A Hybrid Agent Architecture for Dynamic and Unpredictable Environments*

Yang Yang, Xiaoping Chen

AI Research Center, Computer Science Department
University of Science and Technology of China,
Hefei 230027, Anhui, P.R.China
blade@mail.ustc.edu.cn

Abstract. It is believed that hybrid architecture from reactive and deliberative ones will help to solve some difficulties of reasoning about and acting in highly dynamic and unpredictable environments. This paper describes our work in developing such architecture with a multi-layer model.

Keyword: Belief-Desire-Intention (BDI), hybrid architectures, dynamic and unpredictable environments, resource boundedness.

1. Introduction

The design of autonomous multi-agent systems that are required to perform in complex dynamic environments is becoming of increasing realistic importance. A lot of work has been done in this domain in recent years. Among these work is RoboCup Challenge, which offers a standard test bed to develop and demonstrate techniques in this domain.

One of the crucial problems facing designers of RoboCup teams – systems capable of effective, rational behavior in dynamic and unpredictable environments – is to ensure those agents' responses to important changes in the environment are both appropriate and timely. It seems somewhat conflict. On one side, to make rational decisions and act appropriately agent needs large amount of time for deliberation. On the other side, in such a highly dynamic domain, the environment may change at any time,

---

* This work is supported by NSFC (60275024) and High-Tech Project (2001AA422200).

which means during the period that the agent engages in reasoning, the world may change in important ways which would cause the current plan of the agent no longer be realistic or even possible. One solution is to employ reactive architectures that minimize the need for run-time computation by pre-compiling appropriate responses to situations into a form that can be utilized without reasoning. Purely reactive architecture is well-suited to some dynamic environments, but cannot guarantee it in the unanticipated situations. Furthermore, it becomes more and more difficult to design such systems as the complexity of the environment grows. Another solution is real-time reasoning systems which can behave more robustly in such situations. But such systems require more time to perform, which would cause failure in highly dynamic environment. It is clear that there should be a rational balance between deliberative reasoning and reactive acting. It is generally believed [3] that the hybrid from deliberative planning and reactive response will help to solve this difficulty.

In this paper, we present our work on constructing such a hybrid architecture that reconciles deliberation with reactivity. Deliberation is achieved by means of a real-time reasoning system that base on belief-desire-intention models [4], while reactivity attained on basis of structure similar to Brooks' subsumption architecture [6]. The rest of the paper is structured as follows. We first present our agent architecture in section 2. Section 3 and 4 will make further explanation on how this architecture works and some details of its implementation. The experiments and results about this architecture will be given in section 5. The last part of the article is conclusion and our future work.

2. Agent Architecture

To design agent architecture for domains such as RoboCup, which are highly dynamic and unpredictable, we should fully address the problem of resource-bounded and achieve a rational balance between deliberation and reactivity. For these purposes, our approach is realized as four software layers which are arranged into a hierarchy, with different levels in the hierarchy dealing with information about the environment at different levels of abstraction.

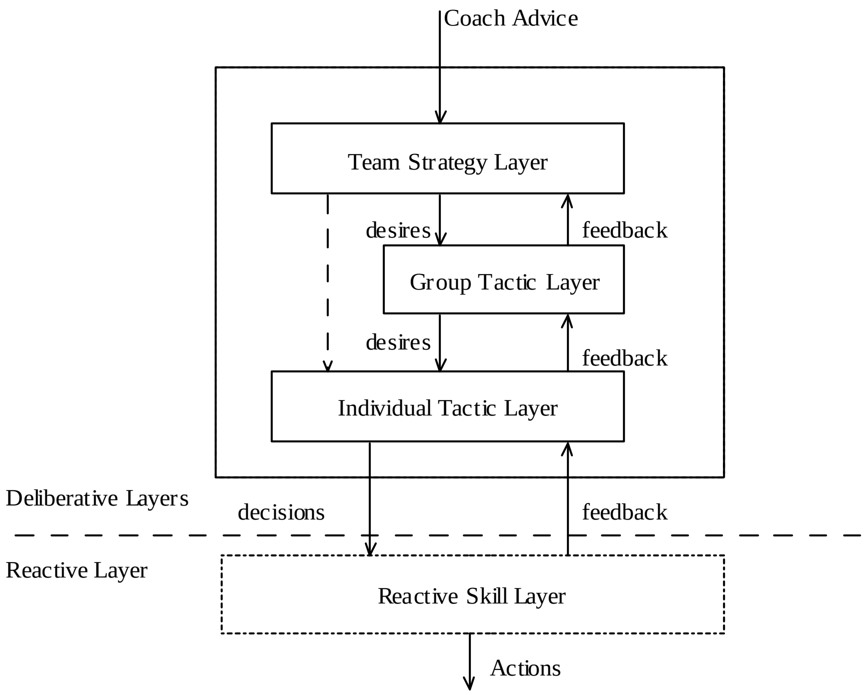

Fig. 1. Layerd Agent Architecture

As we can see in the Figure1, at the lowest in the hierarchy is a "reactive" layer, which makes decisions about what to do by pre-compiling responses. The other three are "deliberative" layers performing deliberative thinking and reasoning at different levels. The reactive skill layer implements typical basic capabilities of the agent, e.g. dribbling, kicking of RoboCup agents. The task of the individual tactic layer is to reason about affairs concern about individual and makes decisions. The top two levels – group tactic layer and team strategy layer – deal with the social aspects of the environment. We thus represent other agents – their goals, beliefs and so on – in these layers. What the team strategy layer does is to reason about the whole team's strategy at each offence or defense and produce cooperative behaviors in the strategy level. The group tactic layer keeps track of tactics of a group of players and ensures these cooperative processes are successful. It must be pointed out that groups here are not statically assigned, they are dynamically formed according to particular situations.

In order to produce global behavior of the agent, these layers should interact with one another. In our approach, the connection between two layers is "desire". Desires are the states that agent hopes to achieve, in other words, an agent's desires are the tasks given to him. Each layer in this architecture gets tasks from its predecessor and produces desires to its successor. At the same time, each layer monitors the environment and produces desires itself. That means desires from the higher level sometimes will not be adopted if there exist better selections. This aspect gives more flexibility to agent though sometimes it would cause incoherent between layers, in which case feedback from lower layers is necessary.

3. Reactive Skill Layer

The Reactive Skill Layer, which is developed on the basis of the WrightEagle2001 simulator team, provides the basic capabilities of the agent, e.g. running, dribbling, intercepting ball etc. As its name implies, this layer is a typically "reactive" layer - there are no symbolic representations or reasoning at all. The whole structure of this layer is implemented using the techniques rather similar to Brooks' subsumption architecture, as showing in Figure 2.

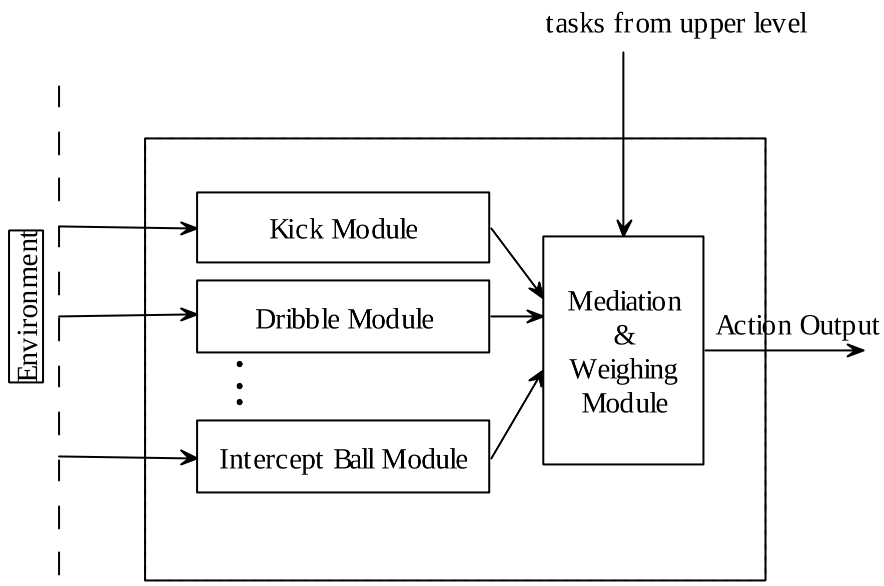

Fig. 2. Structure of the Reactive Layer

All skill modules in this layer are deemed as competing options. On every decision cycle, each of them will produce possible actions, together with their utilities and feasibilities, and send the to the "Mediation & Weighing" module. What the "Mediation & Weighing" module does is to select one of the best as the action output.

4. Deliberative Reasoning Based on BDI Model

In our approach, the three deliberative layers are implemented with identical structure. The only distinction between them is that they are working on different levels of abstraction. Figure 3 is a block diagram of such structure.

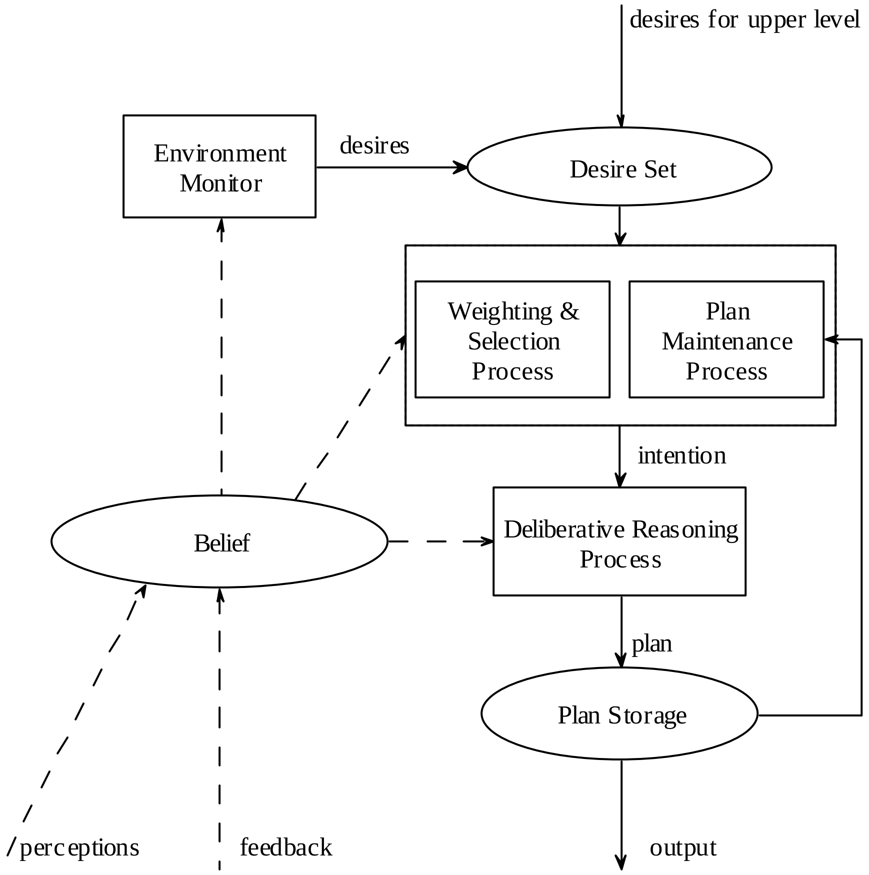

Fig. 3. BDI Model Inside Each Layer

The structure of the deliberative layers can be classified as a belief-desire-intention (BDI) architecture: the agent's mental attitudes of beliefs, desires and intentions are explicitly represented. These mental attitudes determine the agent's behaviors and are critical for achieving adequate or optimal performance when deliberation is subject to resource boundedness [4]. Intuitively, beliefs correspond to knowledge of the agent about the world. Desires are the states that the agent wants to achieve. Intentions represent the states or actions that the agent has chosen and committed resources to. Besides the information stores, which are denoted by ovals, there are four processes – denoted by rectangles: the "Weighing & Selection Process", the "Deliberative Reasoning Process", the "Plan Maintenance Process" and the "Environment Monitor". Together these constitute our deliberative structure by which the agent forms and executes plan.

Before we can start our illustration about how our architecture works, we first explain how plans are represented in our system. In domains such as RoboCup, agents are resource-bounded and knowledge-bounded, that is, they are neither prescient nor omniscient. Hence highly detailed plans about far future are of no use to them. Therefore plans represented in our approach are rather partial ones. That is, agents often decide what end states should be achieve, leaving for later deliberation questions about how to reach those states. An agent may, for example, first decide to pass the ball to a teammate, but whether to pass to him directly or via another teammate is postponed for later decisions.

Now return to the structure illustrated in Figure 3. During the runtime, the system will receive desires from higher level as its tasks. However, not all desires are generated by higher levels. Changes in the world may bring opportunities or threats, which may result in new desires to the agent. The "Environment Monitor" in Figure 3 is the component that perceives significant changes in the environment and produces desires. Another possible source of desires is the current "partial plan" that to be completed. These desires, together with those from higher level and those generated by the "Environment Monitor", will be placed in the "Desire Set".

Once desires have been generated, there are two primary problems that need to be solved. The first one is to weigh these alternative desires and decide on one of them as intention that are committed to. The second problem can be addressed as "planning problem", that is, to form a possible plan that will achieve the ultimate goal.

In our approach, the solution to the first problem is given by two processes: the "Weighing & Selection Process" and the "Plan Maintenance Process". As its name implies, the "Plan Maintenance Process" will examine desires to determine compatibility with the existing plans. Desires that may cause incoherent result with current plan will trigger the "Weighing & Selection Process". This component will analyze every alternative desires and select one of them, during which process the current plan may be abandon or reconsidered if there exists better selection.

Surviving desires from above two processes will be passed to the "Deliberative Reasoning Process" as intentions. What the "Deliberative Reasoning Process" does is establish a reasonable plan that ensures the intention to be achieved.

5. Experiments

In order to demonstrate our architecture and test its effectiveness, we developed a simulated soccer team under this architecture and have conducted a series of experiments.

The preliminary experiments, in which we chose several different styles of simulator teams that participated in RoboCup2002 and let them play games with our team, show that agents who adopt the architecture described above can well adapt to the highly dynamic domains such as RoboCup and play well. Table 1 shows the results.

Table 1.

Opponent	Num of games that played	Average Scores	Average Losses
Robolog2002	10	6.3	0.0
UvA Trilearn 2002	10	1.4	0.8
YowAI 2002	10	7.2	0.4
HelliRespina 2002	10	2.1	0.1

Another point that we interest in is how important that our deliberative reasoning in the process of decision making and whether our hybrid architecture have advantage over our previous approach. Therefore we take a number of comparative experiments. These experiments involve 4 players: a striker and 2 defenders (do not include goalie). The striker and the ball were repeatedly placed in center of the field. The initial positions of other two defenders were set to the two corner of the penalty area. The task of the striker is to break through the defense and shoot. In the experiments, we chose a striker from WrightEagle2001 (WE2001) as contrast and another 2 players as standard defenders. These games, which were over as soon as the striker lost his ball, repeatedly played 100 times for each striker. Figure 4 reveals the results of the experiments.

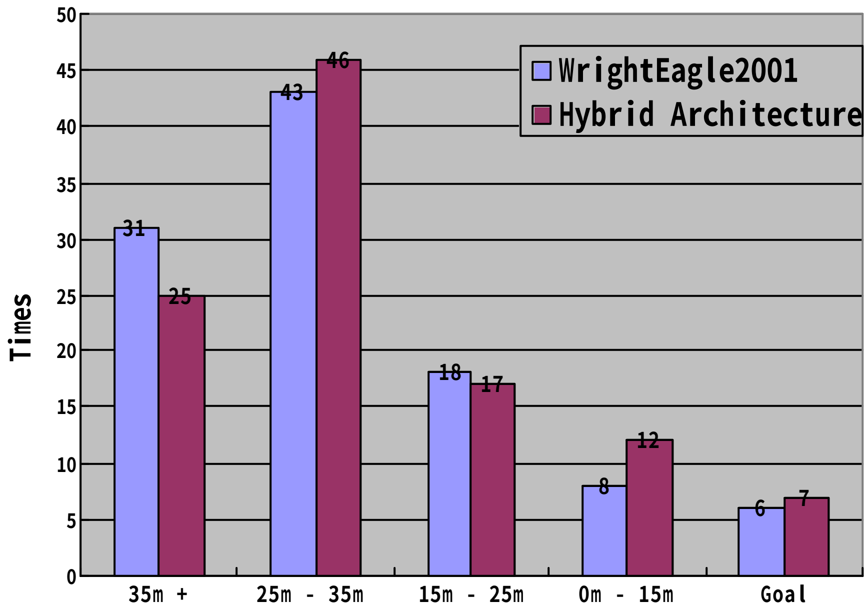

Fig. 4.

As shown in the figure, the striker from the new team successfully reaches the area whose distance to goal is 25m to 35m for 46 times, and the area "15m-25m" for 17 times, 12 times for the area "0m-15m" , while the corresponding times of the other striker is "43, 18, 8". From this result we can see that the striker implemented based on our new architecture is more successfully in attacks – he can get closer to the goal, which brings better chance. As we mentioned above, the new team is developed on the basis of the WrightEagle2001 simulator team. That means the capabilities of basic skills for both strikers are almost the same. One of the most important reasons why the new agent has the advantage in offense is that the WrightEagle2001 simulator team is implemented on the basis of decision-making theory. In many cases, the small incoherence of information may lead to incoherent decisions between two cycles, which may lead to failure. But this is not the case for the new agent. Within the BDI model, agents make commitments to their intentions which will ensure the coherence of their plans.

6. Conclusion & Future work

The experiments show that it is, to some extent, feasible to design agents with our hybrid architecture such that they can work normally in highly dynamic and unpredictable environments. However, our current work is an intermediate step to our ultimate goal of building highly competent agents with "deliberate adaptation"[8] in such environments. Until now our work mainly involves in behaviors of single agent and there are still lots of work needed to do. In the future we will focus our work on agent cooperation within this hybrid architecture.

Acknowledgement

The authors would like to thank Bin Yang and Zhenyu Li for valuable discussions and their helps.

Reference

[1]. A.S.Rao and M.P.Georgeff, Modeling rational agents within a BDI-architecture. In: J. Allen, *et al.* (Eds.) Principles of Knowledge Representation and Reasoning: Proc. of the second Intel. Conf.(KR-91). San Mateo,CA: Morgan Kaufmann Publishers, Inc.,1991. 473~484.

[2]. A.S.Rao and M.P.Georgeff, BDI Agent:from theory to Practice. Proceedings of the First International Conference on Multi-Agent Systems, ICMAS95, San Francisco, 1995

[3]. M. Wooldridge and N. R. Jennings, Intelligent agents: Theory and practice, Knowledge Engineering Review, Vol.10, No.2, 1995

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

[5]. M.E.Bratman, D.J.Israel and M.E.Pollack, Plans and resource-bounded practical reasoning. Computational Intelligence, 4(4): 349-355, 1988

[6]. R.A.Brooks, A robust layered control system for a mobile robot, *IEEE J. Rob. Autom.* 2(1986) 14-23

[7]. R. Kowalski and F. Sadri. *Towards a unified agent architecture that combines rationality with reactivity*. In D. Pedreschi and C Zaniolo, editors, Proceedings of LID-96, volume 1154 of LNAI, pages 137--149, 1996.

[8]. Xiaoping Chen, On deliberate adaptation, Proc. of CACCS-2002, *Computer Science Vol.*29 No.9 (Sup.) 2002, 49ff.