An Agent's Perspective of a Team in a Dynamic World

Madhu Goyal

Faculty of IT
University of Technology
PO Box 123
Broadway NSW 2007
Australia
Email: madhu@it.uts.edu.au

Abstract: In this paper, we propose an Attitude Based Team Model (ABTM) in which we argue that team members' perspective or attitude is a very important attribute for the performance of a team. Our team model presents team as a *collective abstract* attitude, in which is embedded a novel way of solving problems and conflicts in our domain. We argue that this collective attitude is further decomposed into the individual attitudes of the agents towards various team attributes (Team definition, team methods … etc.). We then evaluate the different types of the agent teams in a simulated fire world using teams with and without different types of attitudes. The application and implementation of this model to a virtual fire world has revealed a promising prospect in developing team agents.

1. Introduction

A key issue in multi-agent systems research is how can heterogeneous and homogenous agents interact to form a team [6][9][10]. However, building an effective, realistic and dynamic team is very difficult and cumbersome. If the agents are required to be in a team, they can subsequently fail for many different reasons. For example, a team of agents that agree to cooperate in principle may discover that the assumptions upon which their choices were made do not in fact hold. Alternatively, events beyond the control of the team may make successful completion of their cooperation impossible. The above-mentioned research concentrate mainly on the pre-defined team of agents and to deal with uncertainties of complex dynamic domains in which these teams have precompiled monitoring plans. Examining this, we propose that the team agents need a higher-level approach to deal with complexities of the world or domain to fulfil their responsibilities.

The principal aim of this paper is to find how the team agents can exhibit coherent teamwork in the hostile, complex, dynamic domains. Teams are inherently paradoxical in nature and comprise of apparently contradictory elements, each of which is true. The differences of individual beliefs, intentions and plans cannot be accomplished, yet team demands that these different insights be combined or

integrated so that the agents act as one. To surmount uncertainties in cooperation and maintain coherence in teamwork, we argue that each team members maintain its own "view" or "attitude" about the team it is a member of. In this paper, we present an attitude-based model of teamwork (ATBM) that is suitable for team agent problem solving in a hostile dynamic multi-agent environment. We aim to design teams, which can survive in a virtual, hostile, dynamic fire world and solve problems with other similar agents inhabiting the same world. In addition, they will be capable of deriving plans and negotiating over the sharing of resources in order to achieve a common goal. The simulated fire world domain provides us with an excellent opportunity to design agents, which are capable of exhibiting appropriate team behaviour.

2. Attitude Based Team Model (ABTM)

The teamwork succeeds most dramatically when team members are enthusiastically unified in pursuit of a common objective rather than individual agendas. To achieve this common objective each member of a team should have an explicit model of team particularly when the world is dynamic and the agents are highly autonomous. The problem of modelling the activity of team of agents [5] is a combination of two sub problems: the first is the modelling of the team itself [11] and the second is the modelling of the team activity [8]. Unfortunately, in implemented multiagent systems, team activities and the underlying model of teamwork is often not represented explicitly [6].

The most popular theory on teams is that the team activity is achieved only if the agents have the joint intention towards that activity [7]. It focuses on a team that jointly intends a team action if the team members are jointly committed to completing the team action, while mutually believing they are doing it. To enter into a joint commitment, all team members must establish appropriate mutual beliefs and commitments. Thus the joint action by a team involves more than just the coincidence of simultaneous individual actions [3]. In this paper, we present a team model based on attitudes. We advocate a bottom-up view or participant's view of team instead of designer's view as traditionally done in [3][6][9]. In this *bottom-up* view or the *participant's* view, we claim that team activity is achieved only if the agents have team as a *collective abstract* attitude. From this *collective* attitude, agents derive *individual* attitudes that are then used to guide their behaviours to achieve the team activity. Attitude is probably the most distinctive and indispensable concept in contemporary social psychology [4]. The *Angus & Robertson* English dictionary defines attitude as a mental view or disposition, as it indicates opinion or allegiance. Humans often adopt definite attitudes towards objects in the world while they deal with them during their course of interactions with the objects. We define attitude as follows: an attitude is a **built-in predisposition to respond in a consistently favourable or unfavourable manner with respect to a given object**. Behaviours exhibited by an agent in a multiagent environment can be either individualistic or collective. Accordingly, we can divide attitudes in two broad categories: individual attitudes and collective attitudes. We consider two agent team in this paper i.e. A_1 and

A_2. So the collective attitude of the agent A_1 and A_2 towards the collection team is represented as $\mathbf{Team}_{\mathbf{A1\,A2}}(A_1, A_2)$. But from A1's viewpoint, team is an attitude that it is holding towards the collection (A_1, A_2) and can be denoted as $\mathbf{Team}_{\mathbf{A1}}(A_1, A_2)$. Similarly from A_2's viewpoint, its attitude can be denoted as $\mathbf{Team}_{\mathbf{A2}}(A_1, A_2)$. But the collective attitude $\mathbf{Team}_{\mathbf{A1A2}}(A_1, A_2)$ is decomposed into the individual attitudes only when both the agents mutually believe that they are in the team. In order to establish mutual belief between the agents, the agents have to commit to the team activity by saying so.

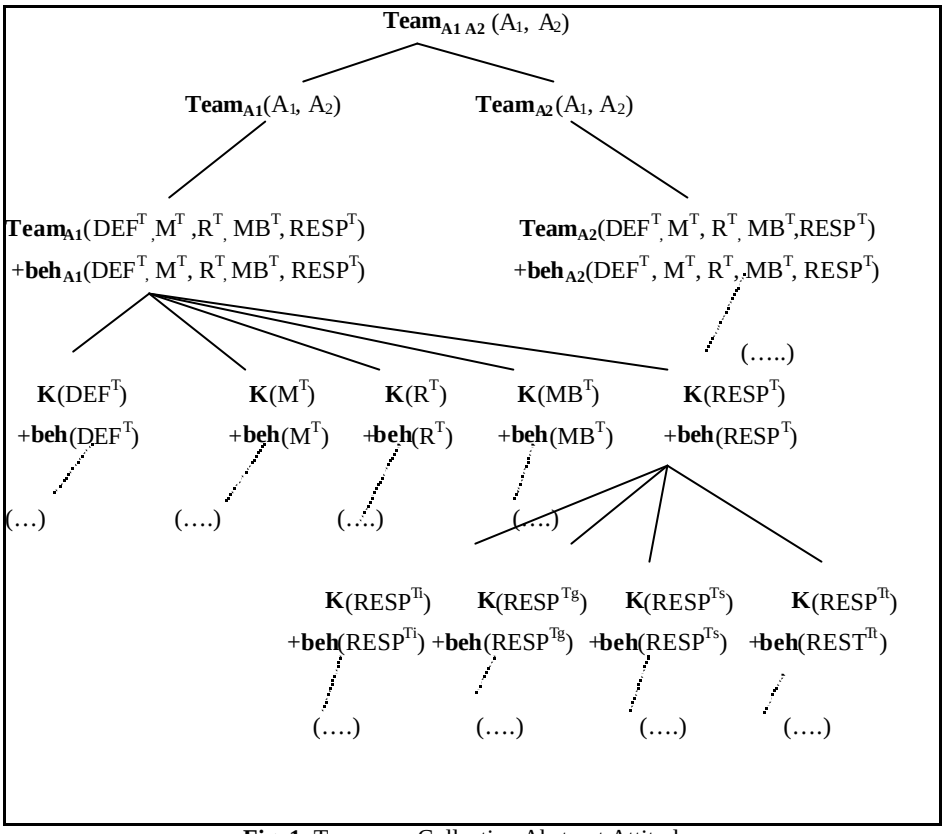


Fig. 1: Team as a Collective Abstract Attitude

The collective abstract attitude of a team can be represented in the form of a hierarchical tree as shown in figure 1. Besides mutual belief, we model a team using a set of five attributes i.e. team definition (DEF^T), team methods (M^T), team rule base (R^T), mutual belief (MB^T) and team responsibility ($RESP^T$). Thus $\mathbf{Team}_{\mathbf{Ai}}(A_1,A_2)$, the abstract attitude of a team agent can be divided into two components (i) attitude towards team definition (DEF^T), team methods (M^T), team rule base (R^T), mutual belief (MB^T) and team responsibility ($RESP^T$) (ii) behaviours towards the components of team structure i.e. towards that team definition (DEF^T), team methods (M^T), team

rule base (R^T), mutual belief (MB^T) and team responsibility ($RESP^T$). The overall attitude towards the team attributes is further divided into the attitudes towards each component using the same attitude decomposition technique described above. The attitude towards the team definition (DEF^T), team methods (M^T), team rule base (R^T) can not be further decomposed, so the team agents holding these attitudes result in different behaviours. But the team responsibility ($RESP^T$) is further divided into four components as described in section 2.5. The attitude tree will keep decomposing until all the attitudes are converted into the resulting behaviours. Thus the team responsibility can be further divided into sub-parts like individual responsibility ($RESP^{Ti}$), group responsibility ($RESP^{Tg}$), social responsibility ($RESP^{Ts}$) and team responsibility ($RESP^{Tt}$). The abstract attitude towards the team responsibility results into (i) attitudes towards the individual responsibility ($RESP^{Ti}$), group responsibility ($RESP^{Tg}$), social responsibility ($RESP^{Ts}$) and team responsibility ($RESP^{Tt}$) (ii) behaviours towards the individual responsibility ($RESP^{Ti}$), group responsibility ($RESP^{Tg}$), social responsibility ($RESP^{Ts}$) and team responsibility ($RESP^{Tt}$). However, all these different attitudes towards the team attributes ultimately convert into different team behaviours.

2.1 Team Definition (DEF^T)

The team definition attribute (DEF^T) specifies the definition of the team. Each agent uses this definition to verify whether the team exists at any point of time. The team definition can be implemented in a number of ways. However, the most important point in the implementation of team definition is that each agent in the team is individually aware of the team, and "consciously" supports the team as an abstract entity by performing appropriate team definition activities. In a hostile dynamic environment, the world state and the behaviours of the agents are unpredictable. In order to keep pace with the changing world an agent has to verify, whether the agents are still in the team or not. In our implementation of team definition, each agent sends a query to other agent asking whether it is in team or not. If the answer is positive the team exists, otherwise it does not. Accordingly, every agent in a team will follow set of rules defined by team eg. to check whether two agents A_1 and A_2 are in a team or not, we have a pair of question answer session between two agents. To establish this session A_1 first asks A_2: Are you in the team? If A_2 is in team it will reply with the answer: Yes, I am in the team. If A_1 does not receive an answer within a reasonable time, or the received reply is "No", it will unform the team. Thus the function of team definition attribute (DEF^T) is to check whether agents are in the team or not. In our model, this team definition attribute (DEF^T) can be invoked periodically or whenever there is a change in the situation.

2.2 Team Rule Base (R^T)

The team rule base (R^T) contains a set of rules, which generates various team behaviours in different conditions. The team rule base (R^T) determines what different team goals or sub-goals are to be generated under what conditions. These behaviours

of the team agents are represented in the form of condition/action pairs where conditions are logical expressions over the inputs and actions lead to behaviours when executed.

2.3 Team Methods (M^T)

In our team model, the team method (M^T) attribute specifies the problem solving methodology for achieving the team goals. A team problem can be solved in a number of different ways. When the team agents examine the problem besetting the team, each team agent invokes the team methods (M^T) attribute to get information about the methods to solve the problem. The team methods (M^T) attribute provides every agent the knowledge of the particular method it is using and details of using this method in order to achieve a particular team task or goal. For example, in the fire world, the fire fighting agents should know how to put out the fire in various situations.

2.4 Mutual Beliefs (MB^T)

The mutual belief attribute (MB^T) provides every team agent with mutual belief in order to achieve the team goal. As discussed above, the teamwork by a team does not merely consist of simultaneous and coordinated individual actions, but team agents should mutually believe the commitment to the team goal. In order to establish this mutual belief, each team agent tries to be responsive and committed to the actions of the other team agent.

2.5 Responsibility ($RESP^T$)

The responsibility attribute ($RESP^T$) specifies the responsibility of each agent towards the team. Without the responsibility attribute the agents in a team cannot work out the preconditions necessary for each team action and how to act in a social or group setting. From the multiagent perspective, an agent's responsibility can be divided into the following major subcategories: (1) individual responsibility ($RESP^{Ti}$) (2) group responsibility ($RESP^{Tg}$) (3) social responsibility ($RESP^{Ts}$) (4) team responsibility ($RESP^{Tt}$). Thus the individual commitments, social commitments and group commitments parameterise our idea of responsibility. All of these agent's responsibilities provide the necessary background knowledge to guide agent's actions and allow them to exhibit appropriate behaviours in a complex, dynamic multiagent world. The benefit of this multi-level structured view of agent's responsibility is that it provides a basis for the agents to work out their activities and conflicts in such a way that they don't violate the social and group norms.

While doing the team activity in the dynamic environment, it is inadequate for the agent to commit to the team activity only. The goal of the team indicates the overall mission that the team is required to accomplish. The team agents are autonomous participants that perform specialised functions within a social setting. But besides team modules concerning with the operation of the entire team, we need to build

modules to perform specific non-global subtasks. Striving towards a team goal, however, does not imply that every action by every agent need to be team action, since conflict at a local level may occur without compromising the global goals. The conflicts in a team can be due to autonomy of agents, group-work and social obligations. In this context, social conventions and individual behaviours provide general guidelines which agents can follow. By adopting a convention, every agent knows what is expected of it, and of every other agent, as a part of collective working towards the goal, and implicitly knows that every other agent has a similar set of expectations. Thus besides doing team activity, the team agent also has to do social actions, group work and individual actions.

3. Fire World Domain

We have implemented our team ideas on a simulation of fire world FFTEAMS using a virtual research campus. The idea of simulated fire world was first given in Phoenix [2], which is a real time, adaptive planner that manages forest fires in simulated environment. The virtual campus is implemented using C++ on Windows98/NT platform, where more than 40 agents share the world via network. FFTEAMS is a dynamic, distributed, interactive, simulated fire environment where agents are working together to solve problems, for example, rescuing victims and extinguishing fire. The fire world FFTEAMS that we have considered in this paper consists of a large number of objects (of the order of hundreds) and several agents. It consists of several buildings, an open ground area, walkways, a car park, and campus gates. Objects in the fire world include walls, buildings, furniture, open areas and LPG gas tanks. Our world is different from others' (like Air Combat [10] and RoboCup [9]) in respect that problems posed to the agents and the changes in the environment are not only caused by the actions of other agents but also by the changes the objects themselves undergo in the world (caused by the fire).

In a world such as this, no agent can have full knowledge of the whole world. Humans and animals in the fire world are modelled as autonomous and heterogeneous agents. While the animals run away from fire instinctively, the fire fighters can tackle and extinguish fire and the victims escape from fire in an intelligent fashion. An agent responds to fire at different levels. At the lower level, the agent burns like any object, such as chair. At the higher level, the agent reacts to fire by quickly performing actions, generating goals and achieving goals through plan execution. This world contains all the significant features of a dynamic environment and thus serves as a suitable domain for our team agents. Agents operating in the domain face a high level of uncertainty caused by the fire. Agents in the fire domain do not face the real time constraints as in other domains, where certain tasks have to be finished within the certain time. However, because of the hostile nature of the fire, there is strong motivation for an agent to complete a given goal as soon as possible. There are three main objectives for intelligent agents in the world during the event of fire: self-survival, saving objects including lives of animals and other agents and put-off fire. Because of the hostile settings of the domain, there exist a lot of challenging

situations where agents do the team activities. Whenever there is fire, the basic team behaviour is exhibited by the fire fighters. The fire fighters perform all the tasks necessary to control an emergency scene. The problem solving activities of the fire fighters are putting out fire, rescuing victims and saving property. Apart from these primary activities there are a number of sub tasks e.g. run towards the exit, move the objects out of the room, remove obstacles, and to prevent the spread of fire.

4. Performance

We have done several experiments for fire fighting in FFTEAMS domain to verify our ideas about the team model. To assess the utility of our Team Model a series of experiments were undertaken. The motivation behind doing these experiments is: (i) to verify the various team behaviours as described in our team model; and (ii) to determine the advantages and disadvantages of our attitude based team model. The agents in a team react to the changes in the world states by generating and achieving new goals. Meanwhile, old goals and plans are constantly being monitored and re-structured if necessary. The attitudes are mainly concerned with how to re-organise plans and goals due to situational changes. Whenever there is a new goal because of the changes in the world, one or more attitudes are usually created along with the team goal. The experiments have concentrated on evaluating the performance of attitude-based team in case of unexpected events. When a problem occurs, the attitude model stipulates a new set of behaviours for the team agents. The team agents with attitude respond to changes in the world by adopting a set of attitudes towards these changes. These sets of experiments demonstrate the significance of employing attitudes when team agents have to deal with individual, group, social and team goals in a changing world. In these experiments, we have tried to find the probability of success in case of three types of teams (i) both agents in the team having attitude-T_{both} (ii) both agents in the team having no attitude -T_{nil} (iii) only one agent in the team having attitude -T_{one}. The performance of three types of team agents is measured in terms of the rate of change of world and time taken to complete the task. Figure 2 shows the graph of the experiment measuring the probability of success of goal when the state of the world is changing rapidly. The y-axis of the graph shows the rate of change of the world, while the x-axis represents the probability of success of goal. We measure the rate of change of world in terms of the pace with which the fire spreads out. We observe severer the situations (from small fire to large fire), the more efforts the team spends in achieving the same goal (put out the fire). In the meantime, the chances of plan failures increase as the fire world becomes more hostile. This is understandable since medium and large fire creates more difficulties, and poses more challenges for team's planning and plan execution. We note that when the fire is small, the team plan seems to do well, because there are less chances of plan failure. So the number of rules fired are also less. We also found that probability of success of goal in case of attitude-based teams are more, while in case of non-attitude based team is less.

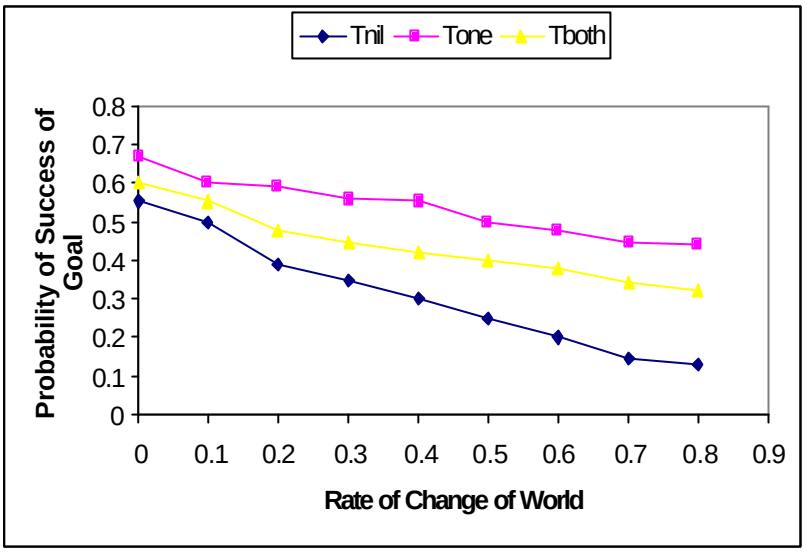


Figure 2: Probability of Success of Team Goal for the Teams T_{nil},T_{one} and T_{both}.

The attitude based team T_{both} performs better than the non-attitude base team T_{nil}, because attitude based agents can easily detect problems and can develop commitments accordingly. The non-attitude based T_{nil} teams are unable to detect problems for themselves and are simply left to complete their actions endlessly. The team T_{one} performs better than the team T_{nil}, because the agent with attitude can partially guide the team to react to unforeseen adversaries. In case of non-attitude based team T_{nil}, the agents would not be able to respond to the changes in the world. They will not be able to detect the safe and unsafe conditions of the world and not able to resolve conflicts when they have to coordinate with other agents to use some common resources. In case of attitude based team, the agents will abandon all the operations and run away from fire. Similarly, the non-attitude based agents will not be able to handle the blockage of the entrance of the room by the agent or by an object (table or chair). There are greater chances of removing the blockage at the entrance in case of attitude based agents The team agents responsibility attribute generates social and group attitudes like **help** and **full-coordination**. These new attitudes adopted will help the team agents to move the table to a safe place. Similarly, if one agent is blocking the exit, the agent has to wait for some time. The team agents' attitude generator will generate social attitudes like **wait**. Thus these experiments show that pure interest in team is not a good basis for cooperation. Participation in team problem solving requires some element of compromise i.e. team interests need to be tempered with the consideration for the individuals, groups and society. The attitude based team model helps agents to resolve conflicts due to individual goals, group goals and social goals and generate appropriate solutions to the local problems of the team agents with the help of various attitudes. Whenever a problem occurs in the fire world, the team attitude model stipulates that a new course of action should be devised. The unexpected events in the fire world cause the team plan violation and

distract the agents from its intended commitments. In order to simulate the team plan violation, varying numbers of unexpected tasks were assigned to the team during the lifetime of the team action. The agent acting as an individual, group and social agent could solve all these additional tasks.

5. Conclusion

The team model (ABTM) presented in this paper refines and formalises the notion of team as viewed as a team member. It grounds the physical and mental activities of a set of agents as a cohesive "team" in the attitudes of its individual members and towards the team attributes. The team model (ABTM) presented in this paper has the following features: (i) Team is a collective attitude (albeit implicit) of the abstract team agents. At the same time, team is also an individual, but explicit attitude of the participating agents (thus team is modelled at two levels). (ii) Team attitude is an abstract attitude decomposed into the individual agent's attitude, depicted as an attitude tree. (iii) This team model is a bottom–up view of the team. (iv) An individual agent uses mutual belief (MB) as an important team attitude. The presence of MB doesn't generate team behaviour automatically, though it is an important component present in a teamwork model. In our team model, a group of agents may perform an activity, which is based on mutual beliefs and joint intentions, but unless they are intended as a team activity, they cannot be viewed as a team activity. In a hostile and dynamic world, it is necessary that individual agents have their own attitudes towards such important attributes as MB.

Finally, we note that our team model is a deliberative model. The teams can be reactive or the agents can be simply in a team without knowing as in ant colonies [1]. The collective behaviour of social insects is not only decentralised, it is also flexible and robust: flexibility allows adaptation to changing environments, while robustness endows the colony with the ability to function even though some individuals may fail to perform their tasks. Such systems offers another way of designing "intelligent" systems, where autonomy, emergence, and distributed functioning replace control, preprogramming, and centralisation. Our team model is totally different from these type of systems i.e. our team agents are based on autonomy, deliberation and distributed functioning.

References

[1] Eric Bonabeau, Marco Dorigo, and Guy Theraulaz. Swarm Intelligence: From Natural to Artificial Systems, Oxford University Press, July, 1999.

[2] P.R.Cohen, M.L.Greenberg, D.M. Hart, and A.E. Howe. Trial by fire: Understanding the design requirements for agents in complex environments. AI Magazine, 10(3), 1989.

[3] P.R.Cohen and H.J.Levesque. Teamwork. Special Issue on Cognitive Science and Artificial Intelligence, 25(4), 1991.

[4] M.Fishbein and I.Ajzen. Belief, Attitude, Intention and Behaviour: An Introduction to theory and research. Reading, MA, USA: Addison-Wesley, 1975.

[5] Barbara J. Grosz and Sarit Kraus. Collaborative plans for complex group actions. Artificial Intelligence, 86, 1996, pp. 269-358.

[6] N.R.Jennings. Controlling Cooperative problem solving in Industrial Multi-agent Systems using Joint Intentions. Journal of Artificial Intelligence 74 (2), 1995

[7] H.J.Levesque, P.R.Cohen and J.H.T.Nunes. On Acting Together. In Proceedings of the 8th National Conference on Artificial Intelligence, 94-100. San Mateo, California, 1990.

[8] D.Kinny, M.Ljungberg, A.Rao, E.Sonenberg, G.Tidhar, and E.Werner. Planned Team activity. In Proceedings of the Fourth European Workshop on Modelling Autonomous Agents in a Multi-Agent World, MAAMAW'92, Italy, 1992

[9] H.Kitano, M.Tambe, P.Stone, M.Veloso, S.Coradeschi, E. Osawa, H. Matsubara, I. Noda, and M. Asada. The RoboCup Synthetic Agent Challenge,97, International Joint Conference on Artificial Intelligence (IJCAI97).

[10] M.Tambe. Towards Flexible Teamwork. In Journal of Artificial Intelligence Research 7(1997), pp 83-124.

[11] G.Tidhar. Team Oriented Programming: Preliminary report. Technical report 41, Australian Artificial Intelligence Institute, Melbourne, Australia.