Towards Knowledge-Level Analysis of Motion Planning *

Ronen I. Brafman Jean-Claude Latombe Yoav Shoham
Robotics Laboratory
Stanford University
Stanford, CA 94305
e-mail: {brafman,latombe,shoham}@cs.stanford.edu

Abstract

Inspired by the success of the distributed computing community in applying logics of knowledge and time to reasoning about distributed protocols, we aim for a similarly powerful and high-level abstraction when reasoning about control problems involving uncertainty. Here we concentrate on robot motion planning, with uncertainty in both control and sensing. This problem has already been well studied within the robotics community. Our contributions include the following:

- We define a new, natural problem in this domain: obtaining a sound and complete termination condition, given initial and goal locations.
- We consider a specific class of (simple) motion plans in R^n from the literature, and provide necessary and sufficient conditions for the existence of sound and complete termination conditions for plans in that class.
- We define a high-level language, a logic of time and knowledge, to reason about motion plans in the presence of uncertainty, and use them to provide general conditions for the existence of sound and complete termination conditions for a broader class of motion plans.

Introduction

Much research carried on in computer science in general, and AI in particular, concerns the development of powerful abstractions, and their application to problems of interest. In the context of this article, of particular note is the application of logics of knowledge in distributed computing (e.g., [Halpern and Moses, 1990]). The essential insight behind that line of research was that a formal notion of "knowing," developed initially in philosophy [Hintikka, 1962] and later imported to AI [Moore, 1985], can be coherently and usefully applied to reasoning about (and later also designing) distributed protocols. The reasons for the success of this approach include:

Intuitiveness: The high-level language supported statements of the sort "processor A doesn't know that processor B is faulty," which are precisely the type employed informally by people reasoning about the domain.

Groundedness: The formal notion of knowledge was anything but vague; it was defined precisely in terms of the underlying protocol.

Abstraction and Generality: In principle, the notion of knowledge was dispensable. However, the analysis in terms of knowledge homed in on the essential notion, the knowledge available to the various processors at different points in time, and allowed one to abstract away from the details of how the particular physical protocol implemented that knowledge. This knowledge-level abstraction made it possible to analyze (and later also design) protocols even before their physical implementation was specified; in fact, the same knowledge-level protocol could be implemented differently, without affecting the high-level analysis.

While logics of knowledge have been widely used in AI (e.g., to model human-computer interaction, distributed planning, and nonmonotonic logics), they have so far not been applied in a similar fashion, as a knowledge level corresponding to some specific concrete system. Two exceptions do come to mind – Levesque's knowledge-level analysis of databases [Levesque, 1984], and Rosenschein and Kaelbling's Situated Automata [Rosenschein, 1985, Kaelbling and Rosenschein, 1990]. We claim, however, that there is a much wider arena in which the lessons from distributed computing can be applied, namely planning and control in the presence of uncertainty. In this article we take one step towards exploring this arena, concentrating on robot motion planning.

Robot motion planning with uncertainty is a well researched area [Canny, 1989, Erdmann, 1986, Latombe, 1991, Latombe *et al.*, 1991]. The uncertainty in that domain can arise from several sources, including par-

*The work of the first and third author is supported by DARPA grant USC 598240 (DARPA DABT63-91-C0025) and AFOSR grant #F49620-92-J-0547. The work of the second author is supported by DARPA grant N00014-92-J-1809.

tial information about the location of various objects, sloppy control, and noisy sensing. We argue that this domain exhibits all the 'right' properties: (1) One naturally analyzes the situation by saying that "the robot knows that it is at the goal, since it knows that the current reading could only have been obtained if it were either at the goal or beyond the wall, and it knows its motion plan could not possibly have taken it behind the wall"; (2) the notion of knowledge can be grounded precisely in the motion plan of the robot, as well as some additional parameters such as the slop in control and the noise in sensing.

It would have been convenient to start with a given class of motion planning problems, and delve directly into their knowledge-level analysis. However, we were surprised to find that, although much related research has been conducted in robotics, the simple question *we* would like to pursue has not been addressed. A typical question asked in robotics is "Given that the robot must end up in a particular region, and given bounds on the slop in control and noise in sensing, what is the biggest initial area from which the robot can start, and still be guaranteed to arrive at the goal and recognize that it is at the goal?" This initial area is called the *pre-image* of the goal. In a multi-step motion plan, this question is repeated in a backward-chaining fashion, leading to the method of *pre-image backchaining* [Lozano-Pérez *et al.*, 1984]. In contrast, we consider fixed initial and goal regions and a class of simple motion commands 'Go in direction D, until the termination condition, T, is satisfied'. D can be seen as responsible for reaching the goal, while T is responsible for recognizing it. The seemingly more basic question we ask is: "Given a fixed D, and given bounds on the slop in control and noise in sensing of the robot, does there exist a *good* definition of T?" Of course, one could interpret *good* in many ways. We will interpret it by appealing to standard computer-scientific notions; we will be interested in termination conditions that are *sound* and *complete*, that is, ones that guarantee that if the robot stops, it only stops at the goal, and that it does eventually stop.

Before introducing a high-level language, we will give more feel for the problem by considering it in the context of a particular class of motion plans in R^n; we will present some results on necessary and sufficient conditions for the existence of sound and complete termination conditions in that class. We will then consider general motion planning, define a (fairly standard) logic of knowledge and time for reasoning about them, and then, for a broad class of motion planning problems in R^n, provide a knowledge-level characterization of the conditions in which sound and complete termination conditions exist. Our proofs, contained in a longer version of this paper ([Brafman *et al.*, 1993]), are constructive, that is, when such termination conditions exist, the proofs yield ways of computing them.

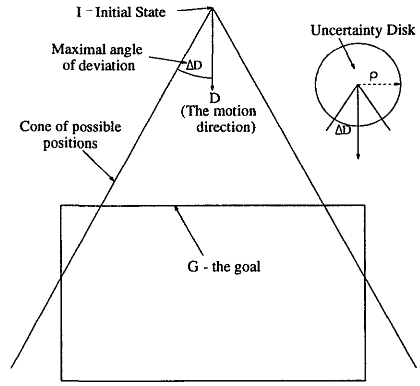

Figure 1: An example domain

We conclude the introduction with a simple path-planning problem in R^1, taken from [Latombe, 1991]. We will use it later to illustrate the various definitions and results, along with more realistic examples, to the extent that space allows.

Example 1 *Assume that our robot is a point moving forward along the positive reals, starting at 0; it moves continuously at finite velocity, until the termination condition is satisfied, at which point it stops. The goal is the interval* $[2,4]$. *There is a position sensing uncertainty of 1, so that if the robot is at location* l, *its sensor may indicate any value between* $l-1$ *and* $l+1$.

In the following let r *denote the current position reading of the robot. Clearly '*$r > 1$*' is a complete termination condition, but not a sound one. Similarly, '*$r = 3$*' is a sound termination condition, but not a complete one (readings need not be continuous: we may have a sequence of readings that are accurate until we reach* 2.5, *at which point they might become consistently off by +1, i.e., start from* 3.5 *and grow). Somewhat surprisingly, there exists a termination condition that is both sound and complete, e.g., '*$r \in [3,5]$*'.*

Termination Conditions in Motion Planning

We introduce a motion planning domain, with particular types of sensing and control uncertainty, in which we investigate the existence of sound and complete termination conditions.

The problem

Our robot starts its motion from a designated set of locations (possibly a singleton), I, within the workspace $\mathcal{W} \subset R^n$. It proceeds along its commanded direction, D, attempting to reach and stop at the goal, G, a compact subset of $\mathcal{W}$ (see Figure 1). However, *control uncertainty* may cause the tangent of the path to deviate by up to ΔD from this direction. This constrains

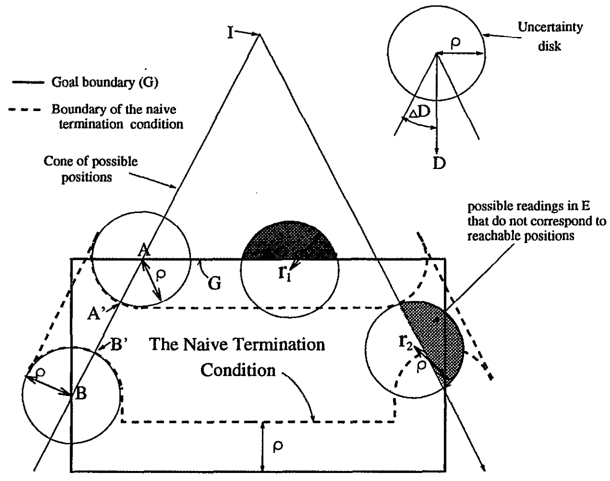

Figure 2: A naive termination condition

the robot to remain within a cone of possible positions defined by the initial state, by D and by ΔD. At each position, q, the robot's sensors supply a position reading, r; but while the robot's motion is continuous, the position readings may not be continuous. The reading, r, must be within a disk of radius ρ, $\Delta R(q)$, centered at q. Consequently, given a position reading, r, the actual position, q, is also within a disk of radius ρ, $\Delta Q(r)$, centered in r, defining the *sensing uncertainty*. If a subset of this disk is outside the cone of possible positions, it can be eliminated as candidates for the actual position. We use the term **motion planning instance** to refer to the domain specified by $\Delta D, \rho, I$, and G. The **termination condition**, T, is a (total) boolean function on the set of possible readings. The first time it evaluates to *true* the robot stops.

An annotated trajectory describes a possible execution of a motion command.

Definition 1 *Given a motion planning instance and a motion command $M=(D,T)$, $\tau=(Q,V)$ is a* **consistent annotated trajectory**, *where Q is a continuous function from $[0,\infty)$ to $\mathcal{W}$ (describing the path), and V is a (not necessarily continuous) function from $[0,\infty)$ to $\mathcal{R}$ (describing the readings), if the the following properties are satisfied:*

- *$Q(0) \in I$;*
- *$\forall s \in [0,\infty)$: $|V(s) - Q(s)| \leq \rho$;*
- *$\forall s \in [0,\infty)$: $|dQ(s) - D| \leq \Delta D$, where $dQ(s)$ is the direction of the tangent to Q at s;*
- *$\forall s \in [0,\infty)$: if $T(V(s))$ is true, then $\forall s' > s$ $Q(s')=Q(s)$.*

Definition 2 *Given a motion planning instance and a motion direction, a termination condition T is* **sound** *if for every consistent trajectory $\tau=(Q,V)$: for $\hat{s} \equiv inf\{s \in [0,\infty)$: $T(V(s))$ is true$\}$ it is the case that $Q(\hat{s}) \in G$. It is* **complete** *if every consistent trajectory $\tau=(Q,V)$ satisfies $\exists s \in [0,\infty)$ such that $T(V(s))$ is true.*

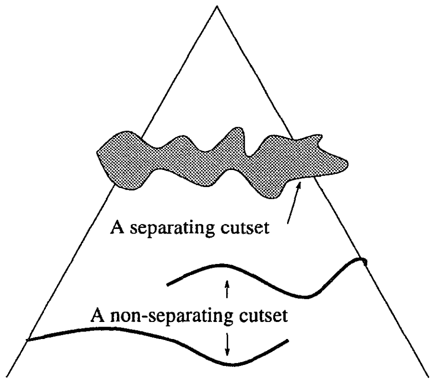

Figure 3: cutsets

That is, if the termination condition is complete the robot eventually stops, and if it is sound, if the robot stops, it stops in the goal. A precise formulation of the problem we wish to investigate, is:

```
Given a motion planning instance and a
motion direction D, does there exist a
sound and complete termination condition?
```

The naive termination condition

Let the **forward projection** of I, $\mathcal{F}(I)$, be the set of positions that can be reached from I by a path consistent with D and ΔD. The termination condition '$\Delta Q(r) \cap \mathcal{F}(I) \subset G$' can only be true at positions within the goal or at positions that cannot be reached from I. Hence it is sound. We call it the **naive termination condition**. The naive termination condition in Example 1 is '$r = 3$'.

Example 2 *In Figure 2 we see an example of a naive termination condition, bounded by the dashed line. The positions from which the reading r_1 can be obtained are contained in the disk $\Delta Q(r_1)$. Some of them, those in the shaded area, are not in the goal. Thus, a reading of r_1 may be obtained from outside the goal, and r_1 is not within the naive termination condition. r_2 is in the naive termination condition, although part of the positions from which we can obtain a reading of r_2 are outside the goal. These positions are outside $\mathcal{F}(I)$, and cannot be reached by any consistent motion.*

The naive termination condition is not complete. In Example 2, because the distance between the points A' and B' is less than 2ρ, a consistent annotated trajectory, in which the termination condition never evaluates to *true* exists.

Conditions for existence of a sound and complete termination condition

We now derive a number of conditions for the existence of a sound and complete termination condition in *domains containing no obstacles, in which the initial states are in a single connected component.*

Definition 3 *A* **motion cutset** *(w.r.t. a motion direction) is a set of positions, at least one of which must be traversed by every consistent trajectory, assuming the termination condition $T \equiv false$.*
A **separating cutset** *is a motion cutset* **cs** *such that $\mathcal{F}(I) \backslash$***cs** *consists of two disjoint connected components, one of which includes the set of possible initial states.*

Figure 3 offers an example of separating and non-separating cutsets.

Theorem 1 *A sufficient condition for the existence of a sound and complete termination condition is that the naive termination condition contains a separating cutset.*

We have a constructive proof of this theorem, which due to length limitations we must omit. The construction relies on the fact that we can disregard positions consistent with our reading if we know that reaching them requires that we pass through earlier positions that satisfy the termination condition. We use the given cutset, cs, to derive a second cutset within the goal, cs', (one that is ρ "farther") and construct a termination condition that guarantees that we never pass cs', thus allowing us to ignore positions that are "beyond" this cutset. Figure 4 illustrates a sound and complete termination condition based on this construction for an example similar to, but more extreme than that of Figure 2 (i.e., A' and B' of Figure 2 are equal).

Example 1 (continued) *The naive termination condition '$r = 3$' is a separating cutset (as is any other point). Using the construction provided by the proof of Theorem 1, a sound and complete termination condition is '$r \in [3, 5]$'. Although a reading in the range of $[3, 5]$ can be obtained outside the goal, the* first *time the termination condition is satisfied will always be within the goal.*

Theorem 2 *A necessary condition for the existence of a sound and complete termination condition is that the naive termination condition be a motion cutset.*

Since for convex goals if the naive termination condition is a motion cutset then it is a separating cutset we obtain:
Consequence 1 *If the goal is convex a necessary and sufficient condition for the existence of a sound termination condition is that the naive termination condition is a motion cutset.*

Using more complex definitions we can prove stronger versions of Theorem 1 (i.e., allow obstacles and relax restrictions on initial states) and Theorem 2 (i.e., relax restriction on initial states).

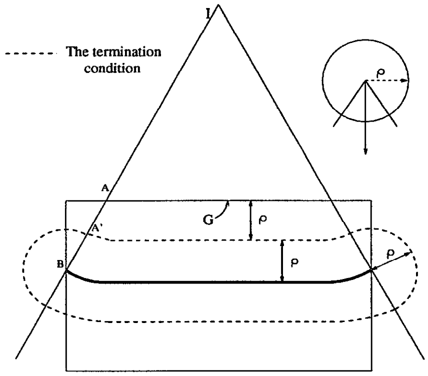

Figure 4: sound & complete termination condition

Knowledge-Level Formalization

We would like to generalize the results of the previous section, by extending our analysis to more general domains, not restricting ourselves to such a limited class of sensors. The notion of knowledge will serve as a powerful abstraction, enabling us to do so.

The motion protocol

We can view our endeavor as the investigation of a class of two-player protocols. Our players are the robot, which follows the motion command, and the environment, which decides nondeterministically, how the robot actually moves and what it senses, within given margins [Taylor *et al.*, 1988]. Let us make this more precise.

Definition 4 *A* **(global) state** *is a pair (e,r), where $e \in \mathcal{E}$ and $r \in \mathcal{R}$. $\mathcal{E}$ is the set of possible (local) states of the environment and $\mathcal{R}$ is the set of possible (local) states of the robot. The set of global states is denoted by $\mathcal{S}$. If a is an agent (i.e., the robot or the environment), the set of local states of a is denoted by $S_{|a}$.*

In the domain of the previous section the local state of the robot consisted of its current sensor reading, while the local state of the environment consisted of the robot's *actual* position.

Definition 5 *A* **run** *R, is a function from $[0, \infty)$ to $\mathcal{S}$. A* **system** *is a set of runs. A run R is* **consistent** *with a system S if $R \in S$.*

A run extends the static description of the world, provided by the state, over time. A system corresponds to some subset of the set of possible runs, those runs that describe behaviors that correspond to the ones we would like to model. The states that are part of a consistent run are called **possible states** (or worlds), and represent states of the world that may hold at some time. To reason about one-step motion plans we look at a parameterized class of systems that includes runs that obey the following restrictions:

$\underline{\mathcal{P}(I, D, \Delta D, \Delta R)}$

The initial state is in I.
The restriction of the run to the local state of the robot describes a continuous trajectory.
The direction of the tangent to the trajectory at each point q, dq, satisfies $dq \in \Delta D(q)$.
The reading r in q satisfies $r \in \Delta R(q)$.

Note that there are no restrictions on the shape or size of $\Delta R(q)$, the set of readings possible at q. The possible initial states are not constrained, either, nor is $\Delta D(q)$ (thus we are not restricted to cone-like domains).

As an example, in the previous section we looked at a case where $\Delta R(q)$ is a disk of radius ρ centered in q, and $\Delta D(q) = \{D' : |D - D'| \leq \Delta\}$ for a constant Δ.

Knowledge-level analysis of the motion protocol

We introduce a language for reasoning about the motion protocol. This language contains temporal and epistemic modal operators and has intuitive semantics. We use it to formulate and prove new, more general, results.

The language We assume that we have a propositional language $\mathcal{L}$. Given an instance of the above motion protocol, let $\mathcal{S}$ be the set of possible states. Let $\mathcal{I}$ be an *interpretation function*, which is a boolean function on $\mathcal{S}$ and the primitive propositions in $\mathcal{L}$. $\mathcal{I}$ tells us whether a certain primitive proposition is true in a certain possible state. We define the notion of satisfiability of a propositional formula in a state s, under interpretation $\mathcal{I}$ of $\mathcal{S}$.

- $\mathcal{I},\mathcal{S},s \models p$ for a primitive proposition p if $\mathcal{I}(s,p) = true$;
- $\mathcal{I},\mathcal{S},s \models \neg\alpha \Leftrightarrow \mathcal{I},\mathcal{S},s \not\models \alpha$;
- $\mathcal{I},\mathcal{S},s \models \alpha \wedge \beta \Leftrightarrow \mathcal{I},\mathcal{S},s \models \alpha$ and $\mathcal{I},\mathcal{S},s \models \beta$.

We assume that the interpretation function $\mathcal{I}$, is fixed and 'natural', e.g., the proposition g, denoting a goal state, will be satisfied *exactly* by those states in which the position is part of the goal. We will use $\mathcal{S} \models \alpha$ when α is satisfied by all states in $\mathcal{S}$, and $\models \alpha$ when $\mathcal{S}' \models \alpha$ for any $\mathcal{S}'$.

Motion is closely connected with time, thus we add temporal operators:[1]

- $\mathcal{S}, s \models \exists_\Box \alpha$ if there exists a run R such that $\exists t \in [0, \infty)$ such that $R(t) = s$ and $\forall t' \geq t$ $R(t') \models \alpha$. I.e., all possible states following s (s inclusive) of a certain run containing s, satisfy α.
- $\mathcal{S}, s \models \exists_\Diamond \alpha$ if for some run R such that $\exists t \in [0, \infty)$ $R(t) = s$, $\exists t' \geq t$ s.t. $R(t') \models \alpha$. I.e., some possible state following s (or s itself) in some run containing s, satisfies α.

[1]Temporal logic with these operators were investigated by Emerson and Halpern ([Emerson and Halpern, 1985]), among others.

These modal operators define two other operators:

$$\forall_\Box \alpha \equiv \neg\exists_\Diamond \neg\alpha \qquad \forall_\Diamond \alpha \equiv \neg\exists_\Box \neg\alpha$$

In the above operators the present state is considered a part of the future. We have similar operators for the past. The 8 operator for all tenses are:

Future	Past
$\forall_\Diamond, \forall_\Box, \exists_\Diamond, \exists_\Box$	$\forall_\diamondminus, \forall_\boxminus, \exists_\diamondminus, \exists_\boxminus$

To deal with uncertainty we define the notion of knowledge.

Definition 6 *Let $\mathcal{S}$ be a given set of states, and let $s, s' \in \mathcal{S}$.*
$s \sim_r s' \Leftrightarrow$ the state of the robot is identical in s and s'.

Note that $\sim_r$ is an equivalence relation.

Definition 7 *Let $\alpha \in \mathcal{L}$. The robot* **knows** *α at $s \in \mathcal{S}$, written $\mathcal{S},s \models K_r\alpha$ if for any state s', such that $s \sim_r s'$, $\mathcal{S},s' \models \alpha$.*

From now on we shall assume that our language, $\mathcal{L}$, is closed under the temporal and epistemic operators. All definitions remain unchanged. Although detailed logical investigation is not the thrust of this paper, we note in passing that K is an $S5$ operator. Note that given $\mathcal{S}$, the satisfiability in s of a formula of the form $K_r\alpha$ depends only on the local state of the robot, and can be interpreted as a predicate on its local state. $K_r\alpha$ thus uniquely defines a termination condition (although not necessarily an easily computable one).

The following is our main theorem:

Theorem 3 *Assume that $\mathcal{S} \models (\neg g \wedge \forall_\diamondminus g) \rightarrow \exists_\Box \neg g$, where g is the proposition satisfied precisely by the goal states. A sound and complete termination condition exists iff $K_r\forall_\diamondminus g$ defines a sound and complete termination condition.*[2]

<u>Discussion</u>: There are a number of important things to note.

Constructive canonical form: The theorem gives a constructive definition of a sound and complete termination condition, if one exists, so we need only check this condition to verify the existence of a sound and complete termination condition.

Optimality: For any run, this termination condition evaluates to *true* no later than any other sound and complete termination condition.

Generality: The use of knowledge to characterize the termination condition means that it applies to sensing uncertainty of *any* type, i.e., while previously $\Delta R(q)$ was a disk, it can now take on any shape. In fact, we

[2]We will assume that $K_r\forall_\diamondminus g$ is a closed set. This is a very weak assumption which we discuss in the longer version of this paper [Brafman *et al.*, 1993].

are *not* constrained to position sensing, and the theorem applies to force sensing, as well as robots with sensing memory. The initial state may be part of a large region or just a point and our domain may contain obstacles. The control uncertainty is not limited to the cone-like behavior of the previous section.

The following lemma shows that Theorem 3 covers a large natural family of domains,

Lemma 1 *For any convex subspace of R^n, if the goal region is convex, $\mathcal{S} \models (\neg g \wedge \forall_{\diamond} g) \rightarrow \exists_{\square} \neg g$.*

Example 1 (continued) *If you recall our robot moved along the positive reals, its goal was to be in $[2,4]$ and its reading uncertainty was $\rho = 1$. The condition $\forall_{\diamond} g$ is satisfied by positions $q \in [2, \infty)$, thus $K_r \forall_{\diamond} g$ is satisfied by readings $r \in [3, \infty)$, corresponding to a sound and complete termination condition.*

Using the above theorem we have been able to prove the following result about the domain of the previous section.

Theorem 4 *In an empty subspace of R^n with $\Delta R(q)$ a disc of radius $\rho \in R$ and a compact goal, in which the condition $\mathcal{S} \models (\neg g \wedge \forall_{\diamond} g) \rightarrow \exists_{\square} \neg g$ of Theorem 3 holds for a robot with sensing history; a sound and complete termination condition based on a complete reading history exists, iff there is one dependent only on the* current *position reading.*

Conclusion and future work

Knowledge is a powerful tool for reasoning about domains in which uncertainty exists. The temporal-epistemic language we used provides a natural and powerful tool in the domain of motion planning with uncertainty, and enabled us to express and prove results more general than when using geometric specifications. One important task for future research will be to look for interesting temporal/epistemic properties of different sensors and domains (relating the knowledge level and the geometric level), and exploit these properties to prove more specific results. We also hope to be able to lift the restrictions of Theorem 3, and find a general characterization of sound and complete termination conditions.

The present paper has studied one particular problem. However, knowledge can be applied to many natural problems in motion planning, especially ones that deal with multiple agents, where purely geometrical reasoning would become even more complicated. In fact, most interesting aspects of knowledge come out when there are a number of agents. For example, one can look at the problem of *mobile coordinated attack*, in which two robots need to halt at their respective goals at synchronized times. Here, a more complicated notion of *common knowledge* is involved (see [Fagin *et al.*, 1993] for definitions and examples). Various geometric settings of this problem can be explored and it seems that in these complex environments knowledge will be an essential tool. In fact, motion planning offers all the problems encountered in distributed systems and more, but in a much richer setting.

Acknowledgement We wish to thank Ronny Kohavi and Moshe Tennenholtz for important discussions about this work.

References

Brafman, R. I.; Latombe, J. C.; and Shoham, Y. 1993. Towards knowledge-level analysis of motion planning. Technical report, Stanford University. in preparation.

Canny, J. F. 1989. On computability of fine motion plans. In *Proc. of the 1989 IEEE Int. Conf. on Robotics and Automation.* 177–182.

Emerson, E. A. and Halpern, J. Y. 1985. Decision procedures and expressivenessin the temporal logic of branching time. *J. of Comp. and Sys. Sci.* 30(1):1–24.

Erdmann, M. 1986. Using backprojection for fine motion planning with uncertainty. *Int. J. of Robotics Research* 5(1):19–45.

Fagin, R.; Halpern, J. Y.; Moses, Y.; and Vardi, M. Y. 1993. *Reasoning about Knowledge.* MIT Press. to appear.

Halpern, J. Y. and Moses, Y. 1990. Knowledge and common knowledge in a distributed environment. *J. ACM* 37(3):549–587.

Hintikka, J. 1962. *Knowledge and Belief.* Cornell University Press, Ithaca, NY.

Kaelbling, L. P. and Rosenschein, S. J. 1990. Action and planning in embedded agents. *Robotics and Autonomous Systems* 6:35–48.

Latombe, J. C.; Lazanas, A.; and Shekhar, S. 1991. Robot motion planning with uncertainty in control and sensing. *Artificial Intelligence* 52:1–47.

Latombe, J. C. 1991. *Robot Motion Planning.* Kluwer Academinc Publishers, Boston.

Levesque, H. 1984. Foundations of a functional approach to knowledge representation. *Artificial Intelligence* 23(2):155–212.

Lozano-Pérez, T.; Mason, M. T.; and Taylor, R. H. 1984. Automatic synthesis of fine-motion strategies for robots. *Int. J. of Robotics Research* 3(1):3–24.

Moore, R. C. 1985. A formal theory of knowledge and action. In Hobbs, J. R. and Moore, R. C., editors 1985, *Formal Theories of the Common Sense World*, Norwood, N.J. Ablex Publishing Corporation.

Rosenschein, S. J. 1985. Formal theories of knowledge in ai and robotics. *New Generation Comp.* 3:345–357.

Taylor, R. H.; Mason, M. T.; and Goldberg, K. Y. 1988. Sensor-based manipulation planning as a game with nature. In Bolles, R. and Roth, B., editors 1988, *Robotics Research 4.* MIT Press. 421–429.