

# ATTITUDE CYCLE FOR PROBLEM SOLVING TEAMS IN A DYNAMIC WORLD

MADHU GOYAL

*Faculty of IT, University of Technology*  
Email: [madhu@it.uts.edu.au](mailto:madhu@it.uts.edu.au)

**ABSTRACT:** In this paper, a mental attribute called *attitude* is introduced and its importance in agent problem solving is discussed. It also presents the various properties of agents describing how the *attitudes* of the agents affect the behaviour of the agents. The paper also discusses how the *attitudes* could be described computationally in terms of various attributes. This paper formalises the team as a *collective* abstract attitude of participating agents. This concept especially has been very useful in formalising the behaviour of complex teams. The team model grounds the team attitude as the individual attitude of its member agents, which in turn is further divided into the attitudes and behaviours towards the various team attributes. In this paper a team problem solving methodology is also presented, which has the notion *attitude* and *team* cycle as its core to allow robust and coherent team behaviour. According to this cycle of problem solving, the team undergoes through the steps of attitude and team formation, attitude and team maintenance and attitude and team dissolution. It also shows how these various attitudes ultimately result into various team behaviours in a fire world. The application and implementation of this methodology to a virtual fire- fighting domain has revealed a promising prospect in developing problem solving team agents.

## 1. Introduction

An agent intending to achieve a problem-solving goal must first commit itself to the goal by assigning the necessary resources, and then carry out the commitment when the appropriate opportunity comes. Carrying out the commitment typically involves planning for the goals, and executing the plans. The planning involves decomposing the complex goals, executing the plans and ultimately executing the primitive physical and communicative actions. The problem solving by a team of agents is based on a set of mutually believed facts, commitments and conventions [4][14]. These researchers have presented an abstract model of cooperative problem solving, which describes all aspects of the process, from recognition of the potential for cooperation to team action. They have developed an explicit representation of joint goals and commitments, along with an implicit (hardwired) response to commitment abandonment. In these cases, the basic idea is that responsibility for updating team members that a commitment has been abandoned

resides with the agent that first drops the commitment.

Typically, an agent that is dropping the commitment warns the others that it is doing so by explicitly sending them messages. But there are problems with this method eg. agent incapacitation and communication failure in stressed situations. This leads to view that it is not unreasonable to assign each agent the job of actively monitoring for itself the commitment of others to the joint goals. In practical terms, this means that agents should be continuously validating their commitments to the joint or team goal, and noticing when this commitment wavers [7]. 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 such uncertainties and maintain coherence in teamwork, we argue that each team members maintain its own “view” about the team it is a member of. In this paper, we propose a new mental construct called *attitude* and discuss its significance in a dynamic world. In this paper, the goal is to develop intelligent team agents having a comprehensive teamwork model based on the notion of attitudes, that identifies the key issues of team problem solving i.e. team formation and dissolution, team maintenance and team plan execution. The overall aim is to design teams, which can survive in a dynamic (virtual) world and solve problems with other similar agents inhabiting the same world.

## 2. Fire World Domain

The team problem solving ideas are implemented and tested on a simulation of fire world *FFTEAMS* using a virtual research campus. The idea of simulated fire world was first given in Phoenix [3], 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. There are two types of agents: victim agents and fire fighting agents. A part of our fire world is shown in figure 1. Our world is different from others’ (like Air Combat [12] and RoboCup [10]) 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. The communication and co-operation are the essentials to accomplish fire-fighting tasks and to achieve proper control of the hostile situation.

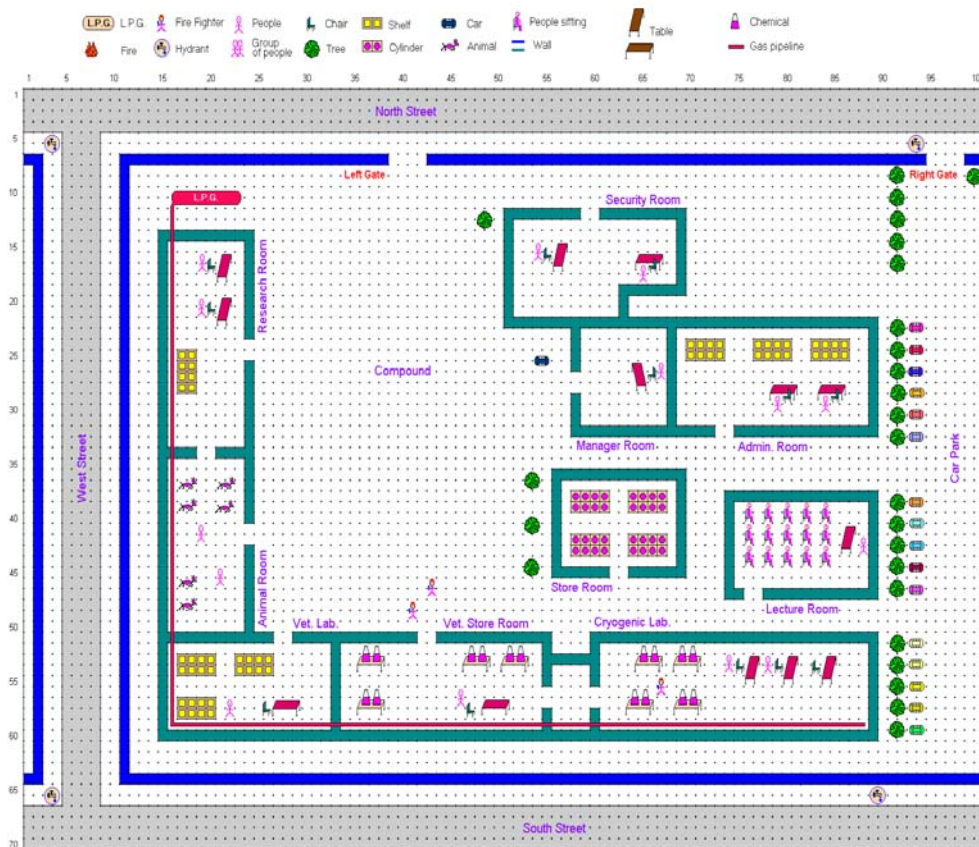


Figure 1: A Fire World

The team behaviours play an integral role in a Fire World. 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 eg. run towards the exit, move the objects out of the room, remove obstacles, and to prevent the spread of fire. In order to carry out the teamwork successfully the team of fire fighters should carry out number of procedures. There are number of preconditions of team behaviour in the fire world. *First*, the firefighters should work in pairs whenever they are in a hazardous or potentially hazardous location on the emergency scene. Firefighters working alone may over exert themselves or be unable to help themselves when trapped. *Second*, a fire attack must be coordinated to be successful. The fire fighters must perform the desired activity at the time when the officer wants them to perform. Depending on the conditions at the fire scene, the fire fighters may choose to perform immediate rescue or to protect exposures rather than attacking the fire. Coordination between team members performing different functions is crucial. For example, when carrying hose-lines the team of fire fighters requires an explicit coordination technique.

### 3. Team Problem Solving

The team's problem solving specifies the conditions under which the agents should reconsider its commitments and describes how the team agents should behave both locally and with respect to its fellow team members if any such problem arises. The team agents are provided with the methods, which clearly distinguish between the situations, in which the commitment to the team action needs to be re-examined and the actions which should be taken in such circumstances. We begin with the same notion of commitment to team goal among members of team and consider conventions or social laws that dictate the actions when an agent abandons the commitment. In addition we propose that, in order to monitor the commitment of team agents towards the team goal, agents should have the appropriate *attitude towards team maintenance*. These attitudes towards team maintenance results in appropriate agent behaviours, which help the agents to maintain the team. In the dynamic worlds, the team state will have to be constantly maintained as the agents' change their attitudes in response to changes in the world and communications received from the other agents. Thus the team problem solving in a dynamic world involves attitude maintenance as well as team maintenance that is crucial to achieving cooperation amongst the various agents, which are otherwise autonomous. In the fire world, problem solving involves the participation of victim agents and fire-fighting agents. The fire fighters in our world exist both as individual agents ( $A_i$ ) and as a team agent ( $A_T$ ) (Figure 2). The problem solving by abstract team agents ( $A_T$ ) is different from problem solving by the individual agents ( $A_i$ ) in the sense that problem solving by the team agent ( $A_T$ ) is a joint activity based on MB while the problem solving by individual agents is a joint activity dealing with the lower level execution issues or details. The team agent ( $A_T$ ) models *team* as a *collective* abstract attitude, while individual agent ( $A_i$ ) models *team* as an *individual* attitude towards the various attributes of team.




Figure 2: Fire-fighting Agents ( $A_i$ ) and Team Agents ( $A_T$ )

### 3.1 Attitude

Attitude is a *learned predisposition to respond in a consistently favourable or unfavourable manner with respect to a given object* [5]. For the purposes of paper, we do not use learning attribute of attitude. So we define attitude as *a predisposition to respond consistently in favourable or unfavourable manner with respect to a given object*. In other words, the attitude is a preparation in advance of the actual response, constitutes an important determinant of the ensuing behaviour. However this definition seems too abstract for computational purposes. In AI, the fundamental notions to generate the desirable behaviours of the agents often include goals, beliefs, intentions, and commitments. Goal is a subset of states, and belief is a proposition that is held as true by an agent. Bratman [1] addresses the problem of defining the nature of intentions. Crucial to his argument is the subtle distinction between doing something intentionally and intending to do something. The former case might be phrased as deliberately doing an action, while intending to do something means one may not be performing the action in order to achieve it. Cohen and Levesque [4], on the other hand, developed a logic in which intention is defined. They define the notion of individual commitment as *persistent goal*, and an intention is defined to be a commitment to act in a certain mental state of *believing throughout* what he is doing.

Thus to provide a definition of attitude that is concrete enough for computational purposes, we model attitude using goals, beliefs, intentions and commitments. From the Fishbein's [5] definition it is clear that when an attitude is adopted, an agent has to exhibit an appropriate behaviour (predisposition means *behave in a particular way*). The exhibited behaviour is based on a number of factors. The most important factor is goal or several goals associated with the object. During problem solving, an agent in order to exhibit behaviour may have to select from one or several goals depending on the nature of the dynamic world.

In a dynamic multiagent world, the behaviour is also based on appropriate commitment of the agent to all unexpected situations in the world including state changes, failures, and other agents' mental and physical behaviours. An agent intending to achieve a goal must first commit itself to the goal by assigning the necessary resources, and then carry out the commitment when the appropriate opportune comes. Carrying out the commitment typically involves planning for the goals, and executing the plans where planning involves decomposing complex goals, and executing the plans ultimately involve executing the primitive physical and communicative actions. Second, if the agent is committed to executing its action, it needs to know how weak or strong the commitment is. If the commitment is weak, the agent may not want to expand too much of its resources in achieving the execution. Thus, the agent needs to know the degree of its commitment towards the action. This degree of commitment quantifies the agent's *attitude* towards the action execution. For example, if the agent considers the action execution to be *higher importance* (an attitude towards the action), then it may choose to execute the action with greater degree of commitment; otherwise, the agent may drop the action even when it had failed at the first time. Thus, in our formulation, an agent when it performs an activity, since the activity is more likely that it will not succeed in a dynamic

world, agents will adopt a definite attitude towards every activity while performing that activity. The adopted attitude will *guide* the agent in responding to failure situations.

Also the behaviour must be consistent over the period of time during which the agent is holding the attitude. Thus attitudes, once adopted, must persist for a reasonable period of time so that other agents can use it to predict the behaviour of the agent under consideration. An agent cannot thus afford to change its attitude towards a given object too often, because if it does, its behaviour will become somewhat like a reactive agent, and its attitude may not be useful to other agents. Once an agent chose to adopt an attitude, it strives to maintain this attitude, until it reaches a situation where the agent may choose to drop its current attitude towards the object and adopt a new attitude towards the same object. We thus, define attitude as follows: *An agent's attitude towards an object refers its persistent degree of commitment towards achieving one or several goals associated with the object, which give rise to an overall favourable or unfavourable behaviour with regard to that object.*

We thus can view goals, beliefs, intentions and commitments as primitive forms of attitudes. In a dynamic world, beliefs are formed from direct observation, messages received from other agents, and through inference. From beliefs, arise attitudes. Depending on what the beliefs are the attitudes may take different polarity. Thus, the attitudes sometimes may be favourable and sometimes unfavourable, but what matters is the overall affect towards the object involved in the attitude. Accordingly as intentions are commitments to persistent goal [4], we can model intentions as always a favourable attitude. While in general attitudes can be specific about how soon or later a chosen activity must begin and end, intention implicitly demands that the activity start within a reasonable time. Also, once an adopted intention is dropped, it cannot be readopted; and if it is, it is usually a different intention. Thus, an adopted intention refers to a continuous and uninterrupted mental state. As opposed to this, attitudes may be continuous or intermittent. That is, an attitude by its definition may demand a mental behaviour that is distributed over time. Also from the above definition it is clear that, an attitude towards an object has the character of a commitment towards that object. In this sense it represents a dynamic assessment of a given situation with reference to an end. One fundamental consequence of such an assessment is that certain facts become relevant, others less relevant and that certain data become crucial, others less important; consequently, certain old plans may have to be revised.

As discussed above, we understand that ultimately attitudes need to be translated into appropriate behaviours. We present a computational model in which we define that every attitude **K** towards an object *x* exhibits a behaviour **beh** consisting of physical, communicative and mental actions. We represent the attitude **K** (*x*) using the following attributes:

*Name of Attitude*: This attribute describes the name of the attitude e.g. *like, hate, cautious* etc.

*Description of Object*: The description of the object contains the name of the object and a description of the internal organization in terms of the components of the object.

*Basic agent behaviour towards x*: This attribute specifies the behaviour that will be performed by the agent with respect to the object *x*.

*Evaluation*: This attribute specifies whether the attitude is favourable or not. If the beliefs of an agent are with favourable attributes, the attitude tends to be positive. Conversely, a negative attitude will result if the beliefs have primarily unfavourable attributes.

*Concurrent attitudes*: This attribute specifies any other attitudes that can coexist with this attitude.

*Persistence of Attitude*: This attribute specifies how long the attitude will persist under various situations. For example, it may specify how the attitude itself will change over time; that is, when to drop it and change it to another attitude, when to pick it up and how long to maintain it.

*Type of Attitude*: This attribute specifies whether the attitude is individual or collective.

All the attributes described above play an important part in the proper understanding of the agents' behaviours in a multiagent dynamic world. Without the knowledge of these attributes, the agent will not be able to respond appropriately to the various situations of the dynamic world.

### **3.2 Team as Attitude**

A team refers to a collection of individuals held in a dynamic and functional relationship by agreement to their mutual and combined benefit [2]. 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 [6] is a combination of two sub problems: the first is the modelling of the team itself [11][13] and the second is the modelling of the team activity [9]. Unfortunately, in implemented multiagent systems, team activities and the underlying model of teamwork is often not represented explicitly [8]. The most popular theory on teams is that the team activity is achieved only if the agents have the joint intention towards that activity [4]. 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. This is done through an exchange of request and conform speech acts [4]. Thus the joint action by a team involves more than just the coincidence of simultaneous individual actions [4]. In summary, the planned team activity, in addition to individual beliefs, goals, plans and intentions will have (a) mutual beliefs about the world and about each other's actions; (b) the joint goals that need to be achieved; (c) joint intentions adopted in response to the occurrence of an external event, or as a consequence of establishment of a joint goal. The commitment protocol synchronises the team, in that all members simultaneously enter into a joint commitment towards a team task.

We claim that team activity is achieved only if the agents have team as a *collective abstract* attitude. From collective attitudes, agents derive individual attitudes that are then

used to guide their behaviours to achieve the team activity. 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  $Team_{A_1 A_2}(A_1, A_2)$ . But from  $A_1$ 's viewpoint, *team* is an attitude that it is holding towards the collection  $(A_1, A_2)$  and can be denoted as  $Team_{A_1}(A_1, A_2)$ . Similarly from  $A_2$ 's viewpoint, its attitude can be denoted as  $Team_{A_2}(A_1, A_2)$ . But the *collective* attitude  $Team_{A_1 A_2}(A_1, A_2)$  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*.

#### 4. Team Problem Solving Methodology

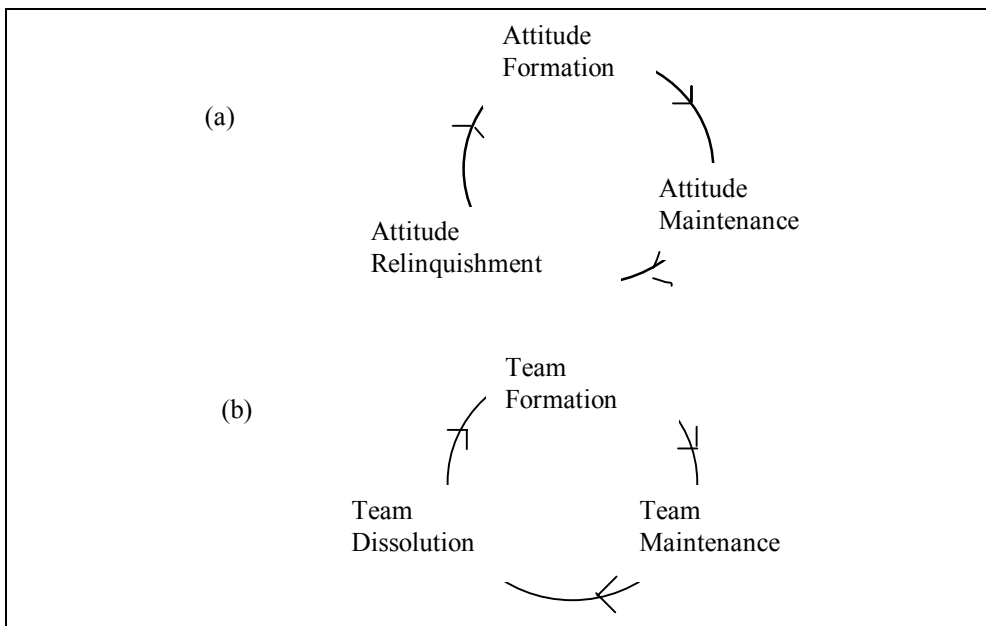


Figure 3 (a) and (b): Attitude and Team Cycle in Team Problem Solving

The team tasks do not remain static i.e. throughout the phases of teamwork, aspects of the issue or problem continues to change, as new facets become apparent. During this change some initial ideas may seem less applicable, or the implementation may be more complex than first predicted. Thus the state of team keeps on changing continuously in cyclic fashion from team formation, to team activity and to team dissolution during the lifecycle of a team. The team problem solving follows a cyclic path (Figure 3(b)) that includes the following steps (1) Team Formation (2) Team Maintenance (3) Team Dissolution. As discussed above, the commitment of agents is extremely essential in a dynamic environment for team work. To maintain this commitment of agents in the team problem solving activity, the agents should have an appropriate *attitude* towards each step of the activity. But this attitude again adopts a cyclic path of attitude formation, attitude maintenance and attitude relinquishment (Figure 3 (a)) during each step of team problem solving.



During any team problem solving activity, the agents in a team first form the attitude towards the particular team activity, then maintain that attitude towards that activity as long as they are doing the activity, and finally they drop this attitude when the team activity is finished. The figure 4 shows all the team states i.e. team formation, team maintenance and team dissolution during a problem solving activity. It also depicts the different attitudes attached to each of these team states. The figure 4 also describes the cyclic path of attitudes along with the each team state. But the attitude formation and attitude relinquishment is only visible to participating agent and not to the outside observing agent.

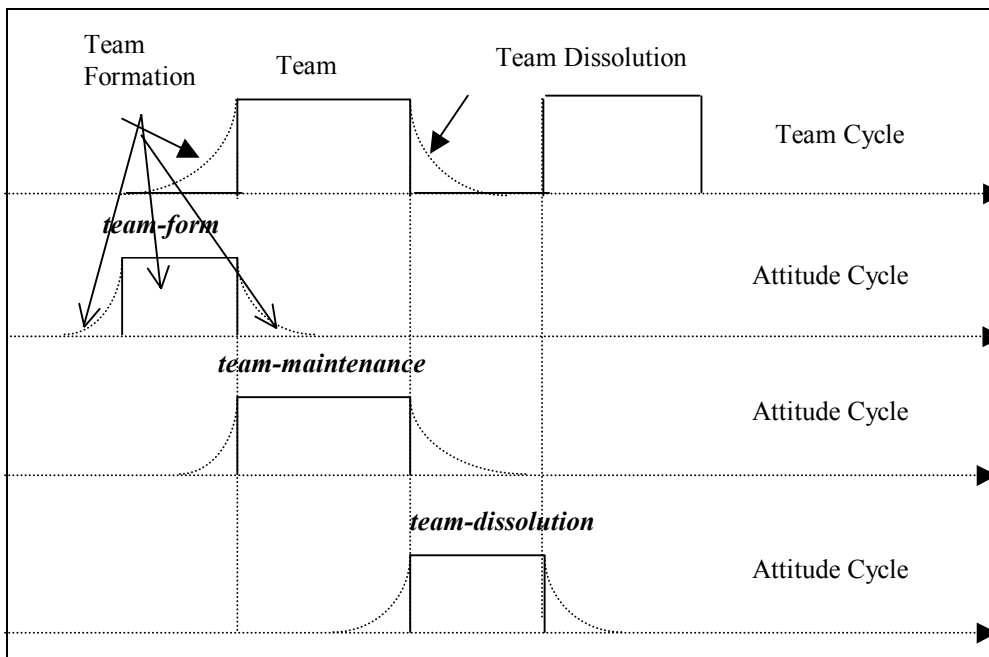


Figure 4: Observing and Participating Agent's View Point of Attitude and Team Cycle

#### 4.1 Team Formation

Impetus for attitude and team formation may arise from the world and a particular domain or from agent's themselves. Having identified the potential for team action with respect to one of its goals, a team agent will solicit assistance from some group of agents that it believes can achieve the goal [14]. Our agents form team, because the inherent nature of the world requires agents to exist as teams. However, a particular situation may force the agents to dissolve the team for some time. The agents will try to come back to the team based existence as the situations permit them. Thus in our case team existence is not motivated by any kind of goals.

The figure 5 demonstrates the progression of events in a Fire World during attitude and team formation. In the fire world, the event triggering the team formation process is a

fire. Whenever there is fire, the security officers call the fire-fighting company to put out the fire. Then the fire fighters arrive at the scene of fire and get the information about when, how and where the fire had started. Suppose there is a medium fire in the campus, which results in the attitudes *medium-fire* and *dangerous-fire* towards the object fire. The attitude *team-form* is also generated, which initiates the team formation process. We propose a dynamic team formation model, in which we consider initially the mental state i.e. the beliefs of all the agents are same. The fire-fighting agents recognise appropriateness of the team model for the task at hand; set up the requirements in terms of other fellow agents, role designation, and structure; and develop attitudes towards the team as well as towards the domain.

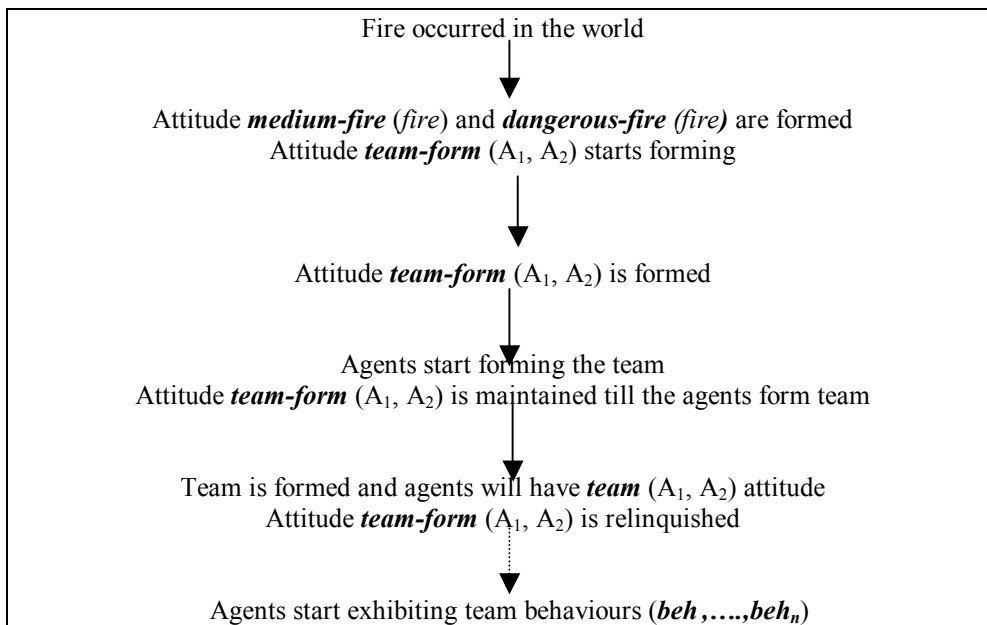


Figure 5: Progression of Events in the Fire World during Attitude and Team Formation

In order to select a member of the team, our agent will select the fellow agent who has following capabilities:

- Has knowledge about the state of other agents.
- Has attitude towards the team formation.
- Can derive roles for other agents based on skills and capabilities.
- Can derive a complete team plan.
- Can maintain a team state.
- Has unified commitment towards team goal.

Our method of forming the team is like this; the agents start broadcasting message to other agents “Let us form a team”. The agents will form a team if two or more than two agents agree by saying, “Yes”. If the agent do not receive the “Yes” message, it will again iterate through the same steps until the team is formed. The *team-form* is maintained as long as the agents are forming the team. Once the team is formed, agents

will drop the *team-form* attitude and form the *team* attitude, which will guide the agents to produce various team behaviours.

#### 4.2 Team Maintenance

As discussed above, while solving a team problem the team agents have also to maintain the team. During the team activity the team agents implement the team plan to achieve the desired team action and sustain the desired consequences. The team maintenance behaviour requires what the agent should do so that team does not disintegrate. In order to maintain the team each agent should ask the other agent periodically or whenever there is a change in the world state, whether he is in the team or not (team definition). So the attitudes like *periodic-team-maintenance* and *situation-team-maintenance* are produced periodically or whenever there is a change in the situation. These attitudes help the agent to exhibit the team maintenance behaviours.

In real world when a team of firefighters put out fire together, each team member is well aware of the team activity – each individual is not merely putting out fire on its own, but also coordinating with others to maintain the team. In order to maintain the team, each agent in the team will remain within a safe distance from the other members of the team and invoke the team definition goal whenever required.

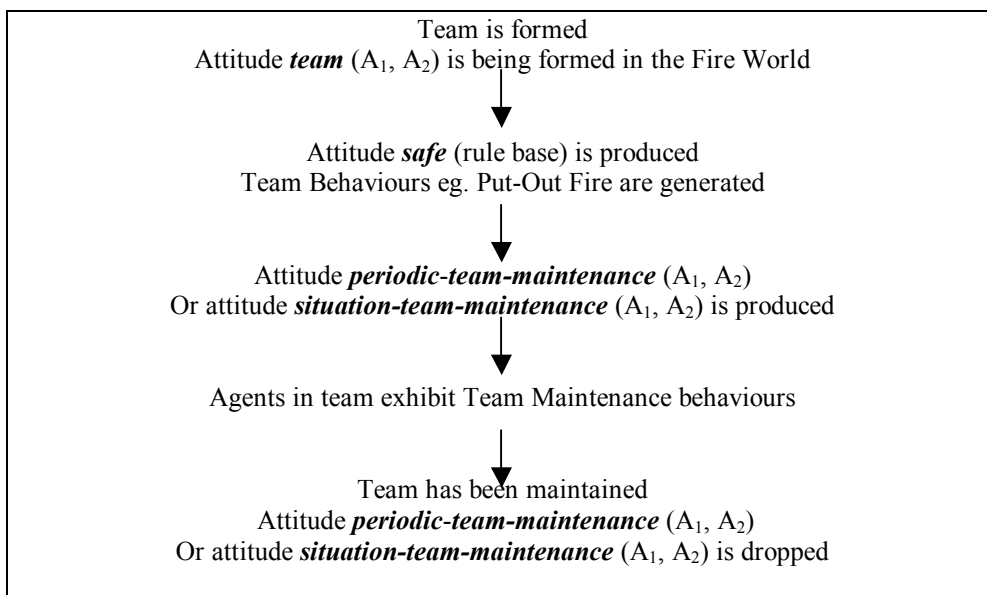


Figure 6: Progression of Events in the Fire World during Attitude and Team Maintenance

Another point of consideration in team problem solving is that attitude towards team activity as well as towards team should not change very quickly. If the attitude of the team agents changes very quickly, the problem solving becomes very cumbersome. In such cases, the team agent will spend most of the time on maintaining the team. Figure 6 describes the progression of events in a Fire World during the team maintenance.

### 4.3 Team Dissolution

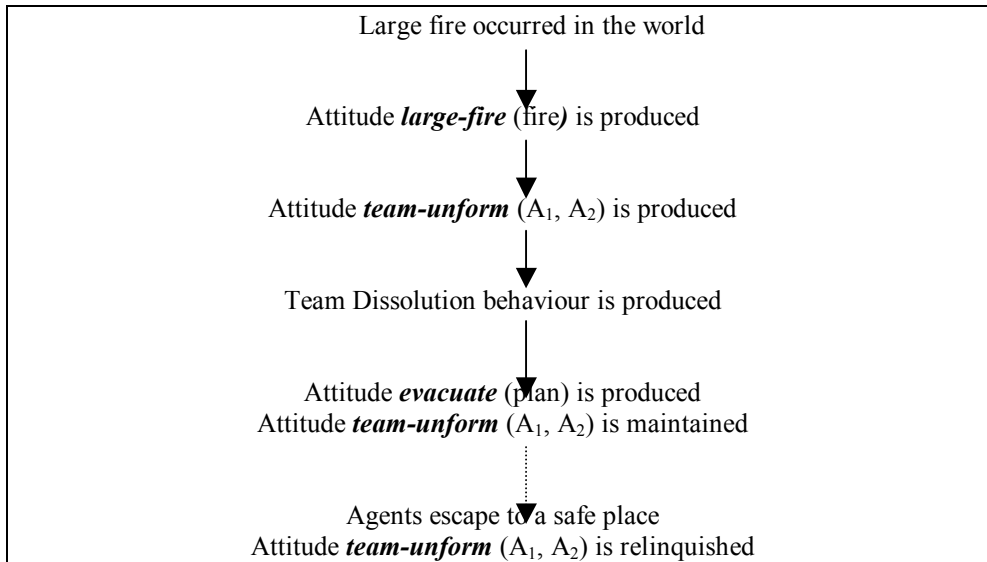


Figure 7: Progression of Events in the Fire World during Attitude and Team Dissolution

When the team task is achieved or team activity has to be stopped due to unavoidable circumstances, the attitude *team-uniform* is generated. This attitude results in the dissolution of the team and further generates attitude *escape*. For example, when the fire becomes very large, the agents have to abandon the team activity and escape. The attitude *team-uniform* is maintained as long as the agents are escaping to a safe place. Once the agents are in the safe place, the attitudes *team-uniform* and *escape* are relinquished. In case the fire comes under control, the agents again form a team by going through the steps of team formation. The figure 7 gives the progression of events occurring while deforming the team.

### 5. Implementation

During implementation, the effect of attitudes on the process of team formation and team dissolution is evaluated. In this experiment, it is shown that how four fire fighting agents ( $ff_1, ff_2, ff_3, ff_4$ ) quickly form a team, adopt roles, select a team plan to handle the fire emergency (figure 8). Two fire fighting agents  $ff_1$  and  $ff_2$  or  $ff_3$  and  $ff_4$  will form a team, if each agent has *team-form* attitude. With this attitude, one of the agents requests the other agent for team formation. If another agent responds positively to this request, the team will be formed, otherwise it will not. In order to evaluate the effectiveness of the team formation, team activity and team dissolution scheme, we carried out several experiments under different fire situations. The goal of the agents in our experiment is that agents form a team to put out the fire and unform the team if the fire is put out. In order to combat the fire, one fire-fighting agent must be supported by another fire fighting agent. In the virtual fire world *FFTEAMS*, the *joint-water* command allows a fire-fighting agent

who is holding a hose to be supported by another fire fighting agent (figure 9).

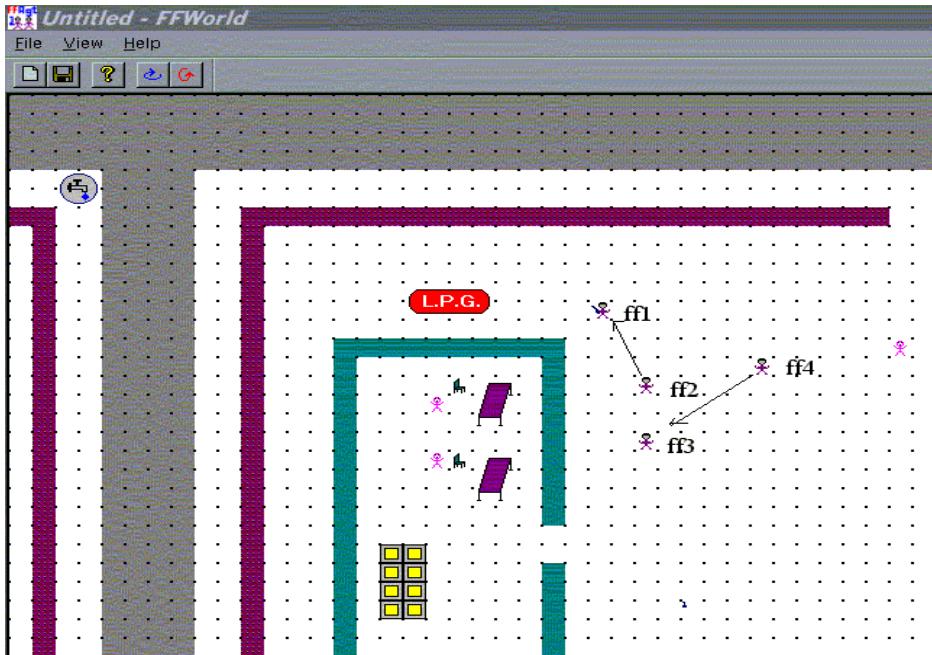


Figure 8: A Section of FFWorld Showing Fire-Fighting Agents Forming Teams

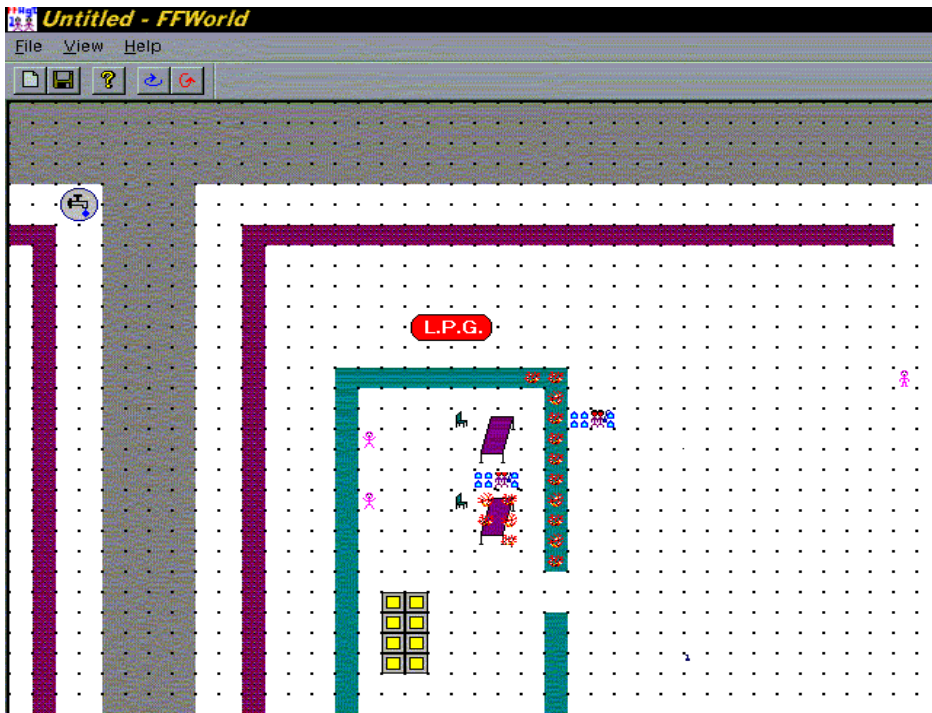


Figure 9: A Section of FFWorld Showing Fire-Fighting Agents Putting out Fire

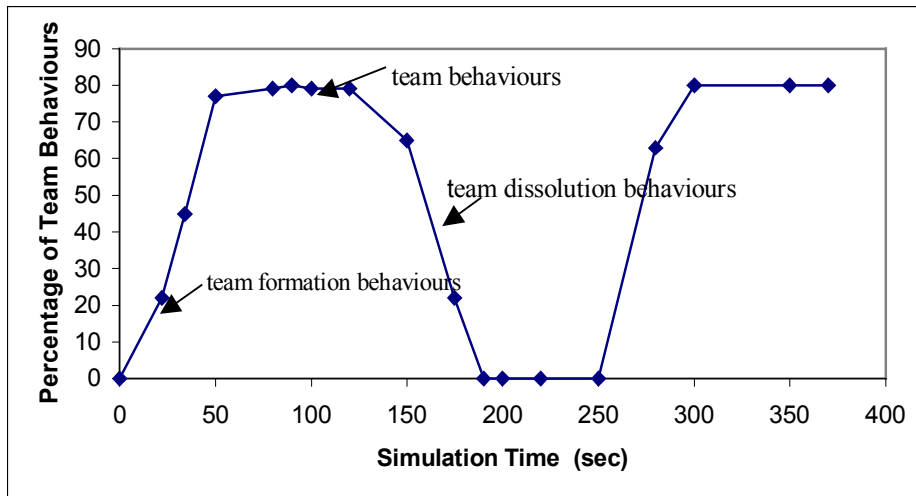


Figure 10: Percentage of Team Formation, Team and Team Dissolution Behaviours on Agent's TLS

After forming the team, the team of fire-fighters start doing the team activity i.e putting out the fire. When the fire is put off, the two agents will dissolve the team. Similarly, if the fire appears again, they will again form a team and perform the team activity. The percentage of team formation, team and team dissolution behaviours (y-axis) vs time(x-axis) are shown in figure 10.

## 6. Conclusion

The 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. This paper have also demonstrated a methodology of team problem solving, which describes all the aspects of the problem solving in a fire world based on attitude and team cycle. We have described how the agents handle team formation, team maintenance and finally team dissolution. We also divide the overall team problem solving process in two levels i.e problem solving by the team agent ( $A_T$ ) and problem solving by the individual agent ( $A_i$ ). This paper identifies several types of *collective* and *individual* attitudes required in team problem solving. The benefit of this approach is that these attitudes guide agents in a team with the aspects of when, why, with whom, and how to interact with ever changing world in a dynamic domain. The when-aspect considers the condition under which a team action takes place; the why and to whom aspects are concerned with the motivation for interactions and the agent interrelationships and dependencies within the team; and the how-aspect covers the methods adopted while performing the team tasks.

The fire-fighting 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 as well as conflict avoidance. Its richness presents numerous possibilities for studying different patterns of individual, group, social and team behaviour.

## 5. References

[1] M.E. Bratman. *Intentions, Plans and Practical Reason*. Harvard University Press, Cambridge, MA, 1987.

[2] Naomi Isgrig Brill. *Teamwork: Working Together in the Human Services*. J.B. Lippincott Company, Philadelphia, 1976.

[3] 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.

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

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

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

[7] Marcus J. Huber, Edmund H. Durfee. On Acting Together: Without Communication. In *Proceedings of AAAI Spring Symposium on Representing Mental States and Mechanisms*, pages 60-71, Stanford, CA, March 1995, AAAI press.

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

[9] David Kinny, Magnus Ljungberg, Anand Rao, Elizabeth Sonenberg, Gil Tidhar, and Eric Werner. Planned Team activity. In *Proceedings of the Fourth European Workshop on Modelling Autonomous Agents in a Multi-Agent World, MAAMAW'92*, Italy, 1992.

[10] 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)*.

[11] G.Tidhar. *Team Oriented Programming: Preliminary report*. Technical report 41, Australian Artificial Intelligence Institute, Melbourne, Australia, 1993.

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

[13] E.Werner. Cooperating agents: A unified theory of communication and social structure. In L.Gasser amd M.N.Huhns, editors, *Distributed Artificial Intelligence:Vol II*. Morgan Kaufmann Publishers, USA,1990.

[14] M. Wooldridge and N. R. Jennings. Towards a Theory of Cooperative Problem Solving. In Proceedings of the Sixth European Workshop on Modelling Autonomous Agents in Multi-Agent Worlds (MAAMAW-94).