and constraints for goal achievement through relationship analysis

Ronald Ashri and Michael Luck
University of Southampton
Southampton, SO17 1BJ, United Kingdom
{ra,mml}@ecs.soton.ac.uk

Mark d'Inverno
University of Westminister
London, W1M 6UW, United Kingdom
dinverm@westminster.ac.uk

Introduction

The ability of agents to interact in order to collectively achieve goals is one of the central arguments for the utility of multi-agent systems. Such interactions take place whenever one agent performs an action which, intentionally or not, affects one or more other agents. Thus, when agents interact we can say that they are *related* by virtue of the fact that they are affecting each other. Identifying, analysing and understanding the implications of the various types of such relationships is of critical importance, since they can have both beneficial and adverse effects on the performance of the overall system and the individuals within it.

Therefore, if coordination and regulation of agents is to be achieved as an agent society evolves, either by external intervention or through interventions by the agents themselves, we require some means of *identifying the relationships between agents at run-time*. Of course, this information is only useful if we are also able to determine how the identified relationships may impact on individual agent operation and the system as a whole. Thus, we also require a principled and comprehensive typology for *characterising agent relationships*. In this paper we outline just such a model.

Model of Agent Interaction

The notions that underpin this model are based on the SMART framework [2], and are discussed in more detail in [1], so they are only briefly described below.

In essence, we consider agents as described by *attributes*. Agents operate within an *environment*, which is also described by attributes, and perform actions that can either change attributes (*actuator* capabilities) or retrieve the value of attributes (*sensor* capabilities). Agents also pursue *goals*, which are desirable environment states described by non-empty sets of attributes.

Given that agents interact with the environment through actuators and sensors, and that the environment as a whole is defined through a set of attributes, we can intuitively think of actuators and sensors as defining *regions of the environment*, or subsets of the entire set of attributes that make up the environment. The attributes that an agent's actuators could *possibly* manipulate define a *Region of Influence* (RoI), while the attributes that an agent's sensors could *possibly* perceive define a *Viewable Environment* (VE).

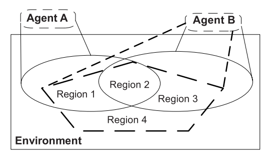

Figure 1. Observable Actions

The *VE* and the *RoI* of an agent provide us with a model that relates an agent and its individual capabilities to the environment. In order to model how two agents may interact, we need to consider how their *VEs* and *RoIs* overlap. The different ways in which these overlaps occur plays a role in determining the possible relationships between them. In the rest of the paper we illustrate these concepts by using an ellipse to represent the *VE* and a pentagon for the *RoI*.

Relationship Analysis

We develop a comprehensive typology of interactions that can provide the building blocks for defining a wide range of different relationships. To do this, we systematically examine all the salient possibilities for interactions between just two agents. First, we consider the possible types of interaction when actions of other agents can be observed, and then consider the possibilities when actions can be directly influenced by other agents due to overlapping *RoIs*.

Mutually Viewable Environment When only the *VEs* of agents overlap, irrespective of the *RoIs*, we can simply identify that there is some region of the environment that is viewable by both agents.

AAMAS'04, July 19-23, 2004, New York, New York, USA.

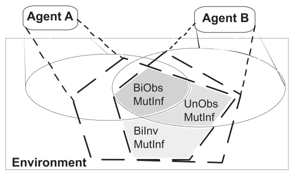

Figure 2. Observable Mutual Influence

Influenced Viewable Environment The next step is to introduce the *RoI* for just one agent, as illustrated in Figure 1. In Regions 1 and 2, the the results of B are visible to A, defined as *Observable Actions*. Region 1 defines *Unilaterally Observable Actions*, since only A can observe the actions of B, while Region 2 defines *Bilaterally Observable Actions*. Similarly, in Regions 3 and 4 we have *Invisible Actions*, since A cannot observe the actions of B. For Region 3, these are *Unilaterally Invisible Actions*, while for Region 4, they are *Bilaterally Invisible Actions*.

Mutual Influence We now move on to examine the situations in which agents can influence actions of each other, by introducing *RoIs* for *both* agents. In the first instance we can say that two agents are able to directly influence each other if their *RoIs* overlap, a situation of *Mutual Influence*.

We define the relationship by which Agent A can observe the region of mutual influence as *Observable Mutual Influence*. Similarly, A may not be able to observe this region of mutual influence, in which case we have *Invisible Mutual Influence*. Having provided definitions from one agent's perspective, we consider the situation in which both agents' *VEs* are examined. The first case is *Bilaterally Observable Mutual Influence*, in which both agents can observe the mutual influence area, as illustrated in Figure 2. The region in question is where both agents' *RoIs* overlap as well as their *VEs*. Similarly, we can also have a situation of *Bilaterally Invisible Mutual Influence*. Finally, we define the situation in which one agent unilaterally observes the region of mutual influence as *Unilaterally Observable Mutual Influence*.

Goals Region Analysis

In addition to identifying relationships based on an analysis of their *VEs* and *RoIs*, we can also gain a more detailed picture by including information about the actual goals that agents are pursuing. With the additional knowledge of these goals, we can narrow or expand the space of possible relationships by identifying interactions that an agent may pursue that are beyond its range in terms of its *RoI* or its *VE*, or by excluding those within its *VE* and *RoI* that it will not pursue.

We distinguishing between *query* and *achievement* goals, with the former representing goals where some information is elicited while the later represent goals where changes in the environment are required.

No control The agent has a goal that describes an environmental state falling outside of both the agent's *VE* and its *RoI*. As a result, this agent has no *control* over satisfying that goal, irrespective of whether it is a *query* or *achievement* goal.

View control In this case, the agent can satisfy a *query* goal but not an *achievement* goal, since the goal is within the agent's *VE*.

Total control A *total control* goal is one that lies both within the agent's *VE* and *RoI*. As a result, regardless of whether it is a *query* or *achievement* goal, the agent can satisfy it.

Blind Control In this case, the goal falls within the agent's *RoI* but not within its *VE*. As a result, the agent is able to satisfy it if it is an *achievement* goal but not if it is a *query* goal. However, the agent is not able to *verify* the results of its actions.

Partial Control Finally, a goal may fall in a region that is partially under the agent's *VE* or the agent's *RoI*. In this case, the agent will have some combination of control based on the four types described above.

Conclusions and Further Work

In this paper we introduced methods for relationship analysis, building on a basic model of interaction between agents and the environment. Furthermore, we related these relationship types to the goals of an agent, by defining goal regions. The combination of relationship analysis with information on goals can provide a useful tool for identifying possibilities for coordination between agents, especially in situations in which we cannot predefine coordination because of incomplete information about an agent's capabilities and goals. Further work will include the development of appropriate tools to allow the automated analysis of agent relationships, so that they can be incorporated within the toolkit of agent application developers as well as the integration of such techniques within existing agent methodologies. Furthermore, we aim to apply the system as a means of providing a general analysis of agent systems, aiming at the definition of metrics revealing issues such as the level of *potential interference* between agents.

References

[1] R. Ashri, M. Luck, and M. d'Inverno. On Identifying and Managing Relationships in Multi-Agent Systems. In George Gottlob and Toby Walsh, editors, *Proceedings of the 18th International Joint Conference on Artificial Intelligence*, pages 743–748. Morgan Kaufmann Publishers, 2003.

[2] M. d'Inverno and M. Luck. *Understanding Agent Systems*. Springer, 2nd edition, 2004.