Opportunities: A unifying framework for planning and execution

Louise Pryor
School of Computer Science
The University of Birmingham
Edgbaston
Birmingham B15 2TT
UK
l.m.pryor@cs.bham.ac.uk

Gregg Collins
The Institute for the Learning Sciences
Northwestern University
1890 Maple Avenue
Evanston IL 60201
USA
collins@ils.nwu.edu

Abstract

A successful agent in the real world must both plan ahead and react to the unexpected. Ideally, both processes should be carried out in a common framework. In this paper we describe such a framework based on the analysis of opportunities. We argue that planning in advance can be viewed as a matter of anticipating opportunities, while responding to the unexpected should be seen as reacting to opportunities when they arise. We present an opportunistic planning agent, PARETO, that operates in a simulated robot delivery world, and implements our approach.

1. Introduction

The real world is regular enough to make advance planning worthwhile, yet unpredictable enough to make planning to the last detail impossible. An autonomous agent must therefore strike a balance between planning ahead and reacting to changes. For example, a robot on a strange planet must determine which areas to explore, where to take soil samples, what routes to take, and so on. The right choices will depend on details concerning the terrain encountered, the atmospheric conditions, and the results of tests performed on earlier samples—factors that cannot, in general, be predicted in sufficient detail to allow firm decisions to be made in advance. On the other hand, undirected wandering makes little sense: enough will be known in advance to make some decisions that will make a productive mission more likely. The distinction is a matter of available information: some of the information that would be required to construct an optimal plan will not be available before plan execution begins.

Since relevant information may become available at any point, even while a plan is being executed, an agent must be prepared to alter its plans to reflect new information. Ideally, this *replanning* process will resemble the process of planning in advance as much as possible, so that the same knowledge and processes can be applied. In other words, it is desirable that planning and replanning can be carried out in a common framework. In this paper we propose such a framework, based on the notion of responding to opportunities. We discuss the implementation of this approach in PARETO,[1] a system that notices and responds to opportunities as it pursues its goals (Pryor 1994).

1.1 Planning in an unpredictable world

Traditional AI planning systems (Fikes and Nilsson 1971; Sacerdoti 1977; Chapman 1987), known as *classical planners* (Wilkins 1988), have effectively decoupled plan construction and plan execution, operating on the assumption that all needed information will be freely available in advance. More specifically, these systems rely on three assumptions about the worlds in which they operate:

- *Simplicity:* it is possible to know everything about the world that might affect the agent's actions.
- *Stasis:* there will be no changes in the world except those caused by the agent's actions.
- *Certainty:* the agent's actions have deterministic results.

Were these assumptions valid, the world would hold no surprises for the agent; hence plans could be specified in exact detail with no fear that an unexpected outcome would force subsequent rethinking. Of course, there are few natural worlds in which the classical assumptions hold up. To take a simple example, consider an everyday human activity like preparing breakfast. For most people, this occurs day after day in the same environment, at the same time, using the same ingredients, and so on. It is an event regular and predictable enough to be described by a *script* (Schank and Abelson 1977). However, the apparent regularity of the breakfast world is only an artifact of our loftily abstract point of view; down at the level of detail at which we must treat the domain in order to execute a plan successfully, we find a world of wild and capricious unpredictability. The plan for making breakfast may involve any number of actions such as grasping, lifting, and pouring containers of milk, cereal, coffee, and so on; dishes and utensils must be manipulated; appliances must be operated; obstacles on the floor must be circumnavigated. To literally make a complete plan in advance, it would be necessary to know the exact position, orientation, and weight of every relevant object in the kitchen. It would be necessary to know, for example, the precise angle at which the box of corn flakes should be tilted to achieve optimal flow into the bowl, the proper position and altitude of the milk container during pouring to minimize splashing caused by the flakes, and so on.

Clearly, this is completely unrealistic. Even assuming the theoretical possibility of gathering such information in advance, the cost of acquiring, storing, and processing such a

[1] Planning and Acting in Realistic Environments by Thinking about Opportunities. Vilfredo Pareto (1848-1923) was an Italian economist, sociologist, and philosopher best known for the notion of *Pareto optimality* and the *Pareto distribution*, neither of which is used in this work.

quantity of information about the world in general is prohibitive. Furthermore, given inaccuracies in sensors and the interference of other agents, much relevant information is likely to be unavailable in principle. In other words, the breakfast world, like most natural environments, displays the following characteristics:

- *Complexity:* it is impossible to know everything.
- *Dynamism:* changes occur as a result of the actions of other agents or of natural phenomena.
- *Uncertainty:* the agent cannot be sure what the results of its actions will be.

Confronted with an environment that displays these characteristics, the classical planning paradigm breaks down. There are three key reasons for this. First, inaccurate information may cause wrong decisions to be made during the construction of a plan. For example, believing that there are clean bowls in the cupboard, you might construct a plan that entails opening the cupboard door. If you are wrong about the bowls, this step is unnecessary. The possibility of a faulty plan implies that agents must monitor their plans during execution and be prepared to recover from failures.

Second, information needed to make some decisions may not be available at the time a plan is chosen. For example, you cannot accurately predict the movements of your roommate, which means that there is potential for interference between your plan and your roommate's actions. In order to make an optimal decision on what path to take across the kitchen, you must know where your roommate will be, or at least know that she is out of the way. In general, decisions that fall into this category should be deferred until enough information is available; since the information will in many cases not be available until after execution of the plan has begun, the agent must prepared to interleave plan construction and plan execution.

Third, decisions may arise that have not been foreseen in the planning process. For example, the telephone might ring while you are pouring a glass of orange juice, forcing you to decide whether to continue pouring or stop and answer the phone. The agent must recognize circumstances under which the need to make an unforeseen decision arises, and must, if necessary, be able to acquire the information needed to make those decisions. The agent must be able to change its plans during their execution to reflect unforeseen situations.

In sum, the inevitability of the unexpected means that plans made in advance will require modification during execution. Expending effort on the construction of elaborate and detailed plans is therefore often unproductive. A more effective approach in the face of unpredictability is to expend some effort on choosing simple plans, and to expend more effort on adapting those plans as unforeseen circumstances are encountered. This is the approach followed in PARETO.

1.2 Plan execution

The emphasis in the design of PARETO is on recognizing the need for and making unforeseen decisions during plan execution. As an example of the kind of reasoning PARETO is meant to perform, suppose you happen to see a sharp knife as you are looking for a pair of scissors with which to cut string. In such a situation, unless there were some clear reason not to do so, you might well use the knife to achieve your goal and abandon the plan to find scissors. PARETO recognizes and takes advantage of such opportunities.

Instead of reasoning in detail about the interaction between plans for its various goals, as a classical planner would, PARETO constructs separate plans for each of its goals and does not expand effort attempting to anticipate potential interactions in advance. In addition, instead of expending a great deal of effort to gather all available information at planning time, PARETO depends on possibly faulty assumptions about the situation in which it will execute its plans. Obviously, the failure of these assumptions can cause PARETO's plans to fail. PARETO is designed to react quickly and flexibly to unexpected circumstances, rather than to minimize the possibility of an uncertainty arising.

2. Planning and opportunities

The cost of achieving a goal and the benefits of doing so can vary wildly over the goal's lifetime. For example, consider a goal to buy gas for your car. During a rush to make an important meeting, the cost of stopping for gas would be very high, since it is likely to make you late for the meeting. The benefit—essentially the reduced probability of running out of gas—is also high, but is likely outweighed by the cost. During the meeting, the cost of pursuing the goal is still higher, since walking out of the meeting and driving off to buy gas would be a most undesirable course of action; the benefit does not change. On the drive home after work, the cost of buying gas is relatively low—assuming you have no urgent plans—while the benefit increases as the likelihood of running out of gas becomes progressively more acute. At home, the cost of buying gas is again high, as it is now inconvenient to make a special journey, while the benefit stays the same. Cost and benefit continue to vary over time, depending on the exact situation in which you find yourself, until you decide to achieve the goal.

An effective planner must thus not only find workable plans to achieve its goal, but must also, insofar as possible, maximize the benefit and minimize the cost of doing so. In short, it must wait until there is a good *opportunity* for achieving a given goal. Planning can be seen as the process of predicting when opportunities will arise and what form they will take, and deciding in advance to take advantage of them. Although it is impractical to perform detailed predictions for all possible circumstances, it is often possible to perform simplified predictions. For example, it is routine to plan when to refuel your car based on your knowledge of the amount remaining in the tank, the locations of gas stations, and your travel plans over the next day or so.

2.1 Adapting the current plan

A plan that is designed to achieve a particular goal should be revised when a predicted opportunity does not in fact arise, or when a better opportunity comes along that was not considered during the planning process. Returning to an earlier example, if you were to receive a call on your car phone informing you that your meeting had been postponed, you would be presented with an unexpected opportunity to buy gas. In general an agent should notice such unexpected opportunities and consider adopting new plans to take advantage of them. As far as possible, the agent should respond to the unexpected opportunity in the same way as it would respond to the opportunity had it been predicted in advance.

The paradigm of noticing and responding to opportunities thus provides a unifying framework within which to approach both the issue of planning and the issue of plan revision during execution. In this approach, an agent's response to an opportunity should be independent of whether it has been predicted or not. *Planning in advance is based on the prediction of future opportunities, while plan revision is based on the recognition of current opportunities.*

2.2 Switching between plans

PARETO's approach means that it is pursuing a number of independent plans at any given time. The management of these diverse plans is thus a critical issue in PARETO's design. To manage them successfully, PARETO must distinguish those plans that are being actively pursued from those that are not. In fact, it can generally be assumed that only a handful of the agent's current plans will be pursued actively at any given time. For example, consider a plan to follow a recipe that says "soak the beans overnight." Clearly, pursuing this plan actively once the beans are put in to soak—for example, by sitting and watching them until morning—would be a tremendously inefficient plan. Instead, the agent should *suspend* the execution of this plan, and turn its attention to the pursuit of other goals. At any one time, most of an agent's goals are suspended. In order to manipulate plans in this way, PARETO must incorporate general mechanisms for deciding when it should change from following the plan for one goal to following the plan for another.

In general, a good time to attend to a particular goal is when there is an *opportunity* for that goal, whether predicted or unpredicted. For example, you should return to the soaking beans when you are in a position to proceed with the next step in the recipe. The existence of such an opportunity—to perform the next step in the preparation of the beans—was predicted in the construction of the overall plan. However, an unpredicted opportunity may similarly trigger a change of attention from one goal to another. For example, suppose you have tried and failed to get hold of a friend on the telephone to make some arrangements with her; if you see her in the supermarket, you may well temporarily suspend your goal of doing your weekly shopping to pursue your goal of making the arrangements.

Thus, management of plans is handled naturally within PARETO's paradigm of responding to opportunities. Decisions about changing plans, or reactivating suspended plans, are based on the recognition of current opportunities, while the plans themselves are based on the prediction of future opportunities.

3. Pareto

In this section we describe PARETO, a working system that illustrates how the framework described in this paper is an effective means of combining the execution of plans with appropriate responses to unexpected situations.

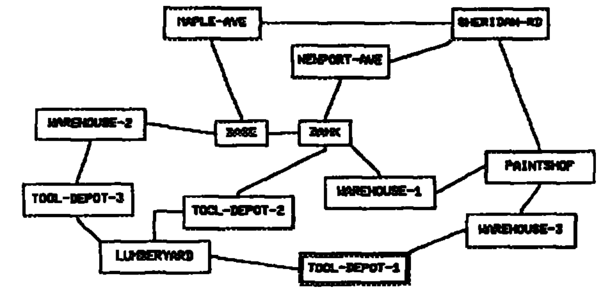

Figure 1 PARETO 's world

3.1 What PARETO does

PARETO operates a simulated robot delivery truck. The simulator was built using TRUCKWORLD (Firby and Hanks 1987; Hanks et al. 1993). In a TRUCKWORLD world, a robot delivery truck travels between locations on a network of roads, encountering and manipulating various objects as it goes. In PARETO's world (described in detail in Pryor, 1994), there are several building sites whose workers use the truck to run delivery errands such as "fetch a hammer," or "fetch something to carry my tools in."

PARETO's world consists of a number of locations linked by roads along which the truck can travel (see figure 1). Three of the locations are building sites, one is the truck's base where it can usually find fuel, and the other locations contain objects that the truck uses to fulfill its delivery goals. Most of these objects are used regularly by the construction workers whose errands the truck runs: hammers, saws, ladders, paint, and so on. There are over 30 different types of object in PARETO's world, of which 20 are used for deliveries. At any moment, there are typically well over 100 different objects at the various locations around the world.

PARETO's world is *unpredictable*: The truck has limited perception, and can sense only those objects that are at its current location, meaning that from PARETO's perspective the world is *complex*. It is also *dynamic*, since objects may spontaneously change location, appear, or disappear. Finally, the results of the truck's actions are *uncertain*: it may drop objects that it is trying to grasp, and neither the time taken to travel between locations nor the amount of fuel used can be predicted.

PARETO receives delivery orders at unpredictable intervals during its operation, with a typical run involving between seven and twelve separate deliveries. Plans that allow for every possible combination of goals would be far too complex; instead it uses a separate plan for each of its goals. These plans are *sketchy*—they do not specify in detail every action that should be performed. Instead, they specify the overall strategy that should be used in terms of a few simple steps, and PARETO decides how each step should be performed as it executes the plans.

PARETO's sketchy plans allow for some of the many contingencies that may arise (for example, they specify what to do when the object being grasped is dropped) but many circumstances cannot be foreseen in the plans. There are two types of situation in which PARETO must respond on the fly to circumstances that it encounters. First, circumstances may dictate that PARETO should switch its attention from one goal to another. For example, PARETO will not continue trying to make a delivery when the truck is running out of fuel, but will instead concentrate on trying to refill its fuel tank.[2] Second, an unforeseen opportunity may arise. This may entail either switching plans, or replacing a current plan with an alternative. For example, suppose the truck has two delivery goals, one for something to carry tools and one for something to cut twine. PARETO may decide to pursue the carry-tools goal first, and set off to the warehouse in which it expects to find a box. If the truck finds a knife or a pair of scissors on the way, however, PARETO will temporarily switch its attention and pick up the cutting tool. If PARETO subsequently encounters a bag that would be suitable to carry tools, it would abandon its plan to find a box and instead pick up and deliver the bag.

PARETO has an efficient mechanism for spotting unpredicted opportunities such as these, and treats them in exactly the same way as it does opportunities that have been predicted in its plans. The next sections explain PARETO's basic operation and characterization of plans and opportunities.

3.2 How PARETO works

PARETO is based on Firby's RAPs[3] plan execution system (Firby 1987; 1989), and is described in (Pryor 1994). When PARETO acquires a new goal, it looks in its library of *RAPs* (sketchy plans) for one that will achieve the goal. The steps in a RAP specify subgoals that the system must achieve in order to execute the plan successfully. PARETO recursively expands sketchy plans by choosing aplna for each subgoal. Eventually, a subgoal will be achievable by performing a simple action, and no further expansion is required.

When PARETO receives a goal, its first action is to place a task aimed at achieving that goal on the task agenda. The task that is placed on the task agenda consists of the goal and the RAP that has been chosen to achieve it. PARETO's execution cycle is summarized in figure 2, and consists of the following steps:

- Choosing a task from the agenda. The task that will be the most productive in furthering PARETO's goals should be chosen. PARETO's task selector operates by looking for opportunities to further its various goals. The ability to recognize unforeseen opportunities is a significant change from the task selector used by Firby's RAPS system.
- Processing of the chosen task to fill in the details of the incompletely specified plan that describes the task. A RAP specifies all the different plans that might be used to achieve its goal. For example, a goal to find fuel might be achieved by going to the location of a fuel drum that PARETO knows about, by going to the base which is a source location for fuel, or by wandering around the world until a fuel drum is found. Each such plan is a method of achieving the goal. PARETO chooses one of the methods of the task that is being processed, based on the state of the world at the time that the processing takes place.
- Adding new tasks to the task agenda for each of new goals created during the previous processing step.
- Reprocessing the original task when each of its subtasks has achieved its goal. The successful execution of the subtasks is not enough to guarantee that the original task will itself have achieved its goal, since PARETO's world is dynamic and some time may pass between the execution of a task's subtasks and the repeat processing of the original task. For example, the truck might succeed in going to the location of a known fuel drum, but the drum might have meanwhile disappeared, or the fuel in it been used by another agent. If the task has succeeded, it is removed from the task agenda, else it is processed as described above.

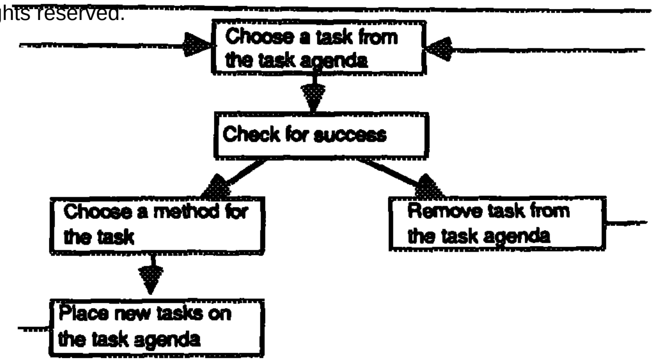

Figure 2 PARETO's execution cycle

To illustrate the execution cycle in action, consider what happens when PARETO receives the two goals in the example above, to deliver something to carry tools and something to cut twine. As each goal is received, the deliver-object RAP is chosen to achieve it and the relevant task is placed on the task agenda. The deliver-object RAP has four steps: PARETO must find a suitable object, load it, travel to the correct location, and unload the object. After PARETO has processed the deliver-object task for the carry-tools goal, there are thus five tasks on the agenda for that goal: deliver-object, find-object, load-payload-object, truck-travel-to, and unload-at. Of these, the deliver-object task is waiting for the other four to complete, and the load-payload-object, truck-travel-to, and unload-at tasks are waiting for their predecessors to complete. If PARETO has no other goals, the next task to be chosen will be the find-object task, which will in turn be expanded and its subgoals placed on the task agenda. If all goes according to plan, all the subtasks will be processed in turn and removed from the task agenda until the unload-at task has been achieved. Finally, PARETO will again process the deliver-object task, find that it has succeeded, and remove it from the task agenda.

[2] As well as delivery goals, PARETO has preservation goals (Schank and Abelson 1977) to ensure that the truck does not run out of fuel and to keep track of its surroundings.

[3] Reactive Action Packages

```
--> New order ORDER-5: something to CARRY TOOLS for ILS
--> New top level goal: DELIVER-OBJECT ::[6:6]
--> New order ORDER-4: something to CUT TWINE for TECH
--> New top level goal: DELIVER-OBJECT ::[7:7]
... Processing task:
      <DELIVER-OBJECT CARRY-TOOLS ILS ORDER-5>::[6:6]
... Processing task: <FIND-OBJECT CARRY-TOOLS ... >::[6:5]
    ... the truck sets off to find a box ...
+++ Potential opportunity for
      <DELIVER-OBJECT CARRY-TOOLS ILS ORDER-9>
      from ITEM-20 (BAG)
+++ Potential opportunity for <FIND-OBJECT CARRY-TOOLS ... >
      from ITEM-20 (BAG)
+++ Potential opportunity for
      <LOAD-PAYLOAD-OBJECT ?OBJECT CARRY-TOOLS>
      from ITEM-20 (BAG)
+++ Potential opportunity for
      <DELIVER-OBJECT CUT-TWINE TECH ORDER-8>
      from ITEM-13 (SCISSORS)
+++  Has already succeeded <FIND-OBJECT CARRY-TOOLS.. >
***  Taking unexpected opportunity:  [6:5]
      <FIND-OBJECT CARRY-TOOLS..>
...  Processing task: <FIND-OBJECT CARRY-TOOLS ... >::[6:5]
...  Processing task:
      <LOAD-PAYLOAD-OBJECT ITEM-20 CARRY-TOOLS>::[6:8]
    ... the truck loads the bag ...
... Task succeeded:
      <LOAD-PAYLOAD-OBJECT ITEM-20 CARRY-TOOLS>::[6:8]
*** Taking expected opportunity:  [7:7]
      <DELIVER-OBJECT CUT-TWINE TECH ORDER-8>
*** Changing goals from 6 to 7
... Processing task:
      <DELIVER-OBJECT CUT-TWINE TECH ORDER-8>::[7:7]
... Processing task: <FIND-OBJECT CUT-TWINE ... >::[7:5]
... Task succeeded:
      <FIND-OBJECT CUT-TWINE => WAREHOUSE-2 ITEM-13>::[7:5]
... Processing task:
      <LOAD-PAYLOAD-OBJECT ITEM-13 CUT-TWINE>::[7:8]
    ... the truck loads the scissors ...
... Task succeeded:
      <LOAD-PAYLOAD-OBJECT ITEM-13 CUT-TWINE>::[7:8]
*** Changing goals from 7 to 6
... Processing task: <TRUCK-TRAVEL-TO MAPLE-AVE>::[6:5]
    ... the truck goes off to deliver the bag...
```

Figure 3 PARETO takes advantage of two opportunities

3.3 Opportunities in PARETO

PARETO chooses which task to process next by considering those tasks that are associated with *opportunities*. When PARETO spots an opportunity, it does not consider whether or not the opportune task is the next step in the plan for the current goal. Thus, it may in effect ignore the existing plan for carrying out a task when an opportunity arises; if it succeeds in taking advantage of the opportunity, the task and its the previous plan are simply removed from the agenda. Furthermore, PARETO is not constrained by any notion of the "current" task in spotting opportunities—thus, it can effectively switch plans whenever an opportunity arises.

In our example, PARETO changes its plan for its current goal by deciding to pick up the bag instead of continuing to look for a box. It also switches its attention to another goal by picking up the scissors. However, the latter switch is only temporary, as picking up the scissors does not achieve the task of finding an object that can carry tools, which task remains on the agenda. After picking up the scissors PARETO returns to the carry-tools goal and goes off to deliver the bag. Figure 3 shows PARETO's output as it recognizes and takes advantage of these opportunities.

PARETO thus characterizes opportunities as tasks on its task agenda that are easy to achieve. It has an efficient mechanism for recognizing opportunities that uses a filtering process based on the functional characteristics of objects in the world (Pryor 1994). There are two ways in which one of PARETO's tasks may be easily achievable, corresponding to the two types of opportunities we discussed earlier: either it has already succeeded, or it is ready for processing (not waiting for any others to be completed). Tasks that have already succeeded may indicate the presence of an unexpected opportunity, while a task that is ready for processing may indicate the presence of a predicted opportunity—one that has been foreseen in a plan. In our example, the finding of a bag is unexpected: PARETO had planned to find a box.

PARETO thus characterizes opportunities in terms of easily achievable tasks on its task agenda. Expected opportunities are associated with tasks that represent the next step in the plan for one of its goals, and unexpected opportunities are associated with tasks that have already succeeded.

4. Related work

An agent should revise its plans when it encounters an unexpected opportunity, but cannot afford to analyze every situation exhaustively. The centrality of opportunities to the execution of plans in an unpredictable world has received little attention by other researchers in the field.

Most current research on the problem of recognizing the need to make unforeseen decisions fails to address the issue of opportunism explicitly (Bresina and Drummond 1990; Ferguson 1992; Lyons and Hendriks 1992; McDermott 1992). In general, this work relies heavily on projecting the agent's current plans to determine when replanning would be desirable. As projection may involve arbitrarily complex reasoning, this approach fails to address the problem of recognizing the need to make decisions quickly enough that the agent can respond appropriately in a dynamic world.

The earliest work on opportunity recognition, by Hayes-Roth and Hayes-Roth (1979), looked at opportunism in plan construction, but did not consider plan execution.

Hammond and his colleagues (Hammond et al. 1993) present a method of opportunity recognition based on recognizing the features involved in a goal's achievement. This approach relies on having specified the plan for the goal in enough detail that the environmental elements involved are already known. It does not allow an agent to recognize opportunities for goals that it has not yet decided how to achieve, and does not allow the recognition of opportunities that require a different method of achievement from that in the current plan. For example, Hammond's approach would not allow an agent to recognize the opportunity discussed above, in which the presence of a bag allows an agent to abandon its plan to find a box. This limited view of opportunity recognition would prevent the use of opportunities as a framework for combining planning and responsiveness.

The impracticality of unlimited replanning has long been recognized as a serious problem in the design of intelligent agents. There are two aspects to the problem: the necessity of limiting the amount of reasoning that is performed, and the necessity of determining when this limited reasoning should occur. Traditional AI approaches have concentrated on limiting the amount of reasoning that is performed, by using either quantitative approximations (Hanks 1990) or qualitative techniques (Wellman 1988; 1990). *Anytime algorithms* (Dean and Boddy 1988; Boddy and Dean 1989) address the issues of designing reasoning algorithms that will produce an answer in limited time. None of this work addresses the issue of *when* this reasoning should be performed.

In Hayes-Roth's GUARDIAN system (Hayes-Roth 1990) *global control plans* are used to direct the agent's reasoning towards important goals. These control plans are changed by global control decisions, which appear to be triggered by the receipt of sensory information. However, Hayes-Roth gives no details of how these decisions are triggered or the process by which they are made. Presumably GUARDIAN's mechanism for making these decisions involves minimal reasoning, as they appear to occur rapidly when necessary, but there is no discussion of this aspect. GUARDIAN thus limits the amount of reasoning that need be done by focusing it on a subset of the agent's goals, but there is no clear answer to the question of when such reasoning should be performed.

Maes (1991) describes a network-based architecture with parameters that adjust the speed with which the agent reacts to changes in its environment. If the environment changes slowly, the agent can perform more reasoning before responding; in very unpredictable environments, the agent must react with little or no reasoning. The parameters change only in response to the unpredictability of the environment, and are not affected by the particular situation in which the agent finds itself. The balance between acting and reasoning changes on a global basis, and no attempt is made to direct the reasoning towards specific goals.

5. Conclusion

To be successful, an agent in the real world must both plan ahead and react to the unexpected. Ideally, both processes should be carried out in a common framework. In this paper we have described such a framework based on the analysis of opportunities, and a computer program, PARETO, that implements our approach.

Acknowledgements

Most of this work was carried out while the first author was at the Institute for the Learning Sciences. This work was supported in part by the AFOSR under grant number AFOSR-91-0341-DEF, and by DARPA, monitored by the AFOSR under contract F49620-88-C-0058. The Institute for the Learning Sciences was established in 1989 with the support of Andersen Consulting, part of The Arthur Andersen Worldwide Organization. The Institute receives additional support from Ameritech and North West Water, Institute Partners, and from IBM.

References

Boddy, M. and T. Dean. 1989. "Solving time-dependent planning problems." In Proceedings of the Eleventh International Joint Conference on Artificial Intelligence, Detroit, MI, AAAI.

Bresina, J. and M. Drummond. 1990. "Integrating Planning and Reaction: A preliminary report." In *Working notes of the Spring Symposium on Planning in Uncertain, Unpredictable, or Changing Environments, Stanford, CA*, AAAI.

Chapman, D. 1987. Planning for Conjunctive Goals. *Artificial Intelligence* 32 : 333-337.

Dean, T. and M. Boddy. 1988. "An analysis of time-dependent planning." In *Proceedings of the Seventh National Conference on Artificial Intelligence, St Paul, MN*, AAAI.

Ferguson, I. A. 1992. *TouringMachines: An architecture for dynamic, rational, mobile agents*. Computer Laboratory, University of Cambridge. Technical Report No. 273.

Fikes, R. E. and N. J. Nilsson. 1971. STRIPS: A new approach to the application of theorem proving to problem solving. *Artificial Intelligence* 2 : 189-208.

Firby, R. J. 1987. "An investigation into reactive planning in complex domains." In *Proceedings of the Sixth National Conference on Artificial Intelligence, Seattle, WA*, AAAI.

Firby, R. J. 1989. *Adaptive execution in complex dynamic worlds*. Department of Computer Science, Yale University. Technical Report YALEU/CSD/RR #672.

Firby, R. J. and S. Hanks. 1987. *The simulator manual*. Department of Computer Science, Yale University. Technical Report YALEU/CSD/RR #563.

Hammond, K., T. Converse, M. Marks, and C. Seifert. 1993. Opportunism and learning. *Machine Learning* 10 : 279-309.

Hanks, S., M. E. Pollack, and P. R. Cohen. 1993. "Benchmarks, testbeds, controlled experimentation, and the design of agent architectures." *AI Magazine*, Winter 1993.

Hanks, S. J. 1990. *Projecting plans for uncertain worlds*. Department of Computer Science, Yale University. Technical Report YALEU/CSD/RR #756.

Hayes-Roth, B. 1990. "Dynamic control planning in intelligent agents." In *Working notes of the Spring Symposium on Planning in Uncertain, Unpedictable or Changing Environments, Stanford University*, .

Hayes-Roth, B. and F. Hayes-Roth. 1979. A cognitive model of planning. *Cognitive Science* 3 (4): 275-310.

Lyons, D. M. and A. J. Hendriks. 1992. "A practical approach to integrating reaction and deliberation." In *Proceedings of the First International Conference on Artificial Intelligence Planning Systems, College Park, Maryland*, Morgan Kaufmann.

Maes, P. 1991. "Adaptive action selection." In *Proceedings of the Thirteenth Annual Conference of the Cognitive Science Society, Chicago, IL*, Lawrence Erlbaum Associates.

McDermott, D. 1992. *Transformational planning of reactive behavior*. Yale University, Department of Computer Science. YALEU/CSD/RR #941.

Pryor, L. 1994. "Opportunities and Planning in an Unpredictable World." PhD dissertation, in preparation, The Institute for the Learning Sciences, Northwestern University.

Sacerdoti, E. 1977. *A structure for plans and behavior*. New York: American Elsevier.

Schank, R. C. and R. P. Abelson. 1977. *Scripts, plans, goals and understanding*. Hillsdale, NJ: Lawrence Erlbaum Associates.

Wellman, M. P. 1988. "Qualitative Probabilistic Networks for Planning under Uncertainty." In *Uncertainty in Artificial Intelligence*, ed. J. F. Lemmer and L. N. Kanal. 2. Amsterdam: Elsevier.

Wellman, M. P. 1990. *Formulation of Tradeoffs in Planning under Uncertainty*. London: Pitman.

Wilkins, D. E. 1988. *Practical Planning: Extending the Classical AI Planning Paradigm*. San Mateo, CA: Morgan Kaufmann.