

Synthesizing Plant Controllers Using Real-time Goals*

M. Barbeau, F. Kabanza, and R. St-Denis

Departement de mathematiques et d'informatique

Universite de Sherbrooke

Sherbrooke, Quebec J1K 2R1 Canada

Email: {barbeau, kabanza, stdenis}@dmi.usherb.ca

Abstract

This paper introduces a novel planning method for reactive agents. Our planning method handles, in a single framework, issues from AI, control theory, and concurrency that have so far been considered apart. These issues are mostly controllability, safety, bounded liveness, and real time. Our approach is founded on the supervisory control theory and on Metric Temporal Logic (MTL). The highlights of our method consist of a new technique for incrementally checking MTL goal formulas over sequences of states generated by actions and a new method for backtracking during search by taking into account uncontrollable actions.

1 Introduction

A *reactive plan* is a program that specifies the actions to be executed by an agent that continuously reacts to the occurrence of discrete events from its environment. The automatic synthesis of reactive plans is a key component in the design of autonomous systems. It is of interest in AI planning [Georgeff and Lansky, 1987; Drummond and Bresina, 1990; Dean *et al.*, 1993], control theory [Ramadge and Wonham, 1989; Makungu *et al.*, 1994], and concurrency theory [Emerson, 1990]. In these areas, the terms *control rules*, *supervisor*, and *reactive programs* are respectively used to denote reactive plans. However, the paradigms used to generate such plans vary widely. In AI planning, the focus is on representing and reasoning about actions in an efficient and natural way. This has led researchers in this field to concentrate on heuristic search techniques. Unfortunately, little research in AI planning has taken into account the concepts of controllable actions and liveness goals.

In the area of control theory, the emphasis has been on reasoning about the issues of controllability and observability of actions [Ramadge and Wonham, 1989]. A

This work was supported by the Canadian Government through its NSERC programs.

key characteristic of models of actions in this field is that they distinguish between controllable and uncontrollable actions. This facilitates reasoning about controllability issues.

In concurrency, modal temporal logic has been recognized as a useful tool for dealing with liveness properties. For instance, real-time modal temporal logic can express bounded-time response properties of the form "each occurrence of event e leads to satisfaction of condition c within 5 time units" [Alur and Henzinger, 1993]. One interesting development in this area is *model checking* [Emerson, 1990]. In this approach, verifying a program is viewed as evaluating the truth of a temporal formula on a temporal structure representing the program.

In this paper, we describe a method for generating reactive plans. Our method borrows concepts and concerns from all the three areas outlined above. From AI planning, we take the concern of reasoning about actions and using heuristic search to generate plans. From control theory, we borrow the concept of controllable action. As in concurrency theory, we are also concerned with the representation and reasoning about real time, safety, and liveness goals.

To be more specific, our planner takes as input a goal, a description of the primitive actions of a reactive agent, which are *controllable*, and those of its environment, which are *uncontrollable*. The planner returns a reactive plan for achieving the goal. Note that both the agent and its environment are referred to as the *process* or *plant* in control theory. We generate plans essentially by searching among sequences of states that represent executions of primitive actions, checking whether they satisfy the goal or not. Reactive plans are more or less immediately extracted from satisfactory sequences. We describe goals using Metric Temporal Logic (MTL) formulas [Alur and Henzinger, 1993].

Recent research has investigated connections between AI planning, control, and concurrency. This includes work on representation of AI control programs [Nilsson, 1994], a recent book by Dean and Wellman on planning and control [Dean and Wellman, 1991], work on the integration of planning and concurrency ideas [Gode-

froid and Kabanza, 1991; Kabanza, 1990], and work on the synthesis of controllers in temporal logic frameworks [Thistle and Wonham, 1986; Fusaoka *et al*, 1983]. We view our contribution to this research in three respects. First, our planner integrates concepts and concerns from the above three fields deeper than any existing comparable framework. Hence, besides its more general capabilities, our planner contributes to a better understanding of issues in the boundaries of these fields.

Second, we introduce a method for checking MTL goals over state sequences *on the fly* (i.e., incrementally). In contrast to other techniques for checking similar real-time formulas (e.g., [Alur and Henzinger, 1993; Alur *et al*, 1993]), our technique does not require an explicit storage of the state-transition graph over which formulas are checked. Rather, our planner generates and checks state sequences *on the fly*. This allows us to generate plans using a standard forward-chaining search. It also lets our planner generate plan control rules incrementally. In this way, our planner can be considered as having *anytime* capability. According to Dean and Boddy [Dean and Body, 1988], an *anytime* algorithm is one that can be stopped at different stages of its processing and yield an approximate, but useful, result.

Third, we introduce a search technique, called *control-directed backtracking*. This technique uses selective backtracking to states from which the planner can expand controllable paths. We use an example to demonstrate that this approach can allow the planner to prune the search space significantly. This search technique is reminiscent of, but quite different from, the mechanisms used in *dependency-directed backtracking* [Stallman and Sussman, 1978] and partial-order search using Mazurkiewicz's *traces* [Godefroid and Kabanza, 1991].

Although our planning approach applies to standard MTL and more general logics, due to space limitations, we will describe it only for a restricted class of MTL formulas called *bounded-time MTL formulas (BMTL)*. Even with this restriction, our planner handles problems that are beyond the scope of comparable AI planning or control-theoretic frameworks. For example, as was mentioned above, our planner treats various types of bounded-time goals.

The remainder of this paper is organized in three parts, beginning with the description of our model of actions and plans. This is followed by a brief overview of MTL and a description of our planning method.

2 Actions and Plans

We adopt the following model of plans from Ramadge and Wonham¹ [Ramadge and Wonham, 1989]. A process is modeled as a spontaneous generator of sequences of actions. The set of actions A of the process are partitioned into two disjoint subsets A_u and A_c . The ac-

¹As was mentioned in the introduction, we use the term "plan" as a synonym for supervisor.

tions in A_u are uncontrollable, while those in A_c are controllable. Hence, the evolution of the process can be controlled by prohibiting the occurrence of some controllable actions at certain points. A *control input* is a subset y of A satisfying the condition $A_u \cap y = \emptyset$. If $a \in y$, then a is permitted to occur. Let L denote the set of sequences of actions that can be generated by the process and $r \subseteq 2^A$ represent the set of all control inputs. A *plan* is a map $f : L \rightarrow r$ yielding, for each sequence $w \in L$ of generated actions, the control input $f(w)$ to be applied to the process at that point. Language L_f denotes the sequences of actions generated by the process under the supervision of f . Let ϵ be the empty string, $a \in A$, and $w \in A^*$. The language L_f is formally defined as follows:

- (1) $\epsilon \in L_f$ and L_f^*
- (2) $wa \in L_f$ iff $w \in L_f, a \in A \setminus f(w)$, and $wa \in L$.

In our framework, the *process* consists of a reactive agent and its environment, both performing primitive actions (see Figure 1). The actions performed by the agent are controllable, whereas those performed by the environment are uncontrollable. We model the interaction between the agent and its environment by interleaving their actions.



Figure 1: A model of control

Actions There are various models of actions in the fields of AI, concurrency, and control theory. Representations of actions in planning are mostly based on STRIPS-like models [Fikes and Nilsson, 1972]. Models of actions in concurrency and control theory are often based on state-transition machines such as automata or Petri nets. Nevertheless, all these models are more or less theoretically equivalent. For example, a state-transition model can be derived from a STRIPS-like formalism by generating the states reachable from one or several initial states. Hence, without loss of generality, we assume a STRIPS-like formalism in this paper.

The STRIPS model of actions includes two basic concepts, namely, world state and primitive action. A *world state* is defined as a set of propositions. Propositions model facts about states of the agent and its environment. Every primitive action is described by a *precondition* (a conjunction of propositions enabling the action), a *delete-list* (a list of propositions retracted by the execution of the action), and an *add-list* (a list of propositions asserted by the execution of the action). Note that, in

world states to compute the formulas to be attached to states. Intuitively, progressing an MTL formula through a world state consists in checking that the *present* requirement conveyed by the formula is satisfied in this state. If so, a formula representing the *future requirement* is returned, otherwise false is returned. Hence, given a current world state and an action, the formula representing the future requirement is paired with the successor world state obtained by applying the action to the current world state.

4.2 Formula Progression

The formula progression algorithm is described in Figure 2. The function *Prog* accepts as input an MTL formula l , a world state s , a real number t , and a function $7r(p, S)$, which returns *true* only if the proposition p holds in the world state s . It returns an MTL formula representing the progression of l through s . This algorithm is characterized by the following theorem.

```
function Prog(f,s,t,π)
1. case /
2.  $p$  ( $p$  a proposition):  $\pi(p, s)$ ;
3.  $\neg f_1$ :  $\neg \text{Prog}(f_1, s, t, \pi)$ ;
4.  $f_1 \wedge f_2$ :  $\text{Prog}(f_1, s, t, \pi) \wedge \text{Prog}(f_2, s, t, \pi)$ ;
5.  $f_1 \vee f_2$ :  $\text{Prog}(f_1, s, t, \pi) \vee \text{Prog}(f_2, s, t, \pi)$ ;
6.  $\Box_{=x} f_1$ : if  $x - t > 0$  then  $\Box_{=(x-t)} f_1$ 
7.           else  $\text{Prog}(f_1, s, t, \pi)$ ;
8.  $\Box_{\leq x} f_1$ : if  $x - t > 0$  then  $\text{Prog}(f_1, s, t, \pi) \wedge \Box_{\leq(x-t)} f_1$ 
9.           else  $\text{Prog}(f_1, s, t, \pi)$ ;
10.  $\Box_{\geq x} f_1$ : if  $x - t > 0$  then  $\Box_{\geq(x-t)} f_1$ 
11.            else  $\text{Prog}(f_1, s, t, \pi) \wedge \Box_{\geq 0} f_1$ ;
12.  $f_1 U_{=x} f_2$ :
13.   if  $x - t > 0$  then  $\text{Prog}(f_1, s, t, \pi) \wedge f_1 U_{=(x-t)} f_2$ 
14.   else  $\text{Prog}(f_2, s, t, \pi)$ ;
15.  $f_1 U_{\leq x} f_2$ :
16.   if  $x - t > 0$  then  $\text{Prog}(f_2, s, t, \pi) \vee$ 
17.      $(\text{Prog}(f_1, s, t, \pi) \wedge f_1 U_{\leq(x-t)} f_2)$ 
18.   else  $\text{Prog}(f_2, s, t, \pi)$ ;
end
```

Figure 2: Progression algorithm for BMTL formulas.

Theorem 4.1 Let $M = \langle S, \pi, \mathcal{T} \rangle$ be any MTL model and let S_i be the i th state in the sequence of states S . Then, we have for any MTL formula f , $(M, s_i) \models f$ if and only if $(M, s_{i+1}) \models \text{Prog}(f, s_{i+1}, \mathcal{T}(s_{i+1}) - \mathcal{T}(s_i), \pi)$.

The time difference $\mathcal{T}(s_{i+1}) - \mathcal{T}(s_i)$ corresponds to the duration of the action that leads to s_{i+1} from s_i . For the sake of simplicity, we assume that each action takes 1 time unit.

4.3 The Planning Algorithm

We can now detail the algorithm of our planner. The input consists of an initial world state x_0 , a set of primitive actions \mathcal{A} , the transition function δ , an MTL goal g , and the function π (see Figure 3). It outputs a plan (F, φ) . For the sake of clarity, the planning algorithm is divided into three parts: the prologue (Figure 3), search process (Figure 4), and epilogue (Figure 5).

function Planner(in $x_0, \delta, \mathcal{A}, g, \pi$, out (R, φ))

```
1.  $f_0 \leftarrow \text{Prog}(g, x_0, 0, \pi)$ ;
2. if  $f_0 = \text{false}$  then return the empty solution;
3.  $[x_0, f_0].\text{pct} \leftarrow \{\}$ ;  $[x_0, f_0].\text{pre} \leftarrow \{\}$ ;
4.  $[x_0, f_0].\text{status} \leftarrow \text{undefined}$ ;
5.  $[x_0, f_0].\text{controllable} \leftarrow \text{false}$ ;
6. CLOSED  $\leftarrow \{\}$ ; OPEN  $\leftarrow \{[x_0, f_0]\}$ ;
```

Figure 3: Prologue to the planning algorithm

```
7. repeat select  $[x, f]$  in OPEN;
8.    $[x, f].\text{controllable} \leftarrow \text{true}$ ;
9.   CLOSED  $\leftarrow$  CLOSED  $\cup \{[x, f]\}$ ; OPEN  $\leftarrow$  OPEN  $\setminus \{[x, f]\}$ ;
10.  for each  $a$  in  $\mathcal{A}$  such that  $\delta(a, x)$  is defined do
11.    /* Expand the state  $[x, f]$  */
12.     $x' \leftarrow \delta(a, x)$ ;  $f' \leftarrow \text{Prog}(f, x', 1, \pi)$ ;
13.    if  $f' = \text{false}$  or  $([x', f'] \in \text{CLOSED and}$ 
14.       $[x', f'].\text{status} = \text{sink})$ 
15.    then if not Controllable( $a$ )
16.      /*  $[x, f]$  becomes sink */
17.       $[x, f].\text{controllable} \leftarrow \text{false}$ ; Sink( $[x, f]$ );
18.      break;
19.    else
20.       $\xi(a, [x, f]) \leftarrow [x', f']$ ;
21.      if  $[x', f'] \notin \text{CLOSED and } [x', f'] \notin \text{OPEN then}$ 
22.        /* Initialization of a new state */
23.         $[x', f'].\text{pct} \leftarrow \{\}$ ;  $[x', f'].\text{pre} \leftarrow \{\}$ ;
24.         $[x, f].\text{status} \leftarrow \text{undefined}$ ;
25.         $[x', f'].\text{controllable} \leftarrow \text{false}$ ;
26.        OPEN  $\leftarrow$  OPEN  $\cup \{[x', f']\}$ ;
27.        if Controllable( $a$ ) then
28.           $[x', f'].\text{pct} \leftarrow [x', f'].\text{pct} \cup \{[a, [x, f]]\}$ ;
29.          else
30.             $[x', f'].\text{pre} \leftarrow [x', f'].\text{pre} \cup \{[x, f]\}$ ;
31.             $[x, f].\text{controllable} \leftarrow \text{false}$ ;
32.            if  $[x, f].\text{controllable} = \text{true then}$ 
33.              /* Check for a trivial satisfactory state */
34.              if not Has_Eventually( $[x, f]$ )
35.                then Satisfactory( $[x, f]$ );
36.            until OPEN =  $\{\}$ ;
37.            /* Check for non trivial satisfactory states */
38.            Check.Non.Trivial.Satisfactory.States();
```

Figure 4: Search process of the planning algorithm

Prologue The prologue consists in initializing a search graph that implements an intermediate automaton $M = \langle X \times F, \mathcal{A}, \xi, [x_0, f_0] \rangle$, where F is a set of bounded-time MTL formulas derived from g using the formula progression algorithm. The set of states $X \times F$ is partitioned into expanded and unexpanded states. Initially, the set of expanded states, CLOSED, is empty; the set of unexpanded states, OPEN, contains the initial state $[x_0, f_0]$, where $f_0 = \text{Prog}(g, x_0, 0, \pi)$. Note that $\text{Prog}(g, x_0, 0, \pi)$ returns a formula equivalent to g , since the time difference used is 0. The purpose is to check that the *present* requirement of g is satisfied in x_0 .

Overview of the Search Process The search process repeatedly selects a state $[x, f]$ from OPEN and proceeds as follows. For each action enabled in x , it

```

/* Compute the reactive plan (supervisor) */
34. if  $[x_0, f_0].status = satisfactory$ 
35. then  $R \leftarrow (CLOSED_{sat}, \mathcal{A}, \xi_{sat}, [x_0, f_0])$ ;
36. else return the empty solution;
/* Compute the feedback function  $\varphi$  */
37. for each  $[x, f]$  in  $CLOSED_{sat}$  do
38.   for each  $a$  in  $\mathcal{A}$  do
39.     if  $\delta(a, x)$  is defined and  $\xi_{sat}(a, [x, f])$  is defined
40.       then  $\varphi([x, f]) \leftarrow \varphi([x, f]) \cup \{a\}$ ;
end

```

Figure 5: Epilogue to the planning algorithm

computes a successor $[x', f']$, where $x' = \delta(a, x)$ and $f' = Prog(f, x', 1, \pi)$.

Epilogue The epilogue derives a plan $\langle R, \langle p \rangle$ from the intermediate automaton M as follows. First, the planner removes the states that are not *satisfactory* from M to obtain the automaton R (lines 34-35). Then $\langle p \rangle$ is obtained by including in $\langle p \rangle$ an action a for each transition (s, a, s') of R (lines 37-40). Hence, a crucial step in our planning process is to determine *satisfactory* states.

Example 4.1 Let us consider the maze example pictured in Figure 6 (taken from [Ramadge and Wonham, 1989]). The maze has five rooms. There are a cat and a mouse moving from one room to another through doorways. Every doorway is either for the exclusive passage of the cat (c_1, \dots, c_7) or the mouse (m_1, \dots, m_6). Doorway c_7 is bidirectional and uncontrollable, while others are unidirectional and controllable (i.e., they can be closed). The goal of this process consists of two requirements. First, the cat and the mouse must never occupy the same room simultaneously. Second, after leaving their initial rooms (room 2 for the cat and room 4 for the mouse), they must return within 5 time units. These requirements are specified by formula f_0 in Figure 7 (this formula is also commented on in Example 3.1).




Figure 6: Maze for cat and mouse

Figure 7 details the expansion of the search graph (i.e., automaton M) for this example. Every state of the graph is composed of a world state, which is a pair consisting of the cat and mouse positions, and an MTL formula. For clarity in the figure, we only indicate formulas for some states of the graph. The initial state is $[(2,4), f_0]$. The formulas $f_i, 1 \leq i \leq 5$, are defined in

Figure 7. A proposition $c(i)$ ($m(i)$) means that the cat (mouse) is in room i . The initial formula f_0 is the result of progressing the goal g_1 g_2 through the world state $(2,4)$ with a time difference of 0. States are expanded by applying actions as explained above. For instance, $[(0,4), f_1]$ is produced by applying action $C3$ to obtain the world state $(0,4)$ and by progressing f_0 through $(0,4)$, using a time difference of 1, to get f_1 . Note that the eventually modality in f_1 is bounded by 4, which conveys the fact that one unit of time has elapsed. The expansion of $[(0,4), f_1]$ into $[(1,4), f_2]$ by applying the action c_1 is similar.

The expansion of $[(1,4), f_2]$ into $[(2,4), f_0]$ by applying c_2 deserves more comment. Following lines 15-18 of the formula progression algorithm, the progression of $\Diamond_{\leq 3}(c(2) \wedge m(4))$, which is a conjunct of f_2 , through $(2,4)$ yields $true \vee \Diamond_{\leq 2}(c(2) \wedge m(4))$, which is simplified as $true$. In this way, the progression of f_2 through $(2,4)$ yields f_0 , hence closing a cycle. Note that the infinite execution represented by this cycle satisfies f_0 . In general, as explained below, each cycle generated by our search process is necessarily *satisfactory*.

In Figure 7, solid lines represent transitions of the plan. Dashed lines represent paths that have been explored then rejected when the planner has determined that they lead to sinks. For instance, state $[(3,3), f_5]$ is a sink because it violates the subgoal g_1 of f_0 . This violation causes the progression of f_3 through $(3,3)$ to return *false*, i.e., f_5 . From such trivial sinks, the planner derives less trivial ones. For example, state $[(1,3), f_3]$ also becomes a sink because action C_7 is uncontrollable.

Dotted lines represent parts of the search space that are not explored by our search process due to the use of the dependency-directed backtracking technique explained below. ■

Determining Satisfactory States and Sinks We extend the notion of controllable action to a notion of *controllable state*. A state is controllable if it is expanded (i.e., it is in $CLOSED$) and all its outgoing transitions are due to controllable actions. Each state is attached with a slot, *controllable*, indicating whether it is controllable or not. This slot is set to false for each newly created state. To each state is also attached a *status* that indicates whether the state is *sink*, *satisfactory*, or still *undefined*. The status of each newly created state is set to *undefined*. Then, the search process updates the status of expanded states based on the following definitions.

Definition 4.1 A *sink* is (1) a state of the form $[x, false]$ or (2) a state from which there is an outgoing uncontrollable transition to a sink or (3) a state involving an *eventually* modality and from which there is no enabled action or all enabled actions lead to sinks.

Condition 1 constitutes a basic case for determining sinks. It is implemented by the lines 14-16 of the plan-

action; $s.pct \subseteq 2^{(\mathcal{A} \times (X \times F))}$ contains the set of controllable transitions incoming to s . Hence, if $[a,s] \in s.pct$, then state s can be avoided by disabling a at s , since a is controllable.

The control-directed backtracking mechanism is then implemented by the function *Sink* (see Figure 8). Given a sink s , this function explores, backwards, all ancestors of s reachable through a path formed only of uncontrollable transitions. Each state met in this traversal is marked "sink" and all its descendants are marked "useless" provided that their predecessors are sink or marked "useless" (the function *Mark-Useless* runs over descendants of s , marking them "useless" until we meet a state already marked "useless" or for which at least one predecessor is not marked "useless"). A state marked "useless" is removed from OPEN.

Note, however, that this process might mark states for which all the incoming transitions have not yet been generated. Hence, when these currently missing transitions are generated, the "useless" mark are removed. The overhead incurred by these removals can be avoided by deleting states rather than marking them. Deleted states would then be re-expanded when reached again by further expansions. In this way, we replace the overhead due to updating the "useless" marks by the overhead incurred by re-expanding states. The problem here is comparable to the management of *backpointers* in the search algorithm A^* [Rich and Knight, 1991] (pp. 76-79): one can either keep and update backpointers on expanded nodes, or re-expand them. Both strategies are incomparable in terms of performance.

5 Conclusion

In this paper, we have described a new planning method based on an original *control-directed backtracking* technique and a procedure for checking MTL goal formulas incrementally. Our framework generates plans of reactive systems that take into consideration safety, liveness, and real-time constraints as well as controllable actions in a uniform way. Note that the computation of the feedback function (p does not need to be done after the construction of the entire automaton R . Instead, we could detect satisfactory states at different stages of expansion and immediately update the function $\langle p$ accordingly. This means that the processing done in the function *Check_Non_Tivial_Satisfactory_States* has to be interleaved with the expansion process. In this way, $\langle p$ is produced incrementally. This gives *anytime* capability to our planner [Dean and Body, 1988J.

As a matter of facts, the search process explained in the previous section can be seen as generating the *product* of two *timed automata* (or *timed transition graphs*): one automaton accepting all the world-state sequences, and the other accepting the legal sequences (satisfying the MTL formula). The product automaton is simply a search graph in which world-states correspond to states

in the world-state automaton and MTL formulas corresponds to states in the MTL automaton. Our search process computes this product *on the fly*. Thus, rather than specifying the transitions of the world-state automaton explicitly, we use a STRIPS-like action language to specify them succinctly and then recursively apply actions to generate them only when they become relevant in the search process. Similarly, rather than specifying the transitions of the MTL automaton explicitly, we progress formulas to generate them only when they become relevant in the search process. Recently, Brandin and Wonham have proposed an approach in which the transitions of the two automata are specified explicitly [Brandin and Wonham, 1994]. Our approach has at least two advantages with respect to their work. First, in our planner only world-states and MTL formulas that are relevant in the product are generated. In the Brandin and Wonham's approach, sink states and all potential time transitions must be specified explicitly in the input automata since their relevance can only be determined when constructing the product. Second, MTL specifications are declarative and are maintained as such during our search process. We think that this will facilitate debugging of controller specifications. For example, sink states rise when an MTL *always* modality is made false or the deadline of an *eventually* modality expires. These modalities provide information in a declarative form about undesired states in the process being controlled.

We are investigating different complementary strategies for coping with state explosion. Recent work in AI planning shows that forward-chaining search can be made effective using such strategies [Bacchus and Kabanza, 1995]. At the same time, work in control theory has demonstrated that a forward-chaining exploration can become effective with real-world problems using efficient implementation or modular techniques [Balemi *et al.*, 1993; Brandin, 1994]. In the same trend, we intend to investigate the application of AI planning *goal-regression* techniques ([Rich and Knight, 1991]) to generate search control knowledge that specifies the relevance of actions to a given goal. With this knowledge, forward-chaining will explore the more relevant parts of the search space before the less relevant ones. The usual *goal-regression* techniques in AI planning deal only with goals of achieving state conditions. Our approach to handle general MTL goals is being influenced by two observations. On the one hand, an MTL *eventually* modality expresses a condition to be achieved. Thus, we could parse an MTL formula to extract the sub-formulas expressing conditions to be achieved and then apply *goal-regression* techniques to determine the relevance of actions to these conditions. On the other hand, an MTL *always* modality expresses a condition to be maintained. Each uncontrollable action that could falsify such a condition is relevant since it might lead to sinks. Such relevant actions could also be determined using *goal-regression*. It thus

seems possible to interleave the *goal-regression* processes of computing relevant actions and the *forward-chaining* one of generating accessible states. The challenge will be to study the tradeoffs between these processes and the effectiveness of their combination on real-world problems. In the same line of inquiry, we are also interested in abstraction over large numbers of states during search.

References

- [Alur and Henzinger, 1993] R. Alur and T. Henzinger. Real-time logics: Complexity and expressiveness. *Information and Computation*, 104(1):35-77, 1993.
- [Alur et al., 1993] R. Alur, C. Courcoubetis, and D. Dill. Model-checking in dense real-time. *Information and Computation*, 104(1):2-34, 1993.
- [Bacchus and Kabanza, 1995] F. Bacchus and F. Kabanza. Control strategies in planning. *AAAI Spring Symposium Series. Extending Theories of Action: Formal Theory and Practical Applications*, March 1995.
- [Balemi et al, 1993] S. Balemi, G.J. Hoffman, P. Gyugyi, H. Wong-Toi, and G.F. Frankli. Supervisory control of rapid thermal multiprocessor. *IEEE Transactions on Automatic Control*, 38(7):1040-1059, 1993.
- [Brandin and Wonham, 1994] B. A. Brandin and W. M. Wonham. Supervisory control of timed discrete-event systems. *IEEE Transactions on Automatic Control*, 39(2):329-42, 1994.
- [Brandin, 1994] B. A. Brandin. Supervisory control of an experimental manufacturing cell. In *Proceedings of Thirty-second Annual Allerton Conference on Communication, Control, and Computing*, pages 699-708, Urbana-Champaign, 1994.
- [Dean and Body, 1988] T. Dean and M. Body. An analysis of time-dependent planning. In *Proc. of Sixth National Conference on Artificial Intelligence*, pages 49-54, 1988.
- [Dean and Wellman, 1991] T. L. Dean and M. P. Wellman. *Planning and Control*. Morgan Kaufmann, 1991.
- [Dean et al., 1993] T. Dean, L. P. Kaelbling, J. Kerman, and A. Nicholson. Planning with deadlines in stochastic domains. In *Proc. of Eleventh National Conference on Artificial Intelligence*, pages 574-579, 1993.
- [Drummond and Bresina, 1990] M. Drummond and J. Bresina. Anytime synthetic projection: Maximizing probability of goal satisfaction. In *Proc. of Eighth National Conference on Artificial Intelligence*, pages 138-144, 1990.
- [Emerson, 1990] E. A. Emerson. Temporal and modal logic. In J. van Leeuwen, editor, *Handbook of Theoretical Computer Science*, volume B, pages 995-1072. MIT Press/Elsevier, 1990.
- [Fikes and Nilsson, 1972] R. Fikes and N. J. Nilsson. Learning and executing generalized robots. *Artificial Intelligence*, 3(4):251-288, 1972.
- [Fusaoka et al, 1983] A. Fusaoka, H. Seki, and K. Takahashi. A description and reasoning of plant controllers in temporal logic. In *Proceedings of the IJCAI*, pages 405-408, 1983.
- [Georgeff and Lansky, 1987] M. P. Georgeff and A. L. Lansky. Reactive reasoning and planning. In *Proceedings of Fifth National Conference on Artificial Intelligence*, pages 677-682, 1987.
- [Godefroid and Kabanza, 1991] P. Godefroid and F. Kabanza. An efficient reactive planner for synthesizing reactive plans. In *Proc. of Ninth National Conference on Artificial Intelligence*, pages 640-645, 1991.
- [Kabanza, 1990] F. Kabanza. Synthesis of reactive plans for multi-path environments. In *Proc. of Eighth National Conference on Artificial Intelligence*, pages 164-169, 1990.
- [Makungu et al, 1994] M. Makungu, M. Barbeau, and R. St-Denis. Synthesis of controllers with colored Petri nets. In *Proceedings of Thirty-second Annual Allerton Conference on Communication, Control, and Computing*, pages 709-718, Urbana-Champaign, 1994.
- [Nilsson, 1994] N. J. Nilsson. Teleo-reactive programs for agent control. *Journal of Artificial Intelligence Research*, 1:139-158, 1994.
- [Ramadge and Wonham, 1989] P. J. G. Ramadge and W. M. Wonham. The control of discrete event systems. *Proceedings of the IEEE*, 77(1):81-98, 1989.
- [Rich and Knight, 1991] E. Rich and K. Knight. *Artificial Intelligence*. McGraw Hill, 1991.
- [Stallman and Sussman, 1978] R. M. Stallman and G. J. Sussman. Forward reasoning and dependency-directed backtracking in a system for computer-aided circuit analysis. *Artificial Intelligence*, 9, 1978.
- [Thistle and Wonham, 1986] J. G. Thistle and W. M. Wonham. Control problems in a temporal logic framework. *International Journal on Control*, 44(4):943-976, 1986.