Anticipatory Planning

1 Background

One of the most difficult problems in the design of autonomous agents is not only how to make them *behave rational* from some point of view, but how to make them *stay rational* when the environment is subject to change. There has been a number of approaches addressing various aspects of the problem; Pengi [Agre-Chapman87], PRS [Georgeff-Lansky87] and [Kaelbling87] are some of the most well-known.

We will argue for the following important abilities of an autonomous planning agent that we believe are not sufficiently addressed by earlier approaches, but should be present if the problems of rationality, adaptability and robustness are exposed to increasing demands from interactions with real-world environments.

(1) The agent should be *active* in its interaction with the environment. In order to deal with events in the real world the agent should be able, not only to predict what will happen, but it should be able to pre-adapt itself for the occurence of a crucial or time-critical event. Moreover, it should be able to issue actions that prevents an undesired event to happen, or to issue actions that brings the event or its consequences under control.

(2) The agent should have an *model* of the environment and of *itself* as a part of that environment. When the environment (or itself) changes the agent should be able to update or exchange its model. This implies that the agent must be able to recognize the fact that there is a discrepancy between the model and the environment (or itself).

(3) There should be a clear separation between the model and the reasoning process of the agent. The agent should be able to distinguish between reasoning *in the model* and reasoning *about the model*.

(4) The agent should have an *introspective* ability, i.e. it should have access to its own internal structures, operations and behavioural potential.

(5) The agent should have a *reflective* ability, i.e. it should be able to reason and deliberate about its situation and embedding context.

(6) The agent should be able to relate symbols to objects and phenomena in the environment (extensionality) and, wishfully, have some representation of their meaning (intensionality).

(7) The agent should have an ability for *learning*. If the learning process can be connected to a representation of semantics as outlined in the previous point, then the agent should be able to use *induction* as a basis for learning.

We suggest that the concept of *anticipatory system* is used as a framework for specifying, analyzing and designing rational agents in real-world environments. The approach itself will not solve the problems but we believe it offers interesting potentials for the inclusion of the abilities sketched above.

The understanding of anticipatory systems as presented in the next section has emerged within general systems theory as a model of such biological phenomena as adaption, learning, evolution and other basic organic behaviours. From a mathematical point of view the anticipatory systems form a class that is opposed to the causal systems that is the traditional paradigm to describe systems in science. In fact, it can be proved [Winde68] that the class of causal systems is precisely the class of non-anticipatory systems.

2 Anticipatory Systems

We will now give a very brief description of a class of anticipatory systems according to Robert Rosen. For an introduction to the concept we refer to [Rosen74], while [Rosen85] gives a more elaborate treatment of the concept.

Let us suppose that we are interested in a system S which we call the *object system*. We will assume that S is an ordinary, non-anticipatory, dynamical system. With S we will associate another dynamical system M, which, in some sense, is a *model* of S. However, we require that if the trajectories of S are parameterized by real time, then the corresponding trajectories of M are parameterized by a time variable which goes faster than real time, i.e. the behaviour of M will *predict* the behaviour of S.

Furthermore, the systems M and S shall be coupled in order to allow them to interact in specific ways. The system M will use a set E of effectors, which allow it to operate either on S itself, or on the environmental inputs to S.

If the entire system is put in a single box (see Figure 1), it will appear to us to be an adaptive system in which predicted future behaviours can be used to determine the present changes of state.

Formulated in this way anticipatory systems seems to be the fundamental solution to the planning problem. However, there is a catch; In general we

Figure 1: The Components of an Anticipatory System

can't construct a *perfect* predictive model of the behaviour of a given object system S. There are two main reasons for this; (1) the difficulty to cover all correspondances between the possible states of a complex object system and the states of the model M, (2) the system may interact with a dynamic environment influenced and changed by other systems outside the range of the object system. It will be possible to construct a perfect predictive model only when we have a moderately complex object system interacting with a *constant* (or periodic) environment.

In spite of this, the concept of anticipatory systems forms a very interesting framework for planning as the self-adaption compensates for the use of imperfect models of various degrees. In addition it defines a conceptual framework that supports other desirable properties of an autonomous agent interacting with a dynamic environment.

2.1 Anticipatory Conditions

When an autonomous agent is carrying out a task in an environment its prescence is, in most cases, an essential part of that environment and should be encoded in a way that enables the agent to reason about itself as an object in the environment in terms of the encoded environment. We will call the natural system comprising the environment S_1 and the system comprising the agent S_2. The shared observable properties of the environment and the agent will be encoded by the same formal notation (variables, formulas, rules, etc.) from both systems into the model M (see Figure 2). The properties of the agent that are not part of the environment, such as the agent's reasoning capability, will be encoded in the model X. S_1 and S_2 are *open* systems, i.e. they may be subject to unpredictable influence from other (unknown) systems interacting with them.

Figure 2: The encodings of the environment (S_1) and the agent (S_2).

The following points will express the formal conditions that the components of this scheme must conform to in order to be an anticipatory system in the sense of [Rosen85]:

1. Those aspects of S_2 which comprise or embody the model M of S_1 must not exhaust all of S_2. That is there must be qualities of S_2 which are not related to the encoding E_2 displayed in the figure above. In terms of the diagram, we may say more formally that $E_2^{-1}E_1(S_1) \neq S_2$.

2. The encoding of S_2 is of the form

   $$S_2 \rightarrow M \times X$$

   where X is an encoding of observables of S_2 which are unlinked to the observables encoded in M, i.e. there is an orthoganality between the subsystem of S_2 encoded in M and the subsystem of S_2 encoded in X.

3. The state of the model M will modify the properties of other observables of S_2. There will be a growing discrepancy between the behaviour of X and $M \times X$ representing precisely the effect of the model M on S_2. This change of behaviour as seen in X can be seen as an *pre-adaption* of X.

4. There must be a discrepancy between the interaction between S_1 and S_2 which actually occurs, and the interaction which would have occured had the model M not been present.

5. The model of S_1 which is encoded into M must be a predictive model, i.e. the encoding E_1 should be a temporal encoding of the dynamics in S_1 such that the trajectories in S_1 as encoded in M will be traversable at a faster rate in M.

3 Example: The Rabbit Chase

In this section we will use this formal definition as a basis for a discussion of a very simple but hopefully illustrative example of an anticipatory approach to a dynamic planning problem.

Scenario: A dog is chasing a rabbit with the intention to catch it. In any instant each of them has a certain location, velocity and direction. We will assume that the rabbit, for some reason, will not run straight away from the dog all the time.

Objectives: We will use this scenario to discuss the essential characteristics of an anticipatory planning system.

3.1 Identification of the System Parts

Using the notation outlined above (see Figure 2) we will identify the system S_1 as the rabbit in its environment and S_2 as the dog. The model M is the encoding of the rabbit and the dog as running objects in the environment. The encodings into M of the rabbit and the dog as running objects will be denoted E_1 and E_2, respectively. The dog will have a planning capability for rabbit chasing encoded in X.

The observables of S_1, i.e. the rabbit running in the environment, will be encoded by E_1 as: x_r, y_r (the observed location of the rabbit), v_r (the observed velocity of the rabbit), ϕ_r (the observed direction of the rabbit). The corresponding encoding of S_2, i.e. the dog as a running object, will be encoded by E_2 as: x_d, y_d, v_d, and ϕ_d. The planning process of the dog is encoded by E_x into X.

3.2 The Predictive Model M

If the functions $x_r(t), v_r(t)$ and $\phi_r(t)$ are explicitly known (i.e. are a part of the model M) the dog may precisely determine the location of the rabbit at any future time. As it also is assumed to have a model of itself as a running object it should be able to modify its own velocity and direction in such a way that the rabbit is intercepted at a convenient point.

If the functions $x_r(t), v_r(t)$ and $\phi_r(t)$ are *not* explicitly known, as they wouldn't be in a normal chasing situation, the dog may use some predictive approximation for the functions. For example, it would be straightforward to assume that the rabbit will run tangential with the same speed at the near future of the last observation as expressed by:

$$\left\{ \begin{array}{lcl} x_r(t+h) & = & x_r(t) + h \times v_r(t) \times \cos\phi_r(t) \\ y_r(t+h) & = & y_r(t) + h \times v_r(t) \times \sin\phi_r(t) \end{array} \right. \qquad (3.1)$$

The corresponding encodings for the dog running in a straight line will be:

$$\begin{cases} x_d(t+h) & = & x_d(t) + h \times v_d(t) \times \cos\phi_d(t) \\ y_d(t+h) & = & y_d(t) + h \times v_d(t) \times \sin\phi_d(t) \end{cases} \tag{3.2}$$

A consequence of using an approximation of the rabbit's behaviour is that the dog must monitor the rabbit continuosly or at certain intervals in order to detect when the velocity or direction of the rabbit changes.

Of course, it is possible to express the relations between the observables in other ways.

3.3 The Planning Process X

The planning process X of the dog might have a number of models (M_i) of running rabbits to choose among. It might also have an ability to reason about modifications of the models to better fit the current or future situations.

An essential property of the planning process, seen as a closed subsystem, is that it should express how the dog would behave if there were no model of the running rabbit present. The straightforward strategy in this case would be to direct the dog to the current location of the rabbit.

3.4 Validation of the Anticipatory Conditions

We are now ready to validate this setup in terms of the anticipatory conditions:

1. It is trivially clear that $E_2^{-1}E_1(S_1) \neq S_2$, i.e. due to the planning capability the dog is more than a running object.

2. The aspects of the dog as a running object (S_2) is a subsystem that is unlinked to the subsystem comprising the dog's planning process (X).

3. The fact that the dog uses a model of the rabbit and itself as running objects will affect its behaviour. It will be able to adjust its chasing direction to an intercepting course. This change of the behaviour of the dog can be seen as a pre-adaption of the planning process of the dog. The encoding of the planning process alone defines a closed subsystem of the dog, whose properties describe how the dog would behave if there were no model M.

4. The fact that the dog will behave differently when using the model will cause a discrepancy in the interaction between itself and the rabbit compared to the interaction that would take place otherwise. Had the

model not been present the dog would just run straight towards the current position of the rabbit.

5. The model M is a predictive model of the positions of the dog and the rabbit due to the equations (3.1) and (3.2). The planning process of the dog may execute the equations to get an (approximative) prediction of the positions at any future time.

The conditions may be summed up as: An anticipatory system of a dog chasing a rabbit is one which contains a model of itself and a running rabbit in its environment with which it interacts. This model is a predictive model; its present states provide information about future locations of the rabbit and the dog. Further, the future states of the predictive model causes a change of state in the planning process of the dog; the planning process (a) causes a difference in the interaction between the dog and the rabbit, and (b) is unlinked to the model of the rabbit as a running object. In general we can regard the change of state in the dog as arising from the predictive model as a *pre-adaption* of the dog relative to its interaction with the rabbit.

4 Anticipatory Planning

We will now discuss the abilities for an autonomous planning agent, that were sketched in section 1, in the light of anticipatory planning.

4.1 Action-Orientation

A serious shortcoming of many planning systems working in real-time environments is their passive attitude to the environment. They do not take action until something happens. An anticipatory planning system, on the other hand, may use the predictive model not only to forsee a critical event but to gain information about the consequences. With this information available it may pre-adapt to the upcoming situation either by updating its model(s) M of the environment, or by an active change of behaviour in the reasoning part X. The latter may be in the form of offensive actions to prevent the event to take place, or it may be in the form of defensive actions in order to protect the agent or other objects from the consequences of the event.

4.2 The Analogous Model

The essence of a model is a relation between a natural system and some suitable formal system. The modelling relation itself is essentially a linkage between behaviours in the natural system and inferences drawn in the formal system. When we encode some aspects of two natural systems S_1 and S_2 in the same formal system M we will establish an analogy relation between them as each of them can be regarded as a model for the other (see Figure 3).

Figure 3: The encodings of S_1 and S_2 into M.

In this sense the relation $E_2^{-1}E_2$ can be regarded as an encoding of the states of S_1 into the the states of S_2. Thus, the relation of analogy gives rise to an essential feature of anticipatory systems; namely, that such a system possess a model of another system.

From a computational point of view the encoding relation will link corresponding observables and properties in both systems to another. The environmental part of the agent will in this respect be referenced by the same formal notation as the environment.

4.3 Levels within the Agent

The part of the agent that is not encoded in the environmental model comprises those aspects of the agent that is unrelated to the environment. We will in this respect have a clean separation between the agent's model of itself as a part of the environment and the other aspects of it, i.e. the reasoning part. The reasoning part of the agent may thus use the model as an *object* for its reasoning. The model may be compared to the environment and it may be subject to change by the reasoning part. When the model is used for prediction it is under the control of the reasoning part. The reasoning part "knows" that the model *is* a model and *what* it is a model of, while the model itself has no such knowledge.

4.4 Introspection and Reflection

In [Rosen74] the predictive model M is supposed to be a *realization* of the natural system it models, i.e. it is not an description of the system. In this case a time regression problem arises for perfect planning apart from the difficulty to construct such a model. To remove this difficulty it is suggested [Pattee77] that the model shall be description-based, where the description describes not only the observables and their relations in the natural system, but also their interpretation. However, this leads to self-referential problems

[Lofgren]. To avoid this altogether it is necessary to use a meta-language to describe the interpretation of the model. The use of a meta-level in the reasoning part of the agent and a meta-language to describe the interpretion of the model connects to the problem of introspection.

To address the problem of reflection one must go further, perhaps by constructing a model M' of the agent seen as the system $M+E+X$ in analogy with Figure 1 in section 2. We will then need a meta-reasoner X' that uses the former setup of the agent as an object to reason about. It will have access to the internal structures, operations and behavioural potential of X in order to realize at least some aspects of reflection.

4.5 Learning

In an anticipatory planner learning results in the modification of existing models and the generation of new ones. This is true both for learning at the environmental level and at the reasoning level. The difference is that in the latter case the agent should have access to the higher level model M'. We believe that *induction* is the natural basis for learning for a rational agent designed to interact with a real-world environment.

5 Summary

We have found the concept of anticipatory systems inspiring in our effort to find a framework for reactive planning in real-world environments. Most of the desirable abilities for an autonomous agent operating in such an environment will be properly addressed within this conceptual frame. The predictive model can be seen as the central part of an anticipatory planning system. It describes the environment and the properties of the agent that are linked to the environment. The reasoning part comprises a meta-level in relation to the predictive model and as such it controls the agent.

The approach has openings in several directions. The interaction with an environment seen as an open system makes the approach suitable for real-world applications and inductive learning. The anticipatory scheme seen as the aggregate $M+E+X$ might be modelled at higher levels opening up the prospects of reflection. However, the use of description-based models activates the complementarity between description and interpretation and, in the further perspective, the question of semantics.

References

[Agre-Chapman87] P. Agre and D. Chapman, *Pengi: An Implementation of a Theory of Activity*, In Proc. AAAI-87, pp. 268–272, 1987.

[Georgeff-Lansky87] M. Georgeff and A. Lansky, *Reactive Reasoning and Planning*, In Proc. AAAI-87, pp. 677–682, 1987.

[Kaelbling87] L. Kaelbling, *An Architecture for Intelligent Reactive Systems*, In Proc. 1986 Workshop on Reasoning about Actions and Plans, pp. 395–410, 1987.

[Lofgren] L. Löfgren, *The Partiality of Self-Reference*, (to be filled in later).

[Pattee77] H. Pattee, *Dynamic and Linguistic Modes of Complex Systems* International Journal of General Systems, Vol. 3, pp. 259-266, 1977.

[Rosen74] R. Rosen, *Planning, Management, Policies and Strategies - Four Fuzzy Concepts*, International Journal of General Systems, Vol. 1, pp. 245–252, 1974.

[Rosen85] R. Rosen, *Anticipatory Systems - Philosophical, Mathematical and Methodological Foundations*, Pergamon Press, 1985, ISBN 0-08-031158-X.

[Winde68] T. G. Windeknecht, *Mathematical Systems Theory: Causality*, Mathematical Systems Theory, Vol. 1, No. 4, 1967.